



Sun StorEdge™ Instant Image 2.0.1 Configuration Guide

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

Part No. 806-5318-10
June 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 94303-4900 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. For Netscape Communicator™, the following notice applies: Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, Instant Image, Sun Enterprise, AnswerBook2, Solstice Disk Suite, SPARCstorage, 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.

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 94303-4900 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. La notice suivante est applicable à Netscape Communicator™: Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, the Sun logo, Sun StorEdge, Instant Image, AnswerBook2, docs.sun.com, Solstice Disk Suite, SPARCstorage, 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.

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 v

How This Book Is Organized v

Using UNIX Commands v

Documentation Conventions vi

Related Documentation vii

Accessing Sun Documentation Online vii

Sun Welcomes Your Comments viii

Sun StorEdge Instant Image 2.0.1 Configuration Guide 1-1

Overview 1-1

Introduction 1-2

System Requirements 1-2

Limitations 1-2

Architecture 1-3

 Hardware Components 1-3

 Software Components 1-3

 Instant Image 2.0 1-3

 STE 1.2 1-4

Implementation Overview 1-4

1. Prepare the target server 1-4
2. Install STE on the target server 1-4
3. Install Instant Image on the target server 1-4
4. Connect the initiator to the target server 1-4
5. Perform Instant Image shadow volume activities 1-4
6. Back up the data 1-6

Configuration Considerations 1-6

Why do I want to export a volume? 1-6

Which volumes do I want to export? 1-7

What level of redundancy should I use to protect the data? 1-7

What is the overhead on the target host with secondary accesses? 1-7

Can I mount the shadow volume on the target host? 1-7

When should I resynchronize the master to the shadow? 1-7

When should I resynchronize the shadow to the master? 1-8

How does this work with Sun Cluster? 1-8

How does this compare with the Sun StorEdge Network Data Replicator solution? 1-8

When should I use dependent mode? 1-8

When should I use independent mode? 1-8

Where should I put the bitmaps? 1-9

Do I need Fast Write Cache? 1-9

GLOSSARY Glossary-1

Preface

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

How This Book Is Organized

Chapter 1 describes the architecture and configuration considerations for Instant Image 2.0.1.

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	<i>ii_control</i> (1M)	N/A
	<i>ii_config</i> (1M)	
	<i>ii_health</i> (1M)	
	<i>esm_cli</i> (1M)	
	<i>steconf</i> (1M)	
	<i>steadm</i> (1M)	
	<i>ste.cf</i> (4M)	
	<i>sftm</i> (7D)	
	<i>ste</i> (7D)	
	<i>socal</i> (7D)	
	<i>svadm</i> (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 2.0.1 Installation Guide</i>	806-5317
	<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

Accessing Sun Documentation Online

The `docs.sun.comsm` web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

`http://docs.sun.com`

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.

Sun StorEdge Instant Image 2.0.1 Configuration Guide

Overview

Sun StorEdge™ Instant Image is a point-in-time copy and snapshot facility for the Solaris™ operating environment. It enables point-in-time shadow volumes to be created and independently accessed once the shadow volume has been separated from its master. Instant Image tracks differences between the master and the shadow volumes to accelerate later resynchronization.

Instant Image 2.0.1 provides the capability to export a snapshot or shadow 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 the *Sun StorEdge Instant Image 2.0.1 Installation Guide*.

This document contains the following main topics:

- Introduction

- System Requirements
- Limitations
- Architecture
- Implementation Overview
- Configuration Considerations

Before reading this guide, you should have already installed the Instant Image and STE software according to the instructions in the *Sun StorEdge Instant Image 2.0.1 Installation Guide*.

Introduction

Instant Image enables more than one application to access a point-in-time copy of the data on your system. With Instant Image 2.0.1, you can utilize a point-in-time copy for application testing or any situation where it may be advantageous to use a copy of a complete data set. This configuration would be useful for the following applications:

- Backup- backup on-line data
- Application testing - Improve testing by using current data
- Sharing data - Make a point-in-time copy available to an alternate host

System Requirements

Sun StorEdge™ Instant Image 2.0.1 has been qualified on the following platforms:

- Operating Systems: Solaris 2.6, 7, 8
- Target servers: Sun Enterprise ES3500, ES4x00-ES6x00, ES10000
- Initiator servers: Any Sun server that supports either a PCI or SBus card
- Storage device: Any storage system supported by Sun

Limitations

Sun StorEdge™ Instant Image 2.0.1 is not supported in a clustered environment.

Sun StorEdge Instant Image 2.0.1 does not support:

- PCI-based target servers
- VERITAS Dynamic Multipathing (DMP)
- Dynamic Reconfiguration (DR)
- Alternate Pathing (AP)

Architecture

Hardware Components

Instant Image 2.0.1 requires two or more servers connected by FC-AL cables. The target server, where the Instant Image and STE software reside, must have a SOC+ FC-AL interface card.

Software Components

Instant Image 2.0.1 is comprised of the following software components:

- Sun StorEdge Instant Image 2.0
- Sun StorEdge Target Emulation 1.2

Instant Image 2.0

The Instant Image data service is a point-in-time volume copy facility for Solaris. It creates a volume group containing an original logical volume (the master), a point-in-time copy of the original volume (the shadow) and a volume used to track the differences between the other two (the bitmaps or scoreboards). Once the Instant Image group has been established, the master and shadow volumes can be accessed independently through the Instant Image module. Two types of shadow volumes can be created, an independent shadow volume or a dependent shadow volume. An independent shadow volume is a complete physical copy of the master volume. A dependent shadow is a point-in-time copy of the master volume that relies on the master volume for accesses to all unmodified data blocks. When data blocks are modified by the master, the old data is written to the shadow.

STE 1.2

The STE software exports volumes on the target server to the initiator as physical devices. STE allows the initiator access to Sun StorEdge data services software on the target server.

Implementation Overview

This section provides a brief overview of the procedure to install the software, create an Instant Image point-in-time copy of the data and then perform a back-up of the data. Details of the installation procedure can be found in the *Sun StorEdge Instant Image 2.0.1 Installation Guide*.

1. Prepare the target server

Ensure that the target server is at the proper revision for all firmware and driver patches.

2. Install STE on the target server

Install STE 1.2 in accordance with the *Sun StorEdge Instant Image 2.0.1 Installation Guide*. You must also update the `ste.cf` file with the appropriate Target Mode Fibre Channel Ports, Virtual Disks and Phantom Partitions. Then, run the `/usr/opt/SUNWesm/SUNWte/tools/stesocal.ksh` script to update the `socal.conf`.

3. Install Instant Image on the target server

Install Instant Image 2.0 on the target server, in accordance with the *Sun StorEdge Instant Image 2.0 Installation Guide*. The master volume must be under SV control while the shadow volume must be under STE control. After updating the `sv.cf` file, use the `svadm` utility to reconfigure the SV driver to add the new volumes in the configuration file to the SV layer.

4. Connect the initiator to the target server

The two servers are connected by FC-AL cables. The target server where the Instant Image and STE software reside must have a SOC+ FC-AL interface card. The initiator can have either a SOC+ or PCI FC-AL interface card.

5. Perform Instant Image shadow volume activities

Dependent Shadow Volume Activities

The dependent shadow volume provides a virtual view to master data frozen at a specific point in time. It appears as a usable copy of a master volume's data. Users can read from and write to the shadow volume just like any other volume.

You can minimize the downtime associated with application backup by using a dependent shadow volume.

a. Create a dependent shadow volume

At the selected point-in-time, initiate a dependent copy that creates a dependent shadow volume. From this point-in-time forward, writes to either the master or shadow will drive a copy on write as required to preserve the point-in-time data.

A bitmap is used to keep track of differences that occur after the established point-in-time. This bitmap records the blocks that have changed between the master and shadow after the established point-in-time and also tracks what has been copied during a copy or update operation.

b. Access the dependent shadow volume

Accesses to the shadow volume of data blocks that have not been modified since the point-in-time snapshot was initiated are serviced from the master volumes. Accesses to the shadow volume of data blocks that have been modified are serviced from the shadow to help ensure the consistent snapshot of data.

c. Resynchronize the shadow to the master volumes

Resynchronization is the process of re-establishing the point-in-time copy. One way to reduce the time it takes to resynchronize is to log the changes that occur after the point-in-time has been established. Then, to resynchronize, the user need only apply the changes rather than do a full disk to disk copy.

To resynchronize the shadow with the master, the user can simply take another dependent snapshot. This clears the bitmaps that are tracking the differences. The data is available during the resynchronization process.

d. Resynchronize the master to the shadow volume

To resynchronize the master to the shadow, apply the shadow to the master volumes that include any updates made to the shadow. The data is available during the resynchronization.

Independent Shadow Volume Activities

Independent Shadow Volumes are complete independent copies of the master data.

a. Create an independent shadow volume

At the selected point-in-time, create an independent copy that initiates a full-volume copy operation between the master and shadow volumes. Once complete, the shadow volumes are identical to the master volumes at the point-in-time the snapshot was initiated.

The shadow volume may be accessed before the copy is complete. During the copy operation, the shadow is accessed as it would be if a dependent copy were initiated. Once the copy is complete, the shadow is accessed as an independent copy of the data.

A bitmap is used to keep track of differences that occur after the established point-in-time. This bitmap records the blocks that have changed between the master and shadow after the established point-in-time.

b. Access the independent shadow volume

Once the copy operation from master to shadow volume is complete, an independent copy of data exists. The secondary application can now access a copy of the data without affecting the performance of the primary application.

c. Resynchronize the shadow and master volumes

After the snapshot is taken, the Instant Image software tracks the differences between the master and shadow volumes. When the next snapshot is taken, only the blocks that have been changed need to be copied from master to shadow (or from shadow to master). The data is available during the resynchronization.

6. Back up the data (optional)

The backup utility will typically be located on a second server. The backup utility may access a dependent shadow volume copy. The backup activities are then performed on a second host without impacting the performance on the target host or the availability of the data on the master volume. Examples of backup utilities include Solstice Backup software, VERITAS Netbackup™, and `ufsdump/restore`.

Configuration Considerations

Why do I want to export a volume?

If you export a volume, a second host can access the data using applications that reside on the second host (initiator). The processing associated with running the application is therefore attributed to that second host. This is useful for running

back-up utilities or performing any data processing. You then have a “normal” production environment on one host while any additional applications (back-ups, testing/validation) can be performed on a second host.

Which volumes do I want to export?

You can export data volumes. You cannot export `/usr`, `/opt`, `/export/home`, `/var`.

What level of redundancy should I use to protect the data?

In the event of a disk failure, it would be difficult to restore the phantom devices (headers and tails) and the bitmap volumes. These volumes are recommended to be mirrored.

What is the overhead on the target host with secondary accesses?

There is some overhead associated with the access by the second host. The primary, as owner of the volumes, is impacted by serving the I/O requests.

Can I mount the shadow volume on the target host?

No. While the volume is under STE control, you cannot mount the volume on the target host.

When should I resynchronize the master to the shadow?

The shadow volume can be resynchronized with the master when you determine a more “current” copy is necessary. You should allow any activities associated with the shadow volume to complete prior to beginning the resynchronization process. Verify that there is no I/O processing on the exported LUN. The master and shadow volumes cannot be mounted on the target or the initiator while performing the resynchronization.

When should I resynchronize the shadow to the master?

You would resynchronize the master with the shadow if it was determined that the master is out-of-date. Typically, you would update the master with the shadow if you thought the integrity of the data on the master was in question.

How does this work with Sun Cluster?

Instant Image 2.0.1 is not supported with Sun Cluster.

How does this compare with the Sun StorEdge Network Data Replicator solution?

Sun StorEdge Network Data Replicator (Sun SNDR) is designed to be an online remote data replicator. It is designed to be active during normal application access to the data volumes, and it will continually replicate the data to the remote site. SNDR can be employed in Disaster Recovery and Business Continuity strategies to provide redundant storage of critical information across physically separate sites.

With Instant Image 2.0.1, you export a point-in-time copy to be used by a remote host (initiator) at the same site. All data resides on the target.

When should I use dependent mode?

Dependent shadow volumes rely on master volumes for accessing all unmodified data blocks. One advantage is that the resynchronization is nearly instantaneous as no copy needs to occur to resynchronize the shadow with the master. You may wish to utilize a dependent copy in conjunction with a backup utility to perform your backup operations. You can minimize the downtime associated with application backup by using a dependent shadow volume. You can stop the application, make a dependent copy and then make the shadow volume available to another host for backup.

When should I use independent mode?

Independent shadow volumes are complete independent copies of the master data. When performing an independent copy, a full volume copy is executed. This is more time consuming than a dependent copy. Once an independent copy is made, you can perform testing without impacting the performance of the primary application.

Where should I put the bitmaps?

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.

There are other options available. Bitmap files in local memory would provide performance enhancements but would be unrecoverable in the event of a system crash. The use of Sun StorEdge Fast Write Cache is recommended to provide access speeds comparable to local memory with the benefit that the bitmaps are available after a system crash.

Do I need Fast Write Cache?

The Fast Write Cache software is an application host-based write accelerator of disk I/O. It improves performance for transaction processing and delivers faster response times to user requests for data by reducing the frequency of disk I/O accesses. Writes are cached in non-volatile memory and then the cached data is destaged to disk at a later time. Fast Write Cache is installed on Solaris servers and consists of NVRAM (non-volatile memory) boards used as cache memory and storage cache management software. The use of the NVRAM boards prevents data loss during a system crash or power failure.

For optimum performance, use Sun StorEdge Instant Image 2.0.1 with Sun StorEdge Fast Write Cache on the target server. Utilizing Fast Write Cache reduces the latency of bitmap writes and allows the bitmaps to be preserved across a system crash. However, there are other options available for storing the bitmaps as discussed in the previous question.

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.