



# Sun StorEdge™ Instant Image 3.0 Configuration Guide

---

Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303-4900 U.S.A.  
650-960-1300

Part No. 806-7676-10  
June 2001, Revision A

Send comments about this document to: [docfeedback@sun.com](mailto: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, AnswerBook2, docs.sun.com, Sun StorEdge, 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, AnswerBook2, docs.sun.com, Sun StorEdge, 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.



# Table of Contents

---

<b>Preface</b>	<b>vii</b>
Using UNIX Commands	vii
Documentation Conventions	viii
Related Documentation	ix
Online Sun Documentation	ix
Sun Welcomes Your Comments	ix
<b>Sun StorEdge Instant Image 3.0 Configuration Guide</b>	<b>1</b>
System Requirements	2
Installation	2
Overview	3
Instant Image Volume Sets	5
Master Volume	5
Shadow Volume	5
Independent Shadow Volume	5
Dependent Shadow Volume	6
Compact Dependent Shadow Volume	6
Overflow Volumes	6
Bitmap Volume	7
Configuration Considerations for all Instant Image Volume Sets	7

When Should I Use a Dependent Shadow Volume Set?	7
When Should I Use an Independent Shadow Volume Set?	8
Where Should I Put the Bitmaps?	9
Multiple Shadow Volumes	10
Configuration Considerations for Multiple Shadow Volumes	10
When Should I Use Multiple Shadow Volumes?	10
When Should I Resynchronize the Shadows to the Master?	11
Exported Shadow Volume	12
Configuration Considerations for Exported Shadow Volumes	12
Why do I Want to Export a Shadow Volume?	12
Which Volumes Can I Not Export?	13
What Level of Redundancy Should I Use to Protect the Data?	13
Can I Mount an Exported Shadow Volume on the Original Host?	13
When Should I Resynchronize the Exported Shadow to the Master?	13
When Should I Resynchronize the Master With the Shadow?	13
How Does This Compare With the Sun StorEdge Network Data Replicator Solution?	14
Exporting, Importing and Joining Instant Image Shadow Volumes	14
Configuration Notes	17
Independent Shadow Volume	17
Dependent Shadow Volume	17
Compact Dependent Shadow Volume	18
Multiple Compact Dependent Shadow Volumes, with Overflow	18
Multiple Shadows of the Same Master	19
Copy or Update Independent Shadow	20
Copy or Update Dependent Shadow	20
Abort Copy or Update - Independent Shadow	20
Abort Copy or Update - Dependent Shadow	21

Suspend or Resume - Shutdown or Startup 21

Set or Get Copy Delay Parameters 22

Group for Atomic Copy or Update 22

Group for Ease of System Administration 22

**Glossary 23**



# Preface

---

This *Sun StorEdge Instant Image 3.0 Configuration Guide* describes the architecture and the configuration considerations for the use of Sun StorEdge™ Instant Image software.

The main section of this guide describes the architecture of and configuration considerations for Instant Image 3.0.

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 or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output.	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
<b>AaBbCc123</b>	What you type, when contrasted with on-screen computer output.	% <b>su</b> Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be <code>root</code> to do this. To delete a file, type <code>rm filename</code> .
[ ]	In syntax, brackets indicate that an argument is optional.	<code>scmadm [-d sec] [-r n[:n][,n]...] [-z]</code>
{ <i>arg</i> / <i>arg</i> }	In syntax, braces and pipes indicate that one of the arguments must be specified.	<code>iiadm -e {ind dep} <i>master shad...</i></code>

---



---

## Related Documentation

Application	Title	Part Number
man Pages	iiadm svadm dscfg pkgrm pkgask	N/A
Release Notes	<i>Sun StorEdge Instant Image 3.0 Release Notes</i>	806-7678
Installation and User	<i>Sun StorEdge Instant Image 3.0 System Administrator's Guide</i>	806-7677
	<i>Sun StorEdge Instant Image 3.0 Configuration Guide</i>	806-7676

## Online Sun Documentation

A broad selection of Sun system documentation is located at:

<http://www.sun.com/products-n-solutions/hardware/docs>

A complete set of Solaris documentation and many other titles are located 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](mailto:docfeedback@sun.com)

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



# Sun StorEdge Instant Image 3.0 Configuration Guide

---

The chapter discusses:

- “System Requirements” on page 2
- “Installation” on page 2
- “Overview” on page 3
- “Instant Image Volume Sets” on page 5
- “Multiple Shadow Volumes” on page 10
- “Exported Shadow Volume” on page 12
- “Configuration Notes” on page 17

Before reading this guide, you should have already installed the Instant Image software according to the instructions in the *Sun StorEdge Instant Image 3.0 Installation Guide*.



---

**Caution** – *Do not* install the Sun StorEdge Version 3.0 core and data services software on servers in a Sun Cluster 3.0 environment.

---

- The Version 3.0 software is *not* coexistent with the Sun Cluster 3.0 environment
- The Version 3.0 software is coexistent in the Sun Cluster 2.2 environment
- The Version 3.0 core and data services software is cluster aware in the Sun Cluster 3.0 Update 1 environment and provides high availability for the Sun StorEdge data services.

---

# System Requirements

---

**Software** Solaris™ 7 or 8 operating environment or a subsequent compatible version

Sun StorEdge data service core software

**Hardware** A CD-ROM drive connected to the host server where the Instant Image software is to be installed.

The Instant Image software is supported on server hosts using the Solaris operating environment. Hosts include but are not limited to:

- Sun Enterprise™ Server models 2x0 through 4x0
- Sun Enterprise Server models 3x00 through 10000

Disk space:

- The Instant Image software requires approximately 1 Mbyte
  - Supporting packages require approximately 3 Mbytes
  - Dual-ported devices are required for exported shadow volumes
- 

---

# Installation

Install the Sun StorEdge Instant Image 3.0 Software according to the procedures in the *Sun StorEdge Instant Image 3.0 Installation Guide* after carefully reading the *Sun StorEdge Instant Image 3.0 Release Notes* for last minute instructions and notices.

---

# Overview

The contents of a database often change from second to second, which can make it difficult to perform any analysis or backup. The Sun StorEdge Instant Image software enables you to take a *point-in-time snapshot* of the database for use in testing, backup, or analysis without affecting ongoing database transactions. Immediately after the point-in-time snapshot is enabled, the database is accessible to the primary business application and the snapshot is independently available to any other application. Changes to the database and to the snapshot are tracked so that the snapshot can be brought up to date at any time in the future. Transactions to the database are largely unaffected by analysis of the snapshot and snapshot operations affect database transactions only minimally.

The Sun StorEdge Instant Image software gives you instant access to your entire database for any purpose without degrading the performance of your business system. Imagine snipping one frame, the one currently being displayed, out of a running movie film, stopping the projector so briefly to make the cut that the audience barely notices; that is Instant Image.

The minimum unit of Instant Image operations is the volume set, which consists of a *master volume*, a *shadow volume*, a *bitmap volume*, and an optional *overflow volume*. The master volume is the database; the shadow volume is the point-in-time snapshot of the master volume; the bitmap volume tracks changes to both volumes; and the overflow volume accepts data sent to shadow volumes that are full. There are three types of shadow volumes: *independent shadow volume*, *dependent shadow volume*, and *compact dependent shadow volume*. All types of shadow volume are discussed in this chapter.

Both the business software and the analysis software can be running on one host, neither interfering with the other beyond sharing resources. If you need to, Instant Image allows you to *export* a shadow volume to another host so that the business transactions or the analysis can be carried out with virtually no impact on the main host. You might need to dedicate the main host to diagnostics, for instance, so you can offload the entire business function to another host with no observable affect to the outside world.

You can also include a single master volume in any number of volume sets, which is called *multi-shadowing* and gives you the ability to perform more than one other operation on a copy of the database. In theory, you could have the main business application running on the first host, an accounting application analyzing the day's transactions on an exported shadow volume at a separate host, and any number of other analysis or backup applications running on both. You can export more than one shadow volume so more than one other hosts can have access to point-in-time copies of the business database. Practically speaking, you are resource bound, but within limits Instant Image provides a flexible way to solve business problems.

As an example of the capability of Instant Image, consider a daily backup. Normally, you would have to quiesce the business application, make your backup, then return the business application to operation. This could take minutes or hours. Instead, with Instant Image, create a volume set and backup your database from the shadow volume. The business application continues with little interruption and the backup can be taken from the shadow volume at any time.

---

# Instant Image Volume Sets

An Instant Image volume set consists of a master volume, a shadow volume, a bitmap volume, and an optional overflow volume. From the Command-Line Interface (CLI), `iiadm`, you enable a volume set, which creates the association between the named volumes and makes the set available for Instant Image operations. Instant Image operations, including updating and copying volumes, are explained in the *Sun StorEdge Instant Image 3.0 System Administrator's Guide*.

## Master Volume

The master volume is the original database. The master volume gets write and read requests from the main business application and maintains the up-to-the-second business records.

## Shadow Volume

The shadow volume contains a point-in-time snapshot of the master volume. This snapshot represents the exact contents of the database at a defined time. Until the shadow is updated, any read of the shadow returns valid data for that point in time.

The shadow volume may also be written to, such that the volume's point-in-time data moves forward in time independently of the master volume's data.

The shadow can be of three types: independent, dependent, and compact dependent.

## Independent Shadow Volume

An independent shadow volume is a full point-in-time copy of the master volume. It can stand alone as a database referenced to a specific time. If the volume set is disabled, meaning that the association of the volumes is broken, the independent shadow volume still contains valid point-in-time data. It can be used for any purpose and can be updated from the master at any time, as long as you have not disabled it, without overly affecting business operations.

## Dependent Shadow Volume

A dependent shadow volume is not a full point-in-time copy of the master. In fact, it contains no data at all until the shadow is written to or the master is modified, at which time the master's old data is written to the shadow before the new data is written to the master volume. Reads of the shadow are served from the master unless the requested block has changed at the master, in which case the shadow serves the read request. In this way, the dependent shadow provides a point-in-time copy of the data without needing a full copy from master to shadow.

However, if a dependent shadow volume set is disabled, the shadow volume is marked as invalid data because it has no access to the master and cannot respond to reads.

The dependent shadow volume can be updated in an instant. All that is required is to clear the bitmaps that track changes to the shadow and master, which means that all reads to the shadow are served from the master until the master changes.

## Compact Dependent Shadow Volume

Often, a database is very stable with only a small percentage of its contents changing in any given period of time. A typical Instant Image volume set requires that the shadow volume be the same size or larger than the master volume; the same size is recommended. A great deal of the capacity of such a shadow volume would be virtually unused and would be an inefficient use of resources.

Instant Image permits you to enable a volume set with a shadow volume that is smaller than the master volume. You can save a lot of storage space by enabling compact dependent shadow volume sets. You should do this only when you know how much activity the database will sustain and can predict what volume is needed in the shadow.

By specifying a volume that is smaller than the specified master volume, Instant Image assumes that you are creating a compact dependent shadow volume set; there is no need to specify this.

## Overflow Volumes

Compact dependent shadow volumes can become filled up under unexpected circumstances. If this happens and no other provision has been made, the write fails and the data is marked as invalid.



Instant Image provides the ability to attach overflow volumes to compact dependent shadow volume sets to receive any writes that exceed the capacity of the shadow. Any number of compact dependent shadow volumes can be assigned to a single overflow volume as long as space is available.

Overflow volumes must be initialized before attaching them to compact dependent shadow volumes. The Instant Image CLI provides the command to do this.

## Bitmap Volume

As the master volume and the shadow volume change over time, each block of each volume is tracked in a bitmap, or *scoreboard*. The bitmap is used later for updating the independent shadow volume efficiently by copying only the changed blocks of data, as noted in the bitmap, not the whole volume.

Bitmaps are not used in this fashion to update dependent shadow volumes because all changed master data are written to the shadow before the master is written to, preserving the point-in-time snapshot data at the shadow volume. To update a dependent shadow volume, the bitmaps are simply cleared. The bitmap of a dependent shadow is used to decide whether to serve a write from the shadow or from the master.

## Configuration Considerations for all Instant Image Volume Sets

The Instant Image software provides a way to make immediately-available, point-in-time copies of a database while the database is active. These copies are called shadows and can be created, updated, and made available in a number of ways.

You can enable two primary types of shadow volume sets: dependent and independent. Whether you want to use one type or the other depends on a number of issues such as redundancy, availability, and the expected use of the shadow volume.

### When Should I Use a Dependent Shadow Volume Set?

Dependent volume sets can be of two types: a dependent shadow volume set or a compact dependent shadow volume set. The difference between these is that the compact dependent shadow volume set has only enough space allocated to the shadow volume to contain the anticipated data writes from the master volume, not the entire contents of the master volume. If you can predict the amount of data writes expected between updates, you can use compact dependent shadow volumes

and reduce the volume required for shadow volumes. If you cannot accurately predict the required volume, but you still want to use a compact dependent shadow volume, you can attach an overflow volume to the volume set. If the data written to the shadow exceeds the compact shadow volume's capacity, it will be directed to the overflow volume. If you cannot predict the data writes to the shadow volume or you do not know how often updates will happen, you should not use a compact dependent shadow volume.

If you use the dependent shadow volume for another application, any reads to the shadow will be served from the master volume unless the block of data requested has been changed at the master, in which case the request is served from the shadow. This means that the master volume's overhead will consist of its own requirements plus some portion of those of the application that is using the shadow. If the master volume is subject to a heavy load from its primary application, do not use a dependent shadow volume.

One use for a dependent shadow volume set is with a backup application. Usually, you have to quiesce the master volume to allow a backup application to make a full copy of the data. Instead, enable a dependent shadow volume set and let the backup application make its copy from the shadow volume. In this way, the backup starts at once using the shadow volume, but the master volume does not need to be quiesced for the duration of the backup. In essence, you are using Instant Image as a window into a point-in-time snapshot of your database.

An excellent use for compact dependent shadow volumes is with a database that changes little over time. The compact shadow volume is sized to accept the expected changes, which saves a lot of storage space. By contrast, a standard dependent shadow volume would represent a storage volume equal to the master volume, but that contained virtually no data.

One advantage of dependent shadow volumes is that they are available as dependent shadows immediately upon being enabled, whereas an independent shadow volume, while available, is tagged as a dependent shadow volume until a full copy is complete.

## When Should I Use an Independent Shadow Volume Set?

There is only one type of independent shadow volume set. You cannot enable an independent shadow volume set with a compact shadow and you cannot attach an overflow volume to it. The shadow volume is required to be the same size as the master, or larger, and there is no purpose to attaching an overflow volume even if you could.

Independent shadow volumes are independent point-in-time copies of the master volume. When you enable an independent volume set, a full volume copy is executed. The shadow volume is considered to be a dependent copy of the master

until the full copy is complete. This is more time consuming initially than enabling a dependent shadow volume set, but after the copy is complete, the independent shadow volume needs no support from the master volume.

An independent shadow volume set can be disabled after the full copy, which results in two complete, usable copies of the same data. By contrast, if you disable a dependent shadow volume set, the shadow volume is marked as being unusable.

Using an independent shadow volume, you can perform testing without impacting the performance of the primary application. If you use a dependent shadow for this same task, the primary application will be degraded to whatever extent I/O requests have to be served.

The master volume can be restored from the shadow volume in an independent shadow volume set, at least back to the point of the last update.

Use independent shadow volume sets when storage space is plentiful and when a full copy of the master volume's data is needed for testing. The full copy will not require the master volume to be quiesced during the copy so the business data flow can continue.

Note that the Instant Image software includes a way to perform *throttled* copies. A throttled copy is one in which you program the maximum amount of data transferred before pausing a programmed length of time for other applications to use the system. In this way, a copy proceeds to completion, but still allows other processes to continue.

## Where Should I Put the Bitmaps?

The Instant Image software's bitmaps must be kept on a volume, not in a file; must be located on a volume that is mounted at bootup; and should be mirrored. This volume should be used exclusively for Instant Image bitmaps.

---

# Multiple Shadow Volumes

You can enable any number of Instant Image volume sets of any kind including the same master volume in all of them. This has the effect of giving the master volume multiple shadows. Each shadow volume is treated exactly as it would be if it were not one of multiple shadows. Independent shadow volumes require full copies upon being enabled and all dependent shadow volumes get a write from the master before new data is written into the master.

You can update the master volume from any of the shadow volumes and you can update the shadow volumes from the master at any time.

Each shadow volume is independent of all the others and can be used for any applicable purpose.

## Configuration Considerations for Multiple Shadow Volumes

### When Should I Use Multiple Shadow Volumes?

Multiple shadow volumes are useful for evaluating software applications using current copies of data. You can, for instance, evaluate software for data analysis using your current business data while having no impact at all on the master volume of that data. The master volume and another of its multiple shadow volumes can be the main business data storage and can continue business operations unaffected by your analysis.

The shadow volumes can be of any type: independent, compact dependent, or dependent. If any shadow is a dependent shadow or a compact dependent shadow, it will get a copy-on-write from the master volume, which maintains its point-in-time snapshot validity.

Each of the multiple shadows of a master volume performs like any other shadow volume. It can be updated from the master, it can be used to update the master, and it can be exported.

In short, multiple shadow volumes enable any reasonable number of applications to have access to independent point-in-time copies of business data.

## When Should I Resynchronize the Shadows to the Master?

Any application that is using one of more than one shadow volumes of the same master has access to valid point-in-time data. If you need the point-in-time data to be updated for one or more of the shadows, they must be updated from the master in the usual manner.

If you plan to export an independent shadow volume, you must perform a full master-to-shadow copy first. Remember that a shadow volume must reside on a dual-ported device for you to be able to export it.

---

# Exported Shadow Volume

You can use Instant Image to export a shadow volume so that another host running Instant Image can *import* it. The other host can use the shadow volume for any purpose, such as data analysis or backup. The imported shadow volume must be imported by specifying a bitmap for tracking changes.

A main use for an exported shadow is to transfer the business transactions to another host in order to offload the main host. Export the shadow after a required full master-to-shadow copy, import it at a second Instant Image host, enable a volume set at the second host with the imported shadow volume as the master, then switch the business transactions over to this second host.

Once the second host is finished with the shadow volume, it disables any volume sets that it formed with it. This releases the shadow volume to be re-joined to its master, bringing with it any changes made by using the bitmap from the second host for the join.

You can export only an independent shadow volume, not a dependent shadow, and only immediately following a full copy. The master can continue to process data while the shadow is exported, tracking changes in its bitmap.

## Configuration Considerations for Exported Shadow Volumes

### Why do I Want to Export a Shadow Volume?

Exporting can be useful for offloading Instant Image from one machine to another. By exporting an independent shadow volume after a full update, then importing it to another Instant Image host, the second host can assume the business activities seamlessly, thus relieving the first host of all tasks but I/O request servicing. You can later join the shadow to its original master and include changes made to the shadow and to the master while the shadow was exported, which transfers the business operation back to the first host seamlessly and with up-to-date data.

You can export the shadow volume for tasks that are not related to Instant Image; for example, when the data it contains is needed for accounting purposes and must be stable throughout the operation. The shadow volume, which had a full master-to-shadow copy just prior to being exported, contains the business data up to the point in time when the shadow was exported. The master volume continues to accept new

transactions while the shadow is being analyzed on another host. When the analysis is complete, the shadow can be joined to the master, updated to reflect any changes to the master, and used as usual in its original Instant Image volume set.

## Which Volumes Can I Not Export?

You can export shadow volumes, but not `/` or `/usr`.

## What Level of Redundancy Should I Use to Protect the Data?

You should use the same level of redundancy for the exported shadow volume that you use for the master and the bitmap volumes.

## Can I Mount an Exported Shadow Volume on the Original Host?

No. While the volume is exported, you cannot mount the volume on the first host.

## When Should I Resynchronize the Exported Shadow to the Master?

The shadow volume can be resynchronized with the master when your purpose for exporting the shadow volume is satisfied or when you need up-to-date data in the exported shadow. You should allow any activities associated with the shadow volume to complete prior to beginning the resynchronization process. To resynchronize the shadow volume, disable any Instant Image volume sets that you created on the second host, quiesce operations on the master volume, join the shadow to its master on the first host, start a master-to-shadow copy, and allow business operations to continue. As soon as the master-to-shadow copy is complete, you can export the shadow volume again.

## When Should I Resynchronize the Master With the Shadow?

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

## 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. Remote can mean a continent away and the connecting network uses the IP protocol. Instant Image volume sets are local and are maintained over Fibre Channel Arbitrated Loop (FC-AL).

SNDR is designed to be active during normal application access to the data volumes, and will continually replicate the data to the remote site. Shadow volumes in Instant Image are point-in-time copies that match the master only after a copy or an update.

SNDR can be employed in disaster recovery and business continuance strategies to provide redundant storage of critical information across physically separate sites. Instant Image is not a backup utility, although it can help to implement an efficient one, and would not usually be used for disaster recovery.

Because large databases can require excessive network time over IP networks, tapes are often hand-carried to the remote site, loaded onto a host, and enabled using SNDR without having to copy over the network. Instant Image does not require this because it is locally based and maintained over a high-speed fibre channel network.

## Exporting, Importing and Joining Instant Image Shadow Volumes

The Instant Image functionality implemented via the three `iiadm` options of `-E` (export), `-I` (import) and `-J` (join) allow for a dual-ported shadow volume to be deported back and forth between a primary and secondary host while under Instant Image control. This capability allows shadow volume processing by its associated applications to be off-loaded to a secondary host without impacting the primary host's master volume or its associated applications.

Retaining Instant Image control of the shadow volume while on the secondary host maintains bitmap integrity to facilitate fast resynchronization via copy or update processing at a later time.

While the shadow volume is deported to a secondary node, read and write access to the master volume is still tracked by Instant Image. Read and write access to the shadow volume on the secondary host is also tracked by Instant Image using a second bitmap.

Once secondary host processing has completed, the shadow volume and second bitmap can be deported from the secondary host back to the primary host and rejoined with the master volume such that the master, shadow, and bitmap



consistency is reconstructed. After completing the join processing, the Instant Image set is now in the same state as it would have been if the secondary host processing to the shadow volume, had been performed on the primary host.

The following is a brief outline of how to export/import/join a shadow volume.

- **Create Instant Image shadow volume on primary node and start application(s) using shadow volume.**

Primary Host	Secondary Host	Comments
<i>master</i> volume		existing <i>master</i> volume, mounted, valid data
create <i>shadow</i> and <i>bitmap1</i> volumes		should be the same redundancy (RAID) as the <i>master</i>
<i>shadow</i> volume		same size as <i>master</i> on a dual-ported device
<i>bitmap1</i> volume		size based on <i>master</i> volume set
enable independent shadow set		<code>iiadm -e ind <i>master shadow bitmap</i></code>
mount <i>shadow</i> volume		<code>mount <i>shadow mount_point</i></code>
start application using <i>shadow</i>		

- **Switch application and shadow volume to secondary host.**

Primary Host	Secondary Host	Comments
stop application using <i>shadow</i>		existing <i>master</i> volume, mounted, valid data
unmount <i>shadow</i>		<code>umount <i>mount_point</i></code>
export <i>shadow</i> volume		<code>iiadm -E <i>shadow</i></code>
SV disable		<code>svadm -d <i>shadow</i></code>
create <i>bitmap2</i>		same size as <i>bitmap1</i> , on dual-port device
copy <i>bitmap1</i> to <i>bitmap2</i>		<code>cp <i>bitmap1 bitmap2</i></code>
deport <i>shadow/bitmap2</i>		required for VxVM or SDS(SLVM)
	import <i>shadow/bitmap2</i>	required for VxVM or SDS(SLVM)
	SV enable <i>shadow</i> volume	<code>svadm -e <i>shadow</i></code>

Primary Host	Secondary Host	Comments
	import <i>shadow</i> volume	iiadm -I <i>shadow bitmap2</i>
	mount <i>shadow</i> volume	mount <i>shadow mount_point</i>
	start application using <i>shadow</i>	

● **Switch application and shadow volume to primary host.**

Primary Host	Secondary Host	Comments
	stop application using <i>shadow</i>	
	unmount <i>shadow</i> volume	umount <i>mount_point</i>
	suspend <i>shadow</i> volume	iiadm -s <i>shadow</i>
	disable <i>shadow</i> volume	iiadm -d <i>shadow</i>
	SV disable <i>shadow</i> volume	svadm -d <i>shadow</i>
	deport <i>shadow/bitmap2</i>	required for VxVM or SDS(SLVM)
import <i>shadow/bitmap2</i>		required for VxVM or SDS(SLVM)
SV enable <i>shadow</i> volume		svadm -e <i>shadow</i>
join <i>shadow</i> volume		iiadm -J <i>shadow bitmap2</i>
mount <i>shadow</i> volume		mount <i>shadow mount_point</i>
start application using <i>shadow</i>		

---

# Configuration Notes

## Independent Shadow Volume

```
iiadm -e ind master shadow bitmap
```

---

**Note** – Shadow size is the same as or greater than the master.

---

An independent shadow volume, initiates a (throttled) full-copy of the master to shadow volume. While the background copy is in progress (see `iiadm -` or `iiadm -w` to wait if scripting), the shadow volume is considered to be a dependent shadow. Once the full copy is complete, the shadow volume can be disabled from the volume set and there is no longer any impact to the master volume. Unless there is a reason, you should leave the volume set enabled so that fast-resynchronization can be realized.

## Dependent Shadow Volume

```
iiadm -e dep master shadow bitmap
```

---

**Note** – Shadow size is the same as or greater than the master.

---

A dependent shadow volume, does not require a full-copy, since point-in-time (unmodified) data is obtained from the master and modified data from the shadow. The shadow volume is always dependent on the master. The impact to the master volume is continuous for all unmodified data access. If the volume is disabled, the shadow volume will be invalidated. Since the shadow volume can not be disabled, fast-resynchronization is always an option.

# Compact Dependent Shadow Volume

`iiadm -e dep master shadow bitmap`

---

**Note** – Shadow size is less than the master.

---

A compact dependent shadow volume has the same characteristics as a dependent shadow volume with the additional benefit of requiring less storage. If the amount of change to the master or shadow volumes over the duration that the point-in-time is required is less than the size of the master volume (typical), the unused space in a dependent shadow can be reclaimed by using a smaller compact dependent shadow.

Prior to switching over from an independent or dependent full sized shadow, determine the average bitmap usage and set the size of your compact dependent shadow to this size plus a comfortable margin.

---

**Note** – If the space allocated to a compact dependent shadow is less than the amount of change, the volume will be invalidated and taken offline. The following section on overflow volumes addresses this issue.

---

## Multiple Compact Dependent Shadow Volumes, with Overflow

`iiadm -O overflow`

(Initialize overflow)

`iiadm -L`

(List all overflow volumes)

`iiadm -Q overflow`

(Query status overflow volume)

`iiadm -A overflow shadow`

(attach overflow to shadow volume)

`iiadm -D shadow`

(detach overflow from shadow volume)

One or more compact dependent shadows can share one or more overflow volumes. These shared overflow volumes are used to store above average changes to one or more of the compact dependent shadows which have been attached to the overflow.

Prior to using an overflow volume for the first time, it must be initialized. Once initialized, it can be attached to one or more compact dependent shadows. There can be more than one overflow volume.

Prior to detaching an overflow volume, the compact dependent shadow must not be in the overflowed state. If it is, a new point-in-time (via copy or update) must be taken, so that the overflow data is reclaimed from the overflow volume. The `iiadm -L` command option lists the status of the overflow volume specified.

## Multiple Shadows of the Same Master

`iiadm -e {dep|ind} master shadow1 bitmap1`

`iiadm -e {dep|ind} master shadow2 bitmap2`

Instant Image supports the capability to allow multiple shadow volumes to be created against a single master. The individual shadow volume can be any type of shadow volume (independent, dependent, compact dependent, or compact dependent with overflow). Each shadow volume can have its own unique point-in-time, or it can be moved into a group and have the same point-in-time as other shadow volumes in the group.

This capability allows a single master and its associated shadows to be used for various concurrent purposes. A new point-in-time can be taken at any time, without impacting the other shadow volumes. Note that each dependent shadow volume may be involved in a copy-before-write I/O operation if the data on the master is updated. If you have 10 compact shadows using the same master volume, a write to the master may incur 10 writes to the shadow plus the one write to the master for each I/O operation.

---

**Note** – For the following update or copy operations, the target of the operation must be unmounted for the instant the the command is active.

---

## Copy or Update Independent Shadow

`iiadm -{cu} {ms} shadow`

Copy or Update independent shadow volume.

A copy or update (fast-resynchronization) creates a new point-in-time of either the master-to-shadow, or shadow-to-master volumes. While the background copy or update is in progress, the shadow is considered a dependent shadow.

## Copy or Update Dependent Shadow

Includes copy or update compact dependent shadow and copy or update compact dependent shadow with overflow.

`iiadm -{cu} {ms} shadow`

Copy or Update dependent shadow volume)

For all dependent shadows, a copy operation is turned into an update operation. The update (fast-resynchronization) creates a new point-in-time of either the master-to-shadow or shadow-to-master volumes. For master-to-shadow copies or updates the operation requires only bitmap maintenance. For shadow-to-master the operation requires copying the changed data in addition to bitmap maintenance. Since the operation of shadow-to-master affects the master volume, an interactive query is generated asking whether to proceed or not. You can override this message (for scripting) with the `-n` option. If an overflow volume is associated with a compact dependent shadow and data has overflowed into it, the copy or update removes all overflow data.

## Abort Copy or Update - Independent Shadow

`iiadm -a shadow`

The master and shadow volumes of an independent volume set are complete, consistent volumes that can be used independently after you disable the volume set. Prior to allowing a disable operation to complete, both the master and shadow volumes must be left in a consistent state.

If an Instant Image copy or update operation is in progress, you will be unable to disable the volume set using the `iiadm -d` command option. If you were to abort the copy or update operation using the `iiadm -a` command option, the volumes would still be inconsistent, and again a disable command option would not be allowed.

If you must disable an independent volume set while a copy or update is in progress, and you are not concerned with the consistency of the target volume, first suspend the volume set with the `iiadm -s` command option, then disable the volume set with the `iiadm -d` command option. The target (master or shadow) volume of the copy or update operation is left in an inconsistent state.

If you must disable an independent volume set while a copy or update is in progress, and you are concerned with the consistency of the target volume, reissue the copy or update operation, let it complete (`iiadm -w`), then issue the disable option.

## Abort Copy or Update - Dependent Shadow

Due to the dependency on the master volume for the consistency of the data on the shadow volume, the Instant Image administration utility will not allow an abort of the copy or update of a dependent shadow. Doing so would render both the master and shadow in an inconsistent state.

## Suspend or Resume - Shutdown or Startup

```
iiadm -s shadow
```

```
iiadm -r shadow
```

The Instant Image suspend and resume commands will stop Instant Image processing of the volume set. Prior to suspending a shadow volume, it must be unmounted. During the time that a shadow set is disabled, changes to the master volume are not under Instant Image control, thus the shadow volume's point-in-time will become inconsistent. The combination of suspend and resume are used by the StorEdge Data Services during system shutdown and reboot so that a consistent view of an Instant Image volume set can be maintained.

## Set or Get Copy Delay Parameters

During an Instant Image copy or update operation, the `iiadm` utility completes in the foreground, but the full-copy or fast-resynchronization occurs in the background. Depending on system requirements, a system administrator may wish to adjust the amount of time and the system impact in which a this background processing occurs. The two parameters; delay and units, provide full control over the impact that the copy or update processes have on a per volume basis.

## Group for Atomic Copy or Update

```
iiadm -g group -{cu} {ms} shadow
```

The copy or update command in combination with the `-g group` option allows a single point-in-time to be taken over the entire named group of volume sets. This allows an application that consists of multiple volumes to have the exact same point-in-time.

## Group for Ease of System Administration

```
iiadm -g group [options]
```

In addition to the grouping for copy or update, most Instant Image options operate at the group level to assist in ease of Instant Image volume management.

For instance, to move multiple volumes into a group:

```
iiadm -g group -m shadow shadow2 ... shadowN
```

The Instant Image utility allows one or more previously configured Instant Image sets to be moved into a group in a single command invocation.



# Glossary

---

<b>bitmap volume</b>	The file used to track the location of data changes on the shadow and master volumes for use in later resynchronization.
<b>compact dependent shadow volume</b>	As form of shadow volume that is smaller than its associated master volume, enabling efficient use of available storage capacity where the data writes to the master are limited and quantifiable.
<b>dependent shadow volume</b>	A volume-dependent, point-in-time snapshot. When you enable an Instant Image volume pair, you must specify the shadow volume as dependent or as independent. For a dependent shadow, Instant Image does not perform a full volume copy. The dependent shadow volume relies on the master for all unmodified data blocks. Disabling a dependent shadow volume causes it to be invalidated.
<b>dual-ported device</b>	A dual-ported disk has two ports, making it possible to connect the disk through two cables to two different controllers. Either controller, at any one time, can transmit data to and receive data from that one disk.
<b>fast resynchronization</b>	Occurs when Instant Image updates volumes by copying all 32Kbyte segments flagged as different between the master and shadow volumes. The master can be used to update the master, and the master can be used to update the shadow, depending on which volume is being updated. Updating does not copy all volume data, only changed data, and is quicker than a full volume copy.
<b>independent shadow volume</b>	A volume-independent, physically-separate full volume copy. When you enable an Instant Image volume pair, you must specify the shadow volume as dependent or as independent. For an independent shadow, Instant Image performs a full volume copy to the shadow, which remains a valid copy after Instant Image is disabled.

<b>invalidate</b>	To make invalid. After disabling a dependent volume, Instant Image clears the first 64Kbytes of data in the volume to prevent users from accessing invalid or inconsistent data.
<b>I/O group</b>	A group to which Instant Image volume sets are assigned for ease of management. When a command is executed against an I/O group, it is executed against all included volume sets.
<b>master volume</b>	The volume that contains the original data that is copied to the shadow volume.
<b>multiple shadow volumes</b>	The result when more than one shadow volume is attached to a single master volume so that more than one application can use the master volume in a volume set.
<b>overflow volume</b>	A volume that is attached to a compact dependent shadow volume to accomodate excess data and prevent a failure. One overflow volume can be attached to more than one compact dependent shadow volume.
<b>point-in-time snapshot</b>	A copy or image of volume data captured or copied at a particular moment.
<b>quiesce</b>	To stop operations of a database application momentarily so that Instant Image update and copy operations will not encounter data errors. After the update or copy command is issued, the database application can be restarted.
<b>resynchronization</b>	See <i>fast resynchronization</i> .
<b>segment</b>	A 32Kbyte track of data.
<b>shadow volume</b>	The volume containing a point-in-time snapshot of data from the master volume.
<b>Sun Cluster tag</b>	A specific Sun Cluster name that is used to reference that cluster for action by an Instant Image command.
<b>throttling a copy</b>	To program the maximum number of chunks to be copied before executing a programmed delay to prevent an Instant Image copy from monopolizing system resources.
<b>volume set</b>	A group of three volumes: the volume pair, which comprises the master volume and the shadow volume, and a bitmap volume. A volume set that includes a compact dependent shadow volume can also include an optional overflow volume to receive data when the dependent short-shadow volume is full.
<b>volume pair</b>	A volume pair includes two volumes: the master volume and its associated shadow volume. Volume pairs always have bitmaps associated with them. The combination of all three volumes is called a volume set.