



Sun StorEdge™ Instant Image 2.0.1 Installation Guide

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

Part No. 806-5317-11
February 2001, Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 U.S.A. All rights reserved.

This product or document is 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, Sun StorEdge, Sun Enterprise, AnswerBook2, docs.sun.com, and Solaris 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 Etats-Unis. Tous droits réservés.

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

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

Sun, Sun Microsystems, le logo Sun, Sun StorEdge, Sun Enterprise, AnswerBook2, docs.sun.com, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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.

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



Contents

Related Documentation	vii
Overview	9
Configuration Requirements	10
Installation Overview	11
Procedural Details	12
▼ To Prepare the Target Host	12
▼ To Install and Configure Instant Image	14
▼ To Configure STE on the Target Server	15
Configuration Considerations	20
▼ To Prepare the Initiator	24
▼ To Install the Fibre Channel Host Adapter	25
▼ To Configure STE Devices on the Initiator	25
▼ To Troubleshoot STE	26
▼ To Install the STE Software Packages Automatically	4
▼ To Install the STE Software Packages Manually	5
Post-Installation Patches	6
▼ Complete the Post-Installation	6
▼ To Disable STE	7
▼ To Reconfigure STE	7

- ▼ To Enable STE 8
- ▼ To Remove STE 8
- ▼ To Upgrade or Reinstall STE 10
- ▼ To Upgrade the Operating System 10

Preface

The *Sun StorEdge Instant Image 2.0.1 Installation Guide* describes the installation and configuration procedures for the use of Sun StorEdge™ Instant Image software with Sun StorEdge Target Emulation (STE).

How This Book Is Organized

Chapter 1 describes the installation and configuration procedures for Instant Image 2.0.1.

Appendix A provides detailed information about the use of STE software.

The Glossary contains definitions of terms used in this document.

Using UNIX Commands

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

Refer to the software documentation that you received with your system.

Documentation Conventions

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

Related Documentation

Application	Title	Part Number
man Pages	ii_control(1M)	N/A
	ii_config(1M)	
	ii_health(1M)	
	esm_cli(1M)	
	steconf(1M)	
	steadm(1M)	
	ste.cf(4M)	
	sftm(7D)	
	ste(7D)	
	socal(7D)	
	svadm(1M)	
Release Notes	<i>Sun StorEdge Instant Image 2.0 Release Notes</i>	806-0231
	<i>Sun StorEdge Target Emulation 1.2 Release Notes</i>	806-5527
	<i>Sun StorEdge Instant Image 2.0.1 Release Notes</i>	806-5603
Installation and User	<i>Sun StorEdge Instant Image 2.0 Installation Guide</i>	806-4004
	<i>Sun StorEdge Instant Image 2.0 System Administrator's Guide</i>	806-0230
	<i>Sun StorEdge Instant Image 2.0.1 Configuration Guide</i>	806-5318

Accessing Sun Documentation Online

Browse the <http://www.sun.com/products-n-solutions/hardware/docs/index.html> .com archive or search for a specific book title or subject.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

`docfeedback@sun.com`

Please include the part number of your document in the subject line of your email.

Instant Image 2.0.1 Installation and Configuration

Overview

The *Sun StorEdge Instant Image 2.0.1 Installation Guide* describes the installation and configuration procedures for the use of Sun StorEdge™ Instant Image software with Sun StorEdge Target Emulation (STE).

Instant Image 2.0.1 provides the capability to export the shadow volume of a master volume to an application server. The operation of Instant Image, with the shadow exported to another host through STE, is the same as if you were accessing the master and shadow volumes locally. You can write to either the master or shadow and you can resynchronize in either direction.

Note – Do not build a new filesystem or database on Instant Image 2.0.1 shadow volumes exported through STE.

Instant Image 2.0.1 is not intended to be used to provide additional storage. Do not reconfigure or reformat the shadow volume. If you build a new filesystem or database on the shadow volume, you will create a new infrastructure based on a size larger than the master. This is not a supported use of this product, which could lead to undesirable results. Refer to the “Configuration Considerations” section of this manual.

For late-breaking news about this release, go to the following web site and select StorEdge Instant Image Software Product:

<http://www.sun.com/storage/software>

If you are a Sun™ support or service provider, go to the following web site:

<http://webhome.ebay/networkstorage/products/>

This chapter contains the following main topics:

- Configuration Requirements
- Installation Overview
- General Procedures
- Procedural Details

Configuration Requirements

The hardware and software configurations you can use are defined in the following tables. Instant Image 2.0.1 runs on the Solaris 2.6, Solaris 7 or Solaris 8 operating environments.

Note – Instant Image 2.0.1 is not supported in a Sun Cluster environment.

Note – Instant Image 2.0.1 supports only the target servers indicated below. PCI-based target servers are not supported.

TABLE 1-1 Target Configuration

Target Server	Fibre Channel Host Adapter	Sun StorEdge System
Sun Enterprise™ server models ES3500, ES4x00 to ES6x00, ES10000	SBus - P/N 501-3060	Any storage system supported by Sun

TABLE 1-2 Initiator Configuration

Qualified Initiators	Fibre Channel Host Adapter
Sun Enterprise PCI based servers	PCI - P/N 375-0040 (QLogic 2100)
Sun Enterprise SBus servers	SBus - P/N 501-3060
Microsoft NT 4.0	QLogic 2200

Installation Overview

The general procedure used to install and configure Instant Image with STE is as follows. The detailed steps are given in the sections that follow.

1. Verify qualified configuration as described in TABLE 1-2. Verify that the servers are at the proper revision for all firmware and driver patches.
2. Install the SOCAL driver patch.
3. Install the Fibre Channel Host Adapter.
4. Install STE 1.2.
If a previous version of STE is installed, uninstall the previous version, then, install STE 1.2.
5. Install Instant Image 2.0 (using the Instant Image 2.0 CD), in accordance with the Instant Image 2.0 Installation Guide.
If a 1.x version of Instant Image is installed, uninstall the previous version then install Instant Image 2.0.
6. Put “Master” Volume under SV control by editing the (`/etc/opt/SUNWspsv/sv.cf`) file and running `/usr/opt/SUNWesm/SUNWspsv/sbin/svadm -r`.
7. Create a point-in-time copy (“Shadow”) of the data using the `/usr/opt/SUNWesm/bin/esm_cli` or the `/usr/opt/SUNWesm/bin/esm_gui`. Follow the procedures in the *Sun StorEdge Instant Image 2.0 System Administrator's Guide*.
8. Use the `/usr/opt/SUNWesm/sbin/steconf` utility to get the port device names that will be referenced in the `/etc/opt/SUNWte/ste.cf` file.
9. Add the Target Mode Fibre Channel Ports, Virtual Disks, and Phantom Devices to the STE configuration file (`/etc/opt/SUNWte/ste.cf`).
10. Run the `/usr/opt/SUNWesm/SUNWte/tools/stesocal.ksh` script to create the `/kernel/drv/socal.conf` file.
11. Reboot the server.
12. The target host exports the STE LUNs to the initiator.
13. On the initiator, label the LUNs presented by the target host using `format (1M)`.

Procedural Details

This section describes in detail the procedures outlined in the previous section. Before getting started, make sure you are logged in as the root user.

▼ To Prepare the Target Host

1. Install the appropriate SOCAL driver patch.

You must install the appropriate SOCAL driver patch as indicated below. The patch number indicates the minimum patch revision required. Contact Sun Service or refer to <http://www.sunsolve.sun.com> to get the latest patch information.

TABLE 1-3 Minimum SOCAL Driver Patches Required

Patch Number	Description
105375-21	Solaris 2.6 SOCAL driver patch
107469-06	Solaris 7 SOCAL driver patch
109460-01	Solaris 8 SOCAL driver patch

2. Ensure you have the appropriate version of Java.

Java version 1.2.1_04b or compatible versions is required. To update your version of java, go to <http://java.sun.com>. To verify your version of java:

```
# java -version
```

If you are loading Instant Image for the French, Japanese, or Chinese locales, load the appropriate localized version of Java.

3. Install the appropriate Solaris environment patches.

Note – There are several patches required to use Instant Image. It is recommended that the most recent revision of these patches be loaded before the Instant Image packages are installed.

TABLE 1-4 Solaris Environment Patches for Instant Image

Solaris Environment Version	Patch	Description
2.6	105181-23	Kernel Update patch is to be installed before loading data services software. This patch is included with the Solaris Recommended Patch Cluster.
	108091-03	Sun OS 5.6: ss JDK 1.2.1_04b
7	106541-14	Kernel Jumbo patch is to be installed before loading data services software. This patch is included in the Solaris Recommended Patch Cluster.
	106924-06	5.7 isp driver
	106980-15	5.7 libthread
	107081-25	5.7 Motif 1.2.7 and 2.1.1 runtime library
	108376-18	OpenWindows 3.6.1 Xsun
	107636-05	X input and output metod patcg
8	108528-05	Kernel Update patch is to be installed before loading data services software. This patch is included with the Solaris Recommended Patch Cluster.

4. If you are running Solaris 2.6, add to the `/etc/system` file:

```
set kobj_map_space_len = 0x200000
```

5. If you are using the Sun StorEdge A5x00 array or T300, or VERITAS Volume Manager™ software, or Solstice DiskSuite™ software, verify that the target host is at the proper revision for all firmware and driver patches.

Note – Refer to <http://www.sunsolve.sun.com> to get the latest patch information.

▼ To Install and Configure Instant Image

In the example that follows, two VERITAS volumes are used. It is assumed that these volumes already exist. One volume (`/dev/vx/rdisk/datadg/master01`) will be used as Instant Image "master" volume. Another volume (`/dev/vx/rdisk/datadg/shadow01`) will be the Instant Image "shadow" volume, and will be exported from the target to the initiator host.

1. On the STE target host, Install Instant Image 2.0.

Install the Instant Image 2.0 software in accordance with the *Sun StorEdge Instant Image 2.0 Installation Guide*. If Instant Image 1.x is installed, uninstall Instant Image 1.x, then install Instant Image 2.0.

2. Create the bitmaps.

The location of the bitmaps affects performance. The Instant Image GUI (`esm_gui`) will create the bitmaps in a default location in the root filesystem. The location of the bitmaps is configurable. Refer to the *Sun StorEdge Instant Image 2.0 System Administrator's Guide* for additional information.

For optimal performance, the bitmaps should be created as a volume and cached with Fast Write Cache. The Instant Image bitmap file should be located in a file system or VERITAS volume that is mounted at boot and should be mirrored, if possible. This file system should be used exclusively for Instant Image bitmap files. In the `/etc/vfstab` file, insert a line similar to the following.

If the mount point is `/iibitmaps`, for example, you can then specify the bitmaps from the GUI or CLI.

```
/dev/vx/dsk/datadg/bitmapvol01 /dev/vx/rdisk/datadg/bitmapvol01  
/iibitmaps ufs 1 yes -
```

The bitmap files should be sized per the "Volume Size Requirements" of the *Sun StorEdge Instant Image 2.0 System Administrator's Guide*. In this example, a 32k bitmap file is created.

```
# /usr/sbin/mkfile 32k /iibitmaps/bitmap1
```

3. Put the master volumes under the control of the Storage Volume driver.

The standard Instant Image configuration puts both the master and shadow of an Instant Image pair volumes under SV control. In this configuration, only the master volumes are under SV control. The shadow volumes will be under STE control. Edit the `/etc/opt/SUNWspv/sv.cf` file and add the master volume name to the file followed by the appropriate keyword, raw or cache.

In this example, Fast Write Cache is installed. A sample entry in the `/etc/opt/SUNWspsv/sv.cf` would be:

```
/dev/vx/rdisk/datadg/master01 cache
```

Use `svadm(1m)` to reconfigure the SV driver to control the master volume:

```
# /usr/opt/SUNWesm/sbin/svadm -r
svadm: enabling new sv: /dev/vx/rdisk/datadg/master01 (cache)
```

For more information on configuring Instant Image volumes for the SV driver, refer to the *Sun StorEdge Instant Image 2.0 System Administrator's Guide*.

Considerations:

- **Protecting disk-related data.** Care must be taken not to over-write disk-related data on master or shadow volumes. See “Protecting Disk-Related Data” in Appendix A for detailed information.
- **Shadowing / (root), /usr, /opt, /var or /export/home file systems.** Do not use Instant Image to shadow /, /usr /opt, /var or /export/home file systems or any file system mounted at boot on the target server because software and drivers exist in these partitions.
- **All volumes on target host drives.** All Instant Image volumes (master, shadow, and bitmap) must be on disks under the direct control of the target host.

▼ To Configure STE on the Target Server

Configure the SOCAL driver to support target mode ports by creating a `/kernel/drv/socal.conf` file. A brief overview of the `ste.cf` file and `socal.conf` file is provided. For more detailed information, refer to Appendix A, the `ste.cf(4M)` and `steconf(1M)` manual page.

1. To list the host bus adapters (HBA), run the `/usr/opt/SUNWesm/sbin/steconf` script.

The utility will display information similar to the following:

```
# /usr/opt/SUNWesm/sbin/steconf
Brd Slot Driver Device Name
-----
 1      0  sftm  /devices/sbus@3,0/SUNW,socal@0,0:1
          sftm  /devices/sbus@3,0/SUNW,socal@0,0:0
 1      1  sftm  /devices/sbus@2,0/SUNW,socal@1,0:0 (Initiator)
          sftm  /devices/sbus@2,0/SUNW,socal@1,0:1 (Offline)
 1      2  sftm  /devices/sbus@2,0/SUNW,socal@2,0:0 (Offline)
          sftm  /devices/sbus@2,0/SUNW,socal@2,0:1 (Offline)
 1     13  sftm  /devices/sbus@2,0/SUNW,socal@d,10000:1 (Offline)
          sftm  /devices/sbus@2,0/SUNW,socal@d,10000:0 (Offline)
```

The entries above with “10000” denote internal SOC adapters that cannot be used by STE software. Adapters labeled as “Initiator” cannot be used for STE software. “Offline” indicates that the Fibre Channel loop is down and can be configured to use STE (if it is not a 10000). Ports already configured in target mode are blank and may already be in use by STE. Check the `ste.cf` file to verify if the port is already configured.

2. Edit the target host's `/etc/opt/SUNWte/ste.cf` file to configure the Target Mode Fibre Channel Ports.

Before you use STE, you must add information about Target Mode Fibre Channel Ports, Virtual Disks, and any Phantom Partitions to the STE configuration file (`ste.cf`).

This step specifies the HBA that will be used for STE software. Add the HBA entry to the Target Mode Fibre Channel Ports section of the `ste.cf` file. See “Target Mode Fibre Channel Port Section of `ste.cf` File” of Appendix A for a description of the fields in this section of the `ste.cf` file. Add the HBA entry to the `ste.cf` file for an SBus slot that does not indicate that it is an initiator. For example:

```
# StorEdge Target Emulation Configuration File - ste.cf
# Target Mode Fibre Channel Port section

#Port
#Name  Driver  Device Name                                Loop Id
#-----
tm0     sftm    /devices/sbus@2,0/SUNW,socal@2,0:0         7
tml     sftm    /devices/sbus@2,0/SUNW,socal@2,0:1         7
```


Considerations:

- **Port name.** The port name is `tm` followed by a unique integer value to define the target mode port being configured.
 - **Driver.** This is the driver as identified by the `steconf` command. The entry is “sftm” for SBus systems.
 - **Device name.** This is the device name from the `steconf` command. The recommended practice is to simply copy and paste the device name from the `steconf` output. The final section of the device name provides the port and the location. An entry of 0:0 indicates port 0, right side while an entry of 0:1 indicates port 0, left side.
 - **Loop ID.** This identifies the target loop ID and can be any number from 0 to 15.
3. Run the `/usr/opt/SUNWesm/SUNWte/tools/stesocal.ksh` script to create the server's `/kernel/drv/socal.conf` file. The `socal.conf` file identifies the SOC+ ports that are available for use by STE software to the `SOCAL(7D)` driver. This script uses the `ste.cf` file (step 2) to create the `socal.conf` file.

See “About the `socal.conf` File” in Appendix A for information about this file. In the following example, ports 0 and 1 of the SOC+ device at `sbus@2, 0/SUNW`, `socal@2,0` (board 1, slot 2) are reserved for target mode.

```
name="SUNW,socal" parent = "/sbus@2,0"
reg=0x2, 0x0, 0x10000, 0x0, 0x10000, 0x10000, 0x0, 0x20000, 0x18
port0-loop-id=7 port1-loop-id=7;
```

4. Continue to edit the target host's `/etc/opt/SUNWte/ste.cf` file to configure the virtual disks for STE software that will be exported to the directly attached initiator host.

See Appendix A for a description of the fields in this section of the `ste.cf` file.

Note – Do not put STE volumes in the `sv.cf` file. All volumes and disks can be listed in *either* the `ste.cf` file *or* the `sv.cf` file, but not in both.

For example:

```
# StorEdge Target Emulation Configuration File - ste.cf
#
# Target Mode Fibre Channel Port section
#Port
#Name    Driver    Device Name                                Loop Id
#-----
tm0      sftm      /devices/sbus@2,0/SUNW,socal@2,0:0        7
#
# Virtual Disk section
#
#Virtual disks
#Name    Partition Name                        Port Name SCSI LUN State  Options
#-----
vdisk0   /dev/vx/rdisk/datadg/lun0              tm0      0      online ro
vdisk1   /dev/vx/rdisk/datadg/shadow01          tm0      1      online ph=dev0 pt=dev1
```

Considerations:

- **Multiple virtual disks.** You can have multiple virtual disks through a single Fibre Channel port by providing each virtual disk with its own unique SCSI LUN number, starting at LUN 0.
- **Vdisk name.** Consists of the keyword “vdisk” followed by a unique integer to define which vdisk is being configured.
- **Partition name.** Any volume to be exported as a LUN.
- **Port name.** The port name will match a vdisk to the correct target mode port by specifying an appropriate `tm` name.
- **SCSI LUN number.** Each virtual disk is assigned its own LUN number. This LUN number is the disk number as seen on the remote initiator. Each target mode port must have a LUN0. LUN0 must never be disabled. You can use numbers from 0 to 225 for up to 256Vdisks per target. The maximum number of LUNs per target is 32. However, a PCI initiator has a physical limitation of 16 LUNs per target.

5. Edit the target's `/etc/opt/SUNWte/ste.cf` file to configure phantom headers

Note – Do not include headers and tails when you configure NT LUNs in the `ste.cf` file.

Certain open systems application servers set aside space at the beginning and/or end of each disk for header information. The required size and location (at the beginning or end of the disk are system-specific.

To prevent these initiators from writing their header information over the contents of a shared device, you must prepend a phantom header and append a phantom tail to the initiator's view of the contents of the disk and store the initiator's header information there. You must complete this section of the STE configuration file.

The phantom headers and tails physically reside on a separate volume that contains the phantom devices. The header is the first cylinder of the LUN as viewed by the initiator. To protect against data loss, the phantom partition should be a mirrored or RAID 5 partition.

See “Phantom Header Section of `ste.cf` File” in Appendix A for a description of the fields in this section of the `ste.cf` file. The following provides an example of the `ste.cf` file with a phantom header and tail. Refer to Appendix A and the following section, “Configuration Considerations”, for guidance in specifying phantom headers and tails for your configuration.

```
# StorEdge Target Emulation Configuration File - ste.cf
#
# Target Mode Fibre Channel Port section
#Port
#Name      Driver  Device Name                                Loop Id
#-----
tm0        sftm    /devices/sbus@2,0/SUNW,socal@2,0:0         7
#
# Virtual Disk section
#
#Virtual disks
#Name      Partition Name                    Port Name SCSI LUN State  Options
#-----
vdisk0     /dev/vx/rdisk/datadg/lun0                tm0      0      online ro
vdisk1     /dev/vx/rdisk/datadg/shadow01            tm0      1      online ph=dev0 pt=dev1

# Phantom Header section

#device
#keyword   Partition Name                    Start Block  Size
#-----
dev0       /dev/vx/rdsk/datadg/headsandtails0        0           2268
dev1       /dev/vx/rdsk/datadg/headsandtails0    2268        4536
```

6. Reboot the target server.

A reboot is required after creating the `socal.conf` file. When the server is rebooted, STE will automatically be enabled.

7. Check the STE configuration using the `steadm -c` command.

The output that you see will be similar to this:

```
/usr/opt/SUNWesm/sbin/steadm -c

STE Version DS1_02 5.8 03.23.2000 - 1
STE Node WWN 5020020000013d51

tm0: SOC+ FC-AL Port 0
tm0: Loop ID 7 WWN 2002020000013d51

vdisk0 : tm0 LUN 0 RO
Vendor: SUN      Product: STE20480      Rev: 0000
Serial: 51006F7B
Main: /dev/vx/rdisk/testdg/lun0 with 20480 (512 byte) Blocks

vdisk1 : tm0 LUN 1
Vendor: SUN      Product: STE416404      Rev: 0000
Serial: 51008EA2
Main: /dev/vx/rdisk/datadg/shadow01 with 409600 (512 byte)
Blocks
Header: /dev/vx/rdisk/datadg/headsandtails0 at Offset 0 with 2268
(512 byte) Blocks
Tail: /dev/vx/rdisk/datadg/headsandtails0 at Offset 2268 with
4536 (512 byte) Blocks
```

Configuration Considerations

Since STE presents an interface that makes it look like a SCSI target, STE is required to export geometry attributes that a SCSI direct access device reports. This geometry is comprised of three fields, the number of tracks per zone (number of heads), the number of sectors per track, and the number of cylinders. As STE exports logical volumes that have no geometry attributes, it must contrive one. STE exports a geometry of 108 sectors/tracks, 21 tracks/cylinder, and n cylinders, where $n = \text{capacity of the logical volume} / (21 * 108)$.

The following figure shows the architecture on application server A (the target). The Instant Image shadow volume is exported through STE to application server B (the initiator).

When the only path to this storage is through STE, the unused area (at most 2267 blocks, where 1 block = 512 bytes) produced by the rounding error in the above computation is inconsequential. The operating system formats this storage, never sees the truncated storage, and never uses it. Currently Solaris bounds its partitions on a cylinder granularity, so LUNs exported through STE will be partitioned by

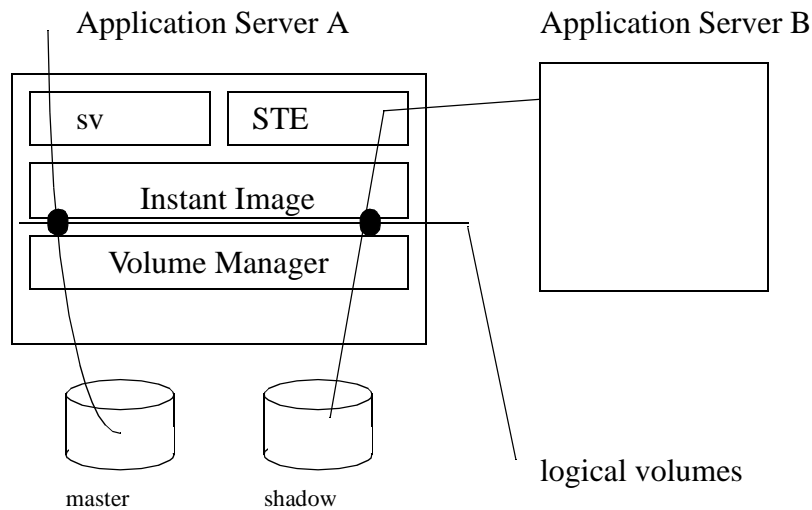


FIGURE 1-1 Instant Image 2.0.1 exports shadow through STE

Solaris's `format (1M)` command into slices with 2268 granularity. A logical volume exported by volume manager software does not necessarily have this size, nor can it always be adjusted to a size that is a multiple of 2268. Similarly, raw partitions are bounded by the geometry of the device on which they are built. This presents a problem for the Instant Image 2.0.1 product which presents an image of the master locally through either a volume manager or raw device, and exports a geometry with 2268 sector granularity from STE. Padding the STE storage with a phantom tail solves the mismatch.

Figure 1-2 depicts the size of the master and shadow volume. STE will truncate the size of the shadow volume to the nearest multiple of 2268. The phantom tail should be adjusted so that the size of the exported volume is an integral multiple of 2268 blocks (512 bytes/block). The size of the exported volume is the sum of the sizes of the phantom header, phantom tail and the shadow volume. Note the following considerations.

- LUNs exported by STE will be partitioned by Solaris format command into slices with 2268 granularity.
- The storage areas allocated for Instant Image master and shadow volumes must be the same size since resynchronization may occur in either direction.
- There are volume manager and raw devices constraints that you cannot create a logical volume with a specific size.




FIGURE 1-2 Volume sizing

Independent of this issue, there is a recommended practice for providing a phantom header to separate a host's meta data (such as VTOC, Volume manager meta data) from the data exported through STE, as well as adding a phantom tail to allow for the 2 cylinders the Solaris format command reserves as alternate cylinders. The padding can be implemented as part of that process.

This entails computing the pad required and adding it to the phantom tail. In this example, your Instant Image shadow volume is `/dev/vx/rdisk/vol02`. You must determine the volume size, compute the pad and then add the pad to the standard phantom tail in `ste.cf`. In this example, volume manager software is used.

1. Determine the volume size (in blocks where 512 bytes/block).

```
# vxprint -v vol02
v vol02          fsgen          ENABLED 18874368 -          ACTIVE  -          -
```

2. Compute the pad. Add the phantom header, phantom tail and shadow volume size. Determine the pad to increase the sum to the next cylinder size.

```
PH = Phantom Header = 2268 blocks
PT = Phantom Tail = 4536 blocks
Shadow Volume = 18874368 blocks
PH + PT + Shadow Volume = 18881172 blocks
18881172/2268 = 8325.03 -> truncates to 8325 cylinders
Next cylinder = 8326
8326 cylinders * 2268 blocks/cylinder = 18883368 blocks
Pad = 18883368 - 18881172 = 2196 blocks
```

3. Add pad to phantom tail in ste.cf.

Note, the phantom tail, in this example is 2 cylinders plus the pad $(2*2268)+2196 = 6732$. This example only shows information relevant to vol02.

```
# StorEdge Target Emulation Configuration File - ste.cf
# Virtual Disk section
#Name      Partition Name          Port LUN  State    Options
#-----
vdisk1 /dev/vx/rdsk/datadg/Vol02 tm0 1  online  ph=dev2 pt=dev3
```

```
# StorEdge Target Emulation Configuration File - ste.cf
# Phantom Devices section
#keyword    Partition Name          Start Block  Size
#-----
dev2        /dev/vx/rdsk/datadg/headsandtails  13608       2268
dev3        /dev/vx/rdsk/datadg/headsandtails  15876       6732
```

Why does this work?

If the shadow is padded with a phantom tail, then the storage presented to the other host is larger than that seen as the master. A file system or database that is built on the master with *n* blocks available, will never seek past the *n*th block. Instant Image shadows this filesystem or database infrastructure to the shadow volume. A filesystem or database utilizing the infrastructure built on the master now shadowed to the shadow volume, will also not seek past the end of the size on which it was built.

Note – Do not build a new filesystem or database on Instant Image 2.0.1. shadow volumes exported through STE.

If you build a new filesystem or database on the shadow volume, you will create a new infrastructure based on a size larger than the master. A resynchronization from shadow to master will truncate the difference between the two volume sizes (up to 2267 blocks). This is not a constraint on the operation of this product. Instant Image 2.0.1 is not intended to be used to provide generic storage. Instant Image 2.0.1 provides the additional capability to export a snapshot or shadow of a master volume. The operation of Instant Image with the shadow exported to another host through STE, is the same as if you were accessing the master and shadow volumes locally. That is, you may write to either master or shadow, resynchronize in either direction, etc.

▼ To Prepare the Initiator

On the initiator system, make sure you have installed the prerequisite Fibre Channel Host Adapters, software packages and Solaris operating environment patches. Refer to <http://www.sunsolve.sun.com> to get the latest patch information.

TABLE 1-5 Minimum SOCAL Driver Patches Required

Patch Number	Description
105375-21	Solaris 2.6 SOCAL driver patch
107469-06	Solaris 7 SOCAL driver patch
109460-01	Solaris 8 SOCAL driver patch

If the initiator is PCI-based, you must install the appropriate ifp patches.

TABLE 1-6 Ifp patches

Patch Number	Description
107280-05	Solaris 2.6 ifp patch
107292-04	Solaris 7 ifp patch
109189-02	Solaris 8 ifp patch

▼ To Install the Fibre Channel Host Adapter

For installation instructions for the Fibre Channel Host Adapter, refer to the *Sun StorEdge PCI FC-100 Host Adapter Installation Manual*. Connect the target and initiator with a Fibre Channel cable.

▼ To Configure STE Devices on the Initiator

To configure the exported LUNs on the initiator host, perform the following three steps.

1. **Add the STE virtual disks to /devices.**

```
# drvconfig
```

2. **Create /dev entries for the STE virtual disks.**

```
# disks
```

3. **Label and optionally partition the virtual disks as you want.**

```
# format
```

The initiator must format the LUN such that slice 0 is the first cylinder (which maps to the size of the phantom header) and the remainder of the drive is in slice 1.

Note – These steps must be performed on the initiator every time the STE configuration changes or STE is stopped. If volume size changes, the volumes will have to be reformatted on the initiator.

▼ To Troubleshoot STE

The drive needs to be labeled if it is the first time the initiator has “seen” the drive. If the `format` command does not show the virtual disks stored on the host server, make sure that the Fibre Channel cable is plugged into the specified ports. Then check that there is not a problem with the cable itself. If necessary, also check that the GBIC connectors are functioning properly. If problems persist, recheck the STE configuration and the output of `steadm -c`. Use `scmadm -S` to check that the STE backend volumes are in use.

STE will not run in singleuser mode. Appendix A contains additional guidance for troubleshooting STE.

Sun StorEdge Target Emulation (STE)

What is STE?

STE is SCSI target emulation software (the STE driver plus one or more target mode drivers) over a Fibre Channel.

STE functionality:

- STE enables an initiator host to connect to a target server and access the attached storage as if it were one or more SCSI target devices.
- Enables directly attached hosts to access Sun StorEdge Data Services, including Instant Image, Network Data Replicator, and Fast Write Cache.
- Provides caching and read-ahead functionality for fast read/write access.
- Provides a variable logical disk partition size.
- Allows multiple LUNs to be configured on a Fibre Channel port.

STE Components

The STE product includes several utilities for managing the STE software and its attached target devices.

Utility	Function
steadm	Provides a command line interface for enabling, monitoring, and disabling STE.
steconf	Prints configuration information about Fibre Channel devices on the system.

Product Considerations

- Do not use the SOC+ Fibre Channel Host Adapter in dual mode. That is, you cannot configure one port for STE and the other port for a storage array.
- The Solaris operating environment reserves space for header information at the beginning of each disk. You may not write to this space. You need to use a phantom header as described in the “Phantom Header Section of `ste.cf` file”.
- VERITAS Dynamic Multipathing (DMP) and Alternate Pathing (AP) are not supported.
- STE 1.2 is not supported in a clustered environment.
- If you perform any reconfiguration, you need to reconfigure the `ste.cf` file.
- STE and SV can not share the same volume. If you have a volume you want under STE control, remove the volume from SV control, add the volume to the `ste.cf` file and then restart Instant Image.
- Java version 1.2.1_04b or compatible versions is required. To update your version of java, go to <http://java.sun.com>. To verify your version of java:

```
# java -version
```

- If you are loading Instant Image for the French, Japanese, or Chinese locales, load the appropriate localized version of Java.
- There are several patches required to use STE. It is recommended that the most recent revision of these patches be loaded before the STE packages are installed.
- The latest Solaris Recommended Patch Cluster is to be installed before loading Data Services software. The following patches, which are available at <http://sunsolve.sun.com> when you have a service contract, are the minimum revisions required:

TABLE A-1 Solaris Environment Patches for STE

Solaris Environment Version	Patch	Description
2.6	105181	Kernel Update patch is to be installed before loading data services software. This patch is included with the Solaris Recommended Patch Cluster.
	105375	5.6 sf and socall drivers
	105356	5.6 /kernel/drv/ssd and /kernel/drv/sd

TABLE A-1 Solaris Environment Patches for STE (Continued)

Solaris Environment Version	Patch	Description
7	108091	Sun OS 5.6: ss JDK 1.2.1_04b
	106541	Kernel Jumbo patch is to be installed before loading data services software. This patch is included in the Solaris Recommended Patch Cluster.
	107458-10	5.7 dad, sd, ssd, uata drivers patch
8	107469	5.7 sf and social drivers
	108528	Kernel Update patch is to be installed before loading data services software. This patch is included with the Solaris Recommended Patch Cluster.
	108974	5.8 dada, uata, dad, sd and scsi patch
	109460	5.8 sf and social drivers
	109524	5.8 ssd driver patch

Installing STE

You may install STE either automatically or manually. The first installation procedure explains how to install the packages automatically and the second explains the manual process.

Note – The installation procedures assume that the system is running Sun Enterprise Volume Manager™. If you are not using Volume Manager to manage diskettes and CDs, refer to the *Solaris System Administration Guide, Volume I* for detailed information about managing removable media with Volume Manager. If you inserted a CD, Volume Manager automatically mounts it.

Note – You should load data services software in a single user state.

▼ To Install the STE Software Packages Automatically

1. **Become superuser (root) on your target system.**
2. **On a system running the Solaris 2.6 environment, add this line to the `/etc/system` file:**

```
set kobj_map_space_len=0x200000
```

If you have not set this, the system may hang on reboot.

3. **Reboot the system.**

You only need to reboot if you are using the Solaris 2.6 operating environment.

4. **Remove STE, if necessary.**

If an earlier version of STE is installed, you must remove it from your target system. See “To Remove STE” on page 8.

5. **If the STE CD is not already in the CD-ROM, insert it into the CD-ROM drive connected to your system.**

6. **Start the Volume Manager daemon (if needed) and start the STE installation.**

To install the package automatically where the value of the variable *Solaris_version* is either *Solaris_2.6*, *Solaris_7* or *Solaris_8*:

```
# /etc/init.d/volmgt start
# cd /cdrom/cdrom0/ste/Solaris_version
# ./install_ste.ksh
```

`install_ste.ksh` installs the data services, `java/framework/sundae` and the STE service.

▼ To Install the STE Software Packages Manually

1. Become superuser (root) on your target system.

2. Remove STE, if necessary.

If an earlier version of STE is installed, remove it from your target system. See “To Remove STE” on page 8.

3. On a system running the Solaris 2.6 environment, add this line to the `/etc/system` file:

```
set kobj_map_space_len=0x200000
```

If you have not set this, the system may hang on reboot.

4. Reboot the system.

You only need to reboot for the Solaris 2.6 operating environment.

5. If the STE CD is not already in the CD-ROM, insert it into the CD-ROM drive connected to your system.

6. Start the Volume Manager daemon (if needed) and start the STE installation.

Follow the order of installation when you add packages manually for a successful installation. Where the value of the variable *Solaris_version* is either *Solaris_2.6*, *Solaris_7* or *Solaris_8*, enter:

```
# /etc/init.d/volmgt start
# cd /cdrom/cdrom0/ste/Solaris_version
# pkgadd -d . SUNWspuni SUNWscm
# pkgadd -d . SUNWte SUNWsftm
```

7. Type `y` in response to all prompts.

If a package to be installed currently exists on your server, make sure you install the most recent version.

Post-Installation Patches

After installing the STE packages, install the appropriate patches.

▼ Complete the Post-Installation

1. Install the appropriate patches.

The patches listed here are required. Patches are included on the software CD, but you should check <http://www.sunsolve.sun.com> to make sure you have the latest revision of the patches.

Solaris Environment	Patch	Description
2.6	109967 ¹	data services core patch
	109983 ¹	Sun StorEdge Target Emulation patch
7	109969 ²	data services core patch
	109985 ²	Sun StorEdge Target Emulation patch
8	109970 ³	data services core patch
	109986 ³	Sun StorEdge Target Emulation patch

¹Load patch 109967 before you load 109983. If later revisions are available, use them.

²Load patch 109969 before you load 109985. If later revisions are available, use them.

³Load patch 109970 before you load 109986. If later revisions are available, use them.

2. If you are installing other Data Services, eject the STE CD and continue.

Otherwise, eject the CD and reboot the system:

```
# cd /
# eject cdrom
# /etc/shutdown -y -g 0 -i 6
```

Operating STE

STE initializes at system startup. To stop STE or change its configuration, you may do so manually with the `steadm(1M)` utility. Refer to the `steadm` man page for more information.

▼ To Disable STE

- To disable STE (all targets and vdisks), issue the following command:

```
# steadm -d
```

- To disable only a specific Target Mode Fibre Channel Port or Virtual Disk, specify the Port Name or Vdisk Name as an argument. For example:

```
# steadm -d vdisk6  
# steadm -d tm0
```

Note – Each port must have a LUN 0; therefore, you cannot disable a virtual disk that is being exported on LUN 0. If the virtual disk on LUN 0 is the only virtual disk on the port, use `steadm -d portname` to disable it.

▼ To Reconfigure STE

Note – To reconfigure an existing Target Mode Fibre Channel Port or Virtual Disk, you must disable, then enable it again.

1. Disable any existing Target Mode Fibre Channel Port(s) and/or Virtual Disk(s) that you want to reconfigure.
2. Using a text editor, modify the STE configuration file `/etc/opt/SUNWte/ste.cf` as desired.
3. Enable the new or modified Target Mode Fibre Channel Port(s) and/or Virtual Disk(s).

▼ To Enable STE

- To enable STE, use the following command:

```
# steadm -e
```

- To enable a disabled Target Mode Fibre Channel Port or Virtual Disk, specify the Port Name or Vdisk Name as an argument. For example:

```
# steadm -e vdisk6  
# steadm -e tm0
```

▼ To Remove STE

1. Stop STE:

```
# steadm -d
```

2. Back out STE and Data Services patches.

If patch 1099983 (Solaris 2.6 environment), 109985 (Solaris 7 environment) or patch 109986 (Solaris 8 environment) is installed, remove it. Also remove any other revisions. For example:

```
# showrev -p | grep 109985  
Patch: 109985-05 Obsoletes: Requires: 109969-05 Incompatibles: \  
Packages: SUNWte, SUNWsftm  
# patchrm 109985-05
```

If there are no other data services on the system, remove patch 109967 (Solaris 2.6 environment) or 109969 (Solaris 7 and 8 environments). Also remove any other revisions. For example:

```
# showrev -p | grep 109969  
Patch: 109969-05 Obsoletes: Requires: 106541-06 Incompatibles: \  
Packages: SUNWspuni, SUNWscm, SUNWspsv, SUNWspcsl  
# patchrm 109969-05
```

3. Remove the STE packages.

When uninstalling, the order in which you remove packages does matter.

```
# pkgrm SUNWsftm SUNWte
```

4. If you are not going to reinstall STE, remove the `/kernel/drv/socal.conf` and `/etc/opt/SUNWte/ste.cf` files.

5. If there are no other Data Services on the system, uninstall the core Data Services packages:

```
# pkgrm SUNWscm SUNWspuni
```

If you are not going to reinstall `SUNWscm` and `SUNWspsv`, and you wish to completely remove them, remove the `/etc/opt/SUNWspsv/sv.cf`, `/etc/opt/SUNWscm/sd.cf` and `/etc/opt/SUNWscm/bitmapfs.cf` files if they exist.

6. If you have other Data Services to remove, continue to do so. If this is the last data service that you are removing, reboot the system now:

```
# /etc/shutdown -y -i 6 -g 0
```

Note – If you are going to reinstall any data services software, reboot the system into a single user state before you load the new data services software.

Note – To return any SOC+ ports to non-STE operation, reboot the system.

▼ To Upgrade or Reinstall STE

1. **Remove the existing version of STE from your target system, as described in Removing STE.**

Multiple instances of the STE packages cannot coexist on a target system.

2. **Follow the procedures in Installing STE.**
3. **If you are upgrading from a previous version of STE, merge configuration file changes into your existing STE configuration file.**

`pkgadd(1M)` does *not* overwrite an existing `/etc/opt/SUNWte/ste.cf` file; therefore, examine the new sample configuration file installed as part of the upgrade (`/usr/opt/SUNWesm/SUNWte/etc/ste.cf.sample`) and merge any changes into your existing STE configuration file.

▼ To Upgrade the Operating System

If you upgrade your operating system you must follow these steps.

1. **Save the `/etc/opt/SUNWte/ste.cf` and the `/kernel/drv/socal.conf` files for use when reconfiguring STE.**
2. **Uninstall STE 1.2 software and drivers.**
3. **Upgrade your operating system.**
4. **Reinstall STE 1.2 software and drivers specific for your operating system. Refer to ‘To Install STE’.**

STE Configuration Files

About the `socal.conf` File

The entries in the server's `/kernel/drv/socal.conf` file (created by executing the `/usr/opt/SUNWesm/SUNWte/tools/stesocal.ksh` script) correspond to the entries in the server's `ste.cf` file as shown below. The `socal.conf` file is required to put the appropriate ports into target mode for use by STE.

Note – Any change to the `socal.conf` file requires a reboot of the target server.

The following is a sample `socal.conf` file.

```
name="SUNW,socal" parent = "sbus@2,0"
reg=0x0, 0x0, 0x10000, 0x0, 0x10000, 0x10000, 0x0, 0x20000, 0x18
port0-loop-id=7 port1-loop-id=7;
name="SUNW,socal" parent = "sbus@a,0"
reg=0x0, 0x0, 0x10000, 0x0, 0x10000, 0x10000, 0x0, 0x20000, 0x18
port0-loop-id=7 port1-loop-id=7;
```

In the example above, ports 0 and 1 of the SOC+ HBA at `sbus@2,0/SUNW,socal@0,0` (board 1, slot 0), and ports 0 and 1 of the SOC+ device at `sbus@a,0/SUNW,socal@0,0` (board 5, slot 0) are reserved for target mode.

There is one entry in the `/kernel/drv/socal.conf` file for each SOC+ adapter. The entries in the `socal.conf` file conform to the standard `driver.conf(4)` and `sbus(4)` configuration file formats as shown below:

TABLE A-2 The `socal.conf` File Properties

Property	Description
name	The name of the SOC+ HBA. It should be set to <code>name="SUNW,socal"</code> .
parent	The full path name of the parent bus. It should be set to <code>sbus@X,0</code> where <i>X</i> matches the SBus name for the SOC+ adapter as displayed by <code>steconf</code> .
reg	An arbitrary length array in which each element consists of a 3-tuple of integers describing the mappable resources on the SBus. It should be set to <code>0xZ, 0x0, 0x10000, 0xZ, 0x10000, 0x10000, 0xZ, 0x20000, 0x18</code> , where <i>Z</i> matches the slot number for the SOC+ adapter as displayed by <code>steconf</code> .
port0-loop-id	The Loop_ID to be used by port 0 in target mode. It should be set to match the Loop ID value in the <code>ste.cf</code> file. Note that once you set this property, you can only use the port in target mode (in other words, you cannot use that port as a SCSI initiator). Loop IDs can range from 0 to 15.
port1-loop-id	The Loop_ID to be used by port 1 in target mode as described above.

About the `ste.cf` File.

An `ste.cf` file includes three sections:

- The first section identifies the target mode Fibre Channel ports.
- The second section identifies the virtual disks to link target mode ports with a disk.
- The third section identifies Phantom Headers and Phantom Tails.

Note – Do not use headers and tails when you configure NT LUNs in the `ste.cf` file because Instant Image does not successfully make a shadow copy of an NT master Volume.

Target Mode Fibre Channel Port Section of `ste.cf` File

The Target Mode Fibre Channel Port section of the `ste.cf` file contains one line per Target Mode Fibre Channel Port. For example:

Sample `ste.cf` File (Target Port Section)

```
# StorEdge Target Emulation Configuration File - ste.cf
# Target Mode Fibre Channel Port section

#Port
#Name  Driver  Device Name                                Loop Id
#-----
tm0     sftm    /devices/sbus@2,0/SUNW,socal@0,0:0        7
tm1     sftm    /devices/sbus@2,0/SUNW,socal@0,0:1        7
tm2     sftm    /devices/sbus@a,0/SUNW,socal@0,0:0        7
tm3     sftm    /devices/sbus@a,0/SUNW,socal@0,0:1        7
```

The four entries in the example above define the four ports in the sample `socal.conf` file and give them the names `tm0`, `tm1`, `tm2`, and `tm3`. You can use the `steconf` utility to determine the driver, device name, and physical location.

Each line of the Target Mode Fibre Channel Port section consists of these four parameters:

TABLE A-3 Properties of the Target Mode Fibre Channel Port section of the `ste.conf` File

Parameter	Description
Port Name	The keyword “tm” followed by a unique integer to define which port is being configured.
Driver	The name of the target mode driver used to control the Fibre Channel hardware. In the example, it is <code>sftm</code> .
Device Name	The path to the Target Mode Fibre Channel Port.
Loop ID	Determines the ID that the Target Mode Fibre Channel Port will present on the Fibre Channel loop. Loop IDs can range from 0 to 15.

Virtual Disk Section of `ste.cf` File

The virtual disk section contains one line per virtual disk. For example:

Sample `ste.cf` File (Virtual Disk Section)

# StorEdge Target Emulation Configuration File - <code>ste.cf</code>					
# Virtual Disk section					
#Virtual disks					
#Name	Partition Name	Port	LUN	State	Options
#-----	-----	---	---	-----	-----
vdisk0	/dev/vx/rdisk/datadg/Vol01	tm0	0	online	ph=dev0 pt=dev1
vdisk1	/dev/vx/rdisk/datadg/Vol02	tm0	1	online	ph=dev2 pt=dev3
vdisk2	/dev/vx/rdisk/datadg/Vol03	tm0	2	online	ph=dev4 pt=dev5
vdisk3	/dev/vx/rdisk/datadg/Vol04	tm1	0	online	ph=dev6 pt=dev7
vdisk4	/dev/vx/rdisk/datadg/Vol05	tm1	1	online	ph=dev8 pt=dev9
vdisk5	/dev/vx/rdisk/datadg/Vol06	tm1	2	online	ph=dev10 pt=dev11

Each line of the Virtual Disk section consists of the following six parameters:

TABLE A-4 Properties of the Virtual Disk section of the `ste.conf` File

Parameter	Description
Name	The keyword “vdisk” followed by a unique integer to define which vdisk is being configured.
Partition Name	The path to the logical volume that will be presented on the Target Mode Fibre Channel Port.
Port	The Target Mode Fibre Channel Port for this virtual disk. The port name must match one of the entries in the Port Configuration section.
LUN	The logical unit number to be presented on the Target Mode Fibre Channel Port. There must exist a LUN0 for each target mode port specified.
State	Enables the user to specify “online” or “offline” for Virtual Disks. If a Virtual Disk’s state is set to “offline”, it will be offline when STE software starts and will only be enabled when the user issues an <code>steadm -e</code> command referencing that specific Virtual Disk as an argument. The recommended practice is to set “online” as the state.
Options	<p>Enables the user to specify additional configuration options for the Virtual Disk. Available options are:</p> <p><i>ph=phantomheader</i> and/or <i>pt=phantom tail</i></p> <p>Phantom headers and tails are used to protect disk-related data from being over-written. See “Protecting Disk-Related Data” for additional information. The <i>phantomdevice</i> variables must match one of the Device Name entries in the Phantom Partition Configuration section.</p> <p><i>ro</i></p> <p>Read Only. For example, if two hosts are sharing a disk you should specify that one is read only. Although the initiator cannot write to the data on the virtual disk, the VTOC can be overwritten.</p>

Phantom Devices Section of `ste.cf` File

The Phantom Devices section of the STE configuration file contains one line per Phantom Device. Phantom headers and tails are used to protect disk-related data from being over-written. See “Protecting Disk-Related Data” for additional information on phantom headers and tails. See “Configuration Considerations” section of this manual for additional guidelines.

For example:

Sample ste.cf File (Phantom Header Section)

# StorEdge Target Emulation Configuration File - ste.cf			
# Phantom Header section			
#device			
#keyword	Partition Name	Start Block	Size
#-----	-----	-----	----
dev0	/dev/vx/rdisk/datadg/headsandtails	0	2268
dev1	/dev/vx/rdisk/datadg/headsandtails	2268	4536
dev2	/dev/vx/rdisk/datadg/headsandtails	6804	2268
dev3	/dev/vx/rdisk/datadg/headsandtails	9072	4536
dev4	/dev/vx/rdisk/datadg/headsandtails	13608	2268
dev5	/dev/vx/rdisk/datadg/headsandtails	15876	4536

Each line of the Phantom Header section consists of these four parameters:

TABLE A-5 Properties of the Phantom Header section of the `ste.conf` File

Parameter	Description
Device Keyword	The keyword “dev” followed by a unique integer to define which phantom device is being configured.
Partition Name	The partition on which the phantom devices reside.
Start Block	The starting block of the phantom device on the partition.
Size	The size (in blocks) of the phantom device.

Certain initiators set aside space at the beginning and/or end of each disk for metadata; the required size and location (at the beginning or end of the disk) are system-specific. However, we recommend a size of 2268 blocks (1 cylinder) for the phantom header and 4536 blocks (2 cylinders) for the phantom tail. Additional blocks may be required to be added to the phantom tail per the guidelines presented in the “Configuration Considerations” section of this manual.

The phantom devices physically reside on a separate phantom partition but appear to the initiator to be part of the shared virtual disk. To protect against data loss, the phantom partition should be a mirrored or RAID 5 partition.

Length of Phantom Headers and Tails

In this example, the phantom header size of 2268 is a result of the STE disk geometry for the LUNs that it exports:

- 512 bytes/sector

- 108 sectors/track
- 21 tracks/cylinder
- n cylinders

where n is adjusted so that total capacity matches the size of the logical volume being exported. Using this geometry, one cylinder is $108 \times 21 = 2268$ sectors (blocks). Therefore, the header = 2268 blocks and the tail = 4536 blocks.

To Protect Disk-Related Data

Certain applications store their private disk-related metadata at the beginning and end of disks. For example, the Solaris operating environment stores the disk label which contains partition information in the first 16 blocks of the disk.

When the initiator first detects a new LUN, it must be labeled. The `format` utility writes the label in the first 16 blocks of the LUN presented to it.

Volume Table of Contents (VTOC)

In a standard STE configuration, there is potential for the initiator to overwrite the original Volume Table of Contents (VTOC). The VTOC is written into the first 16 blocks of cylinder 0. When the target exports the device, a phantom header and tail must be configured. When the initiator formats the exported device, it will then write its VTOC into the first 16 blocks of cylinder 0, in the space allocated, to the phantom header. This preserves the original VTOC.

The original VTOC can be destroyed if the initiator writes data starting at block 0. To avoid corrupting the VTOC, run the `format` command to properly configure the exported device. The initiator formats the LUN such that slice 0 is the first cylinder and the remainder of the drive is in slice 1.

To configure a device to be exported to an initiator:

1. On the target, create a phantom header. This will be used by the initiator for VTOC information.
2. Export the new vdisk that contains the data partition, and the phantom devices.
3. On the initiator, make cylinder 0 slice 0. This cylinder receives the VTOC from the initiator. The data will start on slice 1, which begins on cylinder 1.
4. Make the rest of the volume slice 1. Slice 1 contains all data.

The following figure depicts the use of phantom devices.




FIGURE A-3 Vdisk representation

The Use of Phantom Devices

Certain initiators set aside space at the beginning and end of each disk for header and other system information. The required size and location are system-specific. To prevent these initiators from writing their header information over the contents of the shared device, you must prepend a phantom header and append phantom tail. The phantom headers and tails physically reside on separate partitions. These phantom headers and tails appear to the initiator as part of the shared virtual disk.

Initiators write header and system information to the first cylinder of the exported device. If a phantom header doesn't exist, the system will write to cylinder 0 block 0 of the device. This will result in data loss from the volume you are exporting as a shadow.

The format utility allocates the last 2 cylinders as "alternate" cylinders. These cylinders are allocated at the end of the device. When the initiator does a `format (1M)`, it will take the last two cylinders of the data space and allocate them as alternates. By the inclusion of a phantom tail the initiator will use the phantom tail for its alternate cylinders.

See "Configuration Considerations" section of the Instant Image 2.0.1 Installation Guide for additional guidelines on setting up phantom devices.

Troubleshooting

The following section presents possible problems, their causes and the corrective actions to resolve the issues.

Error: Cannot load module 'sftm'.

STE 1.2 is installed and configured. The system is rebooted and a message similar to the following is displayed.

```
NOTICE: STE: Polling has been started
May 11 10:33:08 ns-east-32 unix: /usr/kernel/drv//sparcv9/sftm:
undefined
symbol 'socal_sftm_attach'
May 11 10:33:08 ns-east-32 root: Error executing
/usr/opt/SUNWesm/SUNWte/sbin/steadm start. Abandoned
May 11 10:33:08 ns-east-32 unix: /usr/kernel/drv//sparcv9/sftm:
undefined
symbol 'socal_sftm_detach'
May 11 10:33:08 ns-east-32 unix: WARNING: mod_load: cannot load
module 'sftm'
May 11 10:33:08 ns-east-32 root: Error executing
/usr/opt/SUNWesm/SUNWte/sbin/steadm start. Abandoned
```

Cause: SOCAL driver patch not loaded.

Corrective Action: Add the proper SOCAL patch as defined below.

Patch Number	Description
105375-21	Solaris 2.6 SOCAL driver patch
107469-06	Solaris 7 SOCAL driver patch
109460-01	Solaris 8 SOCAL driver patch

Error: Error executing steadm start

STE 1.2 is installed and configured. The system is rebooted and a message similar to the following is displayed.

```
NOTICE: STE: Polling has been started
May 11 12:30:13 ns-east-32 root: Error executing
/usr/opt/SUNWesm/SUNWte/sbin/steadm start. Abandoned
May 11 12:30:13 ns-east-32 root: Error executing
/usr/opt/SUNWesm/SUNWte/sbin/steadm start. Abandoned
```

Cause: Either the socal.conf file does not exist or a driver entry is incorrect in the ste.cf file.

Corrective Action: Correct the driver entry in the ste.cf file and run the /usr/opt/SUNWesm/SUNWte/tools/stesocal.ksh.

Error: Unable to detect exported volumes

You are unable to detect the exported volumes while running format (1M) on the initiator.

Possible Causes: Hosts not connected by Fibre Channel Cables.
LUNs not exported.

Corrective Action 1: Verify the Fibre Channel cable is connected to the correct port on the target.

Corrective Action 2: On the target, run `/usr/opt/SUNWesm/sbin/steadm -c` to confirm the LUNs are being exported.

Corrective Action 3: On the initiator, run `drvconfig`, `disks` and `format`.

Error: Drive Not Available

The STE volumes are labeled using the `format (1M)` command. At a later time, you run the `format (1M)` command and a message similar to the following is displayed:

```
c3t7d2 <drive not available:  formatting>
      /sbus@1f,0/SUNW,  socal@2,0/sf@1,0/
```

Cause: Connection lost with target.

Corrective Action: Check the Fibre Channel connection.

Error: Modified volume size not recognized on the initiator

Cause: A valid disk label already exists.

Corrective Action: Execute the `/usr/opt/SUNWesm/SUNWte/tools/stezap.ksh` script to erase the VTOC from the specified phantom header.

This may be necessary when a volume is resized, but the initiator fails to recognize the new size. In order for the initiator to see the new size, the phantom header must be wiped clean.

The usage for this command is:

```
# stezap.ksh [ -f ste.cf_file ] dev# [dev# ...]
```

where `ste.cf_file` is the location of the `ste.cf` file.

If no file is specified, the default location (`/etc/opt/SUNWte/ste.cf`) is used.

To wipe the phantom headers, specify the dev numbers (from the `ste.cf` file) that correspond to the phantom headers.

For instance, if “dev0” in the `ste.cf` file is a phantom header for a volume, and it needs to be wiped, type the following:

```
# stezap.ksh dev0
```

This will automatically read the `ste.cf` file and determine the appropriate device, offset and block count to wipe. After this is done, the initiator is ready to re-label the drive.

Glossary

bitmap	The file or volume that tracks the location of data changes on the shadow or master volume.
dependent	Implies a volume- <i>dependent</i> point-in-time view. When you enable Instant Image 2.0 for a volume pair, you can specify the shadow volume as <i>dependent</i> . This does not cause Instant Image 2.0 to perform a full volume copy; data is not duplicated until you request it from the shadow. The dependent shadow volume is a copy that relies on the master for all unmodified data blocks.
HBA	Host Bus Adapter
independent	A physically-separate full volume copy. It remains so after Instant Image 2.0 is disabled.
initiator host	A device that makes requests to the target device.
LUN	Logical Unit Number. The LUN identifies a unit on a particular device.
master	The volume that contained the original data, copied to the <i>shadow volume</i> .
point-in-time view	A copy or image of volume data captured or copied at a particular moment.
shadow	The volume containing a point-in-copy of data from the <i>master</i> volume.
STE	Sun StorEdge Target Emulation
SV	Storage Volume
target	A device that receives and responds to requests from an initiator. The Instant Image 2.0 and STE 1.2 Data Services reside on the target.
volume	A device representing actual physical disk slices, or a device created by volume manager software.
volume group	The master volume and its associated shadow and bitmap volumes.

volume pair The master and shadow volumes in a volume group.