



# Sun StorEdge™ Availability Suite Software – Maximizing Read-Only Volume Access

---

A Best Practice

Sun Microsystems, Inc.  
4150 Network Circle  
Santa Clara, CA 95054 U.S.A.  
650-960-1300

Part No. 816-7179-10  
June 2002, Revision A

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this 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 in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun StorEdge, Sun Enterprise, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in 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.

Use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth in the Sun Microsystems, Inc. license agreements and as provided in DFARS 227.7202-1(a) and 227.7202-3(a) (1995), DFARS 252.227-7013(c)(1)(ii) (Oct. 1998), FAR 12.212(a) (1995), FAR 52.227-19, or FAR 52.227-14 (ALT III), as applicable.

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 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuelle relatants à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

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

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun StorEdge, Sun Enterprise, et Solaris sont des marques de fabrique ou des marques déposées 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 licenciées 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'ÉTAT" 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

---

Point-in-Time Copy Software Performance Impact	2
Reducing the Impact on Volumes	2
Results of Performing the Procedure	2
▼ To Maximize Read-Only Volume Access	3
Verifying the Data	5
Hardware Configuration	5
Software Configuration	5
Data Verification Procedure	6
▼ To Stage the Data Verification	6
▼ To Validate the Data	7
▼ To Tear Down the Test Setup	8
Code Example	10



# Sun StorEdge Availability Suite Software – Maximizing Read-Only Volume Access

---

The Sun StorEdge™ Availability Suite Software includes the point-in-time copy software and the remote mirror software. These software features affect the performance of storage volumes under their control.

This best practice document provides a procedure for limiting the impact of the Sun StorEdge Availability Suite Software on a volume in a storage configuration that includes the remote mirror software and the point-in-time copy software. Included in this document is an operational procedure for limiting the impact of the point-in-time copy software, the details of how data integrity was verified after running the procedure, and a code example.

This document is also applicable to installations that include the Sun StorEdge Instant Image 3.0.x software (corresponds to the point-in-time copy software) and the Sun StorEdge Network Data Replicator (SNDR software) 3.0.x software (corresponds to the remote mirror software).

---

**Note** – This best practice is applicable *only if* the application uses the shadow volume in a read-only fashion.

---

---

# Point-in-Time Copy Software Performance Impact

The point-in-time copy software introduces a performance impact on a volume under its control. At times, this impact can be significant depending on the type of I/O accessing the volume and on the server configuration.

## Reducing the Impact on Volumes

The procedure described in this best practice helps to reduce the impact of the point-in-time copy software and the remote mirror software on application access to volumes under their control. To limit the impact of the point-in-time copy software and the remote mirror software on a volume during read-only access, remove the volume from the SV layer. Removing a volume from SV control takes the point-in-time copy software and the remote mirror software out of the I/O path for that volume, which removes the overhead required to track the access to the underlying volumes from the user application. Performance returns to levels seen before the introduction of the point-in-time copy software and the remote mirror software. Do not remove the bitmap volume from SV control.

## Results of Performing the Procedure

Performance tests used `vdbench 2.0` to generate data for read I/O done to the shadow volume in a point-in-time copy volume set. The tests used both random and sequential I/O patterns using a 4 Kbyte I/O size. The shadow volume was the target for all I/O.

With the shadow volume under SV control, the tests showed a substantial degradation in performance.

Removing the shadow volume from SV control returned performance to the same level observed before enabling the point-in-time copy software and the remote mirror software.

## ▼ To Maximize Read-Only Volume Access

This procedure makes the following assumptions about the point-in-time copy software and the remote mirror software setup:

- Applications that access the shadow volume have read-only access while the volume is removed from SV control.
- If a file system is part of the application, the file system on the shadow volume is mounted read-only, and no write operations are performed to the volume while it is removed from SV control.
- Point-in-time copy software shadow volume sets are of the independent shadow type.
- The Sun StorEdge Availability Suite Software is being used with the latest patches.

**1. If the shadow volume is mounted or is being accessed, stop application access and unmount the shadow volume.**

**2. Put the remote mirror volume sets into logging mode.**

**3. Choose one of the following operations:**

- Quiesce the master volume application and flush system caches
- Stop master volume application access and unmount the master volume
- Switch the database that is accessing the master volume into hot-backup mode

**4. Update the shadow volume.**

The update command returns immediately, indicating that the background copy is set up and is running. You do not need to wait for the background copy to complete before continuing to the next step.

**5. Choose one of the following operations:**

- Resume master volume application
- Mount master volume and resume master volume application
- Take database out of hot-backup mode

**6. Wait for the point-in-time copy software update of the shadow volume(s) to complete.**

**7. Disable the shadow volume from the SV layer.**

The point-in-time copy software and the remote mirror software are removed from the I/O path to the shadow volume. I/O to the master volume is still under the point-in-time copy software control.

8. Make the shadow volume available to the shadow volume application with read-only access.

```
# mount -o ro block-device mount-point
```

9. Run the shadow volume application against the shadow volume.
10. Return the shadow volume to the state it was in before [Step 7](#).  
If a file system was mounted in [Step 8](#), unmount it.
11. Enable the shadow volume in the SV layer.
12. Put the remote mirror volume sets into replicating mode.
13. Wait until the remote mirror volume sets have completed the sync and are in replicating mode.
14. Return the remote mirror volume sets to logging mode if required, or leave them in replicating mode until the next point-in-time copy software update needs to be performed.



---

# Verifying the Data

This section includes the data verification procedures used to measure the integrity of the data and the performance of the read-only shadow when removed from SV control. The hardware and software configurations are also included.

## Hardware Configuration

Host A:

- Sun Enterprise™ E4500 server with two 400 MHz CPUs and 2048 Mbyte memory
- Sun StorEdge A5200 storage arrays with 9-Gbyte drives for storage

Host B:

- Ultra 60 with two 450 MHz CPUs and 512 Mbyte memory
- Sun StorEdge D1000 array with 9-Gbyte drives for storage

## Software Configuration

Host A:

- Solaris™ 8 (02/02) Operating Environment
- Veritas volume manager 3.2
- Veritas file system 3.4
- Sun StorEdge Instant Image 3.0, SNDR software 3.0, and supporting core software with patches: 111945-05, 111946-05, 111947-03, and 111948-05
- Bitmap volumes residing on disk

Host B:

- Solaris 8 (02/02) Operating Environment
- Veritas volume manager 3.2
- Veritas file system 3.4
- Sun StorEdge Instant Image 3.0, SNDR software 3.0, and supporting core software with patches: 111945-05, 111946-05, 111947-03, and 111948-05
- Bitmap volumes residing on disk

# Data Verification Procedure

The following procedures measure performance and verify data integrity when using the procedure in [“To Maximize Read-Only Volume Access” on page 3](#). The tool, `dex`, verifies data integrity by reading and writing data patterns to block or character devices. The tool, `vdbench`, measures performance.

There are three procedures in the data verification:

- [“To Stage the Data Verification” on page 6](#), in which the test setup is prepared.
- [“To Validate the Data” on page 7](#), in which the procedure is validated.
- [“To Tear Down the Test Setup” on page 8](#), in which the test setup is removed.

## ▼ To Stage the Data Verification

1. **Create two 100-Mbyte volumes on `hostA`, one to be the point-in-time copy master volume, one to be the point-in-time copy shadow volume.**
2. **Create two 1-Mbyte volumes on `hostA` to be used as bitmap volumes, one for the point-in-time copy software and one for the remote mirror software.**
3. **Create a 100-Mbyte volume on `hostB`.**

This volume is used as the secondary volume in the remote mirror volume set.

4. **Create a 1-Mbyte volume on `hostB` to be used as the bitmap volume for the remote mirror volume set.**
5. **Write a `vxfs` file system to the volume to be used as the point-in-time copy master volume.**
6. **Mount the new file system and create an 85-Mbyte file.**

```
# mkfile 85m mount-point/tmpFile
```

7. **Unmount the file system.**
8. **Enable a point-in-time copy independent shadow volume set on `hostA` using the volume with the file system as the master volume.**
9. **Enable a remote mirror volume set with the point-in-time copy shadow volume on `hostA` as the remote mirror primary volume on `hostA`, and enable the volume created on `hostB` as the remote mirror secondary volume on `hostB`.**

For the rest of this process, the remote mirror primary volume and the point-in-time copy software shadow volume are referred to as the `SHADOW_VOLUME`. These two volumes are the same physical volume.

10. **Put the remote mirror software into replicating mode.**

11. When the remote mirror sync has completed, put the remote mirror software into logging mode.
12. Mount the point-in-time copy master volume on `hostA`.
13. Use `dex` to write an I/O pattern to the mounted point-in-time copy master volume.
14. Mount the `SHADOW_VOLUME` on `hostA`.
15. Use `dex` to verify that the I/O pattern written to the point-in-time copy master volume is not on the `SHADOW_VOLUME`.
16. Mount the remote mirror secondary volume on `hostB`.
17. Use `dex` to verify the I/O pattern written to the point-in-time copy master volume on `hostA` is not on the remote mirror secondary volume.
18. Unmount all point-in-time copy and remote mirror volumes on both hosts.

The point-in-time copy master volume now contains unique data. This data is not present on the `SHADOW_VOLUME` (the point-in-time shadow volume on `hostA` and the remote mirror secondary volume on `hostB`).

## ▼ To Validate the Data

1. Update the `SHADOW_VOLUME`.
2. When the update command returns:

This step shows that you can still access the master volume while the copy completes and that the original data pattern is no longer on the master volume.

  - a. Mount the point-in-time copy master volume on `hostA`.
  - b. Use `dex` to write an I/O pattern to the point-in-time copy master volume.

This pattern should be different than the pattern written in [Step 12 of “To Stage the Data Verification” on page 6](#).
  - c. When step 2b completes, use `dex` to verify that the pattern written to the point-in-time copy master volume in [Step 12 of “To Stage the Data Verification” on page 6](#) is no longer present on the point-in-time copy master volume.
3. Mount the `SHADOW_VOLUME` on `hostA`.
4. Use `dex` to verify the pattern written to the point-in-time copy master volume in [Step 12 of “To Stage the Data Verification” on page 6](#) is now copied over to the `SHADOW_VOLUME`.
5. Unmount the `SHADOW_VOLUME`.
6. If there is a point-in-time copy software copy in progress, wait for it to complete.

7. Remove the SHADOW\_VOLUME on hostA from SV control.

```
# svadm -d SHADOW_VOLUME
```

8. Mount the SHADOW\_VOLUME read-only on hostA.

```
# mount -F vxfs -o ro SHADOW_VOLUME mount-point
```

9. Using vdbench, run a series of read performance jobs against the SHADOW\_VOLUME on hostA.

10. Unmount the SHADOW\_VOLUME on hostA.

11. Enable the SHADOW\_VOLUME on hostA in SV.

```
# svadm -e SHADOW_VOLUME
```

12. Mount the SHADOW\_VOLUME on hostA.

13. Use dex to verify that the I/O pattern written to the point-in-time copy master volume in [Step 9](#) of staging is still valid.

14. Unmount the SHADOW\_VOLUME on hostA.

15. Put the remote mirror volume set into replicating mode.

16. Put the remote mirror software into logging mode. This provides access to the remote mirror secondary volume on hostB.

17. Mount the secondary remote mirror volume on hostB.

18. Use dex to verify that the I/O pattern written to the point-in-time copy master volume on hostA during [Step 12](#) of “To Stage the Data Verification” on [page 6](#) is now on the remote mirror secondary volume on hostB.

The original I/O pattern is now on the SHADOW\_VOLUME (the point-in-time copy shadow volume on hostA and the remote mirror secondary volume on hostB). A different I/O pattern should be on the point-in-time copy master volume on hostA.

## ▼ To Tear Down the Test Setup

1. Delete the remote mirror volume sets on both the primary and the secondary hosts.
2. Delete the point-in-time copy volume sets.

### **3. Disable volumes from SV.**

If you are using the point-in-time copy software and the remote mirror software, only the volumes disabled from SV control during the above procedure need to be disabled from SV. All other volumes are disabled from SV automatically.

---

# Code Example

```
[root@hostA:/]
# vxassist -g testdg make ii_master 100m

[root@hostA:/]
# vxassist -g testdg make ii_shadow 100m

[root@hostA:/]
# vxassist -g testdg make ii_bitmap 1m

[root@hostA:/]
# vxassist -g testdg make sndr_bitmap 1m

[root@hostB:/]
# vxassist -g testdg make sndr_vol 100m

[root@hostB:/]
# vxassist -g testdg make sndr_bm 1m

[root@hostA:/]
# mkfs -F vxfs /dev/vx/rdsk/testdg/ii_master
    version 4 layout
    204800 sectors, 102400 blocks of size 1024, log size 1024 blocks
    unlimited inodes, largefiles not supported
    102400 data blocks, 101280 free data blocks
    4 allocation units of 32768 blocks, 32768 data blocks
    last allocation unit has 4096 data blocks

[root@hostA:/]
# mount -F vxfs /dev/vx/dsk/testdg/ii_master /master

[root@hostA:/]
```

```

# mkfile 85m /master/tmpFile

[root@hostA:/]
# umount /master

[root@hostA:/]
# iiadm -e ind /dev/vx/rdisk/testdg/ii_master
/dev/vx/rdisk/testdg/ii_shadow /dev/vx/rdisk/testdg/ii_bitmap

[root@hostA:/]
# sndradm -e hostA /dev/vx/rdisk/testdg/ii_shadow
/dev/vx/rdisk/testdg/sndr_bitmap hostB /dev/vx/rdisk/testdg/sndr_vol
/dev/vx/rdisk/testdg/sndr_bitmap ip sync
Enable SNDR? (Y/N) [N]: y

[root@hostA:/]
# sndradm -u
Refresh secondary with primary? (Y/N) [N]: y

[root@hostA:/]
# sndrstat

Type                                Vols                                s/n state
S P  vx/rdisk/testdg/ii_shadow => hostB:vx/rdisk/testdg/sndr_vol 0.00
REP

[root@hostA:/]
# sndradm -l
Put SNDR into logging mode? (Y/N) [N]: y

[root@hostA:/]
# mount -F vxfs /dev/vx/dsk/testdg/ii_master /master

[root@hostA:/]
# dex -R 80m 1 -w -x 64k -p 4 -e 5 -P4 /master/tmpFile
Pass 1, Errors 0, Elapsed time= 1 min.

```

Testing passed

```
[root@hostA:/]
```

```
# mount -F vxfs /dev/vx/dsk/testdg/ii_shadow /shadow
```

```
[root@hostA:/]
```

```
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /shadow/tmpFile
```

```
tmpFile: Error Block 127362 Bytes 8-15: 4294901760(0xffff0000)  
should be 3067833782(0xb6db6db6)
```

```
COMPARE ERRORS ON BLOCK 127362: 504
```

```
Data unchanged after reread on block 127362
```

TESTING FAILED

```
[root@hostB:/]
```

```
# mount -F vxfs -o ro /dev/vx/dsk/testdg/sndr_vol /secondary
```

```
[root@hostB:/]
```

```
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /secondary/tmpFile
```

```
tmpFile: Error Block 58242 Bytes 8-15: 4294901760(0xffff0000)  
should be 3067833782(0xb6db6db6)
```

```
COMPARE ERRORS ON BLOCK 58242: 504
```

```
Data unchanged after reread on block 58242
```

TESTING FAILED

```
[root@hostA:/]
```

```
# umount /master
```

```
[root@hostA:/]
```

```
# umount /shadow
```

```
[root@hostB:/]
```

```
# umount /secondary
```



---

**Note** – The test failure proves that the point-in-time copy shadow volume and remote mirror secondary volume do not contain the same data that is contained on the point-in-time master volume.

---

```
/** End Staging Phase **/
```

```
[root@hostA:/]
```

```
# iiadm -u s /dev/vx/rdsk/testdg/ii_shadow
```

```
[root@hostA:/]
```

```
# mount -F vxfs /dev/vx/dsk/testdg/ii_master /master
```

```
[root@hostA:/]
```

```
# dex -R 80m 1 -w -x 64k -p 4 -e 5 -P6 /master/tmpFile  
Pass 1, Errors 0, Elapsed time= 1:01 min.
```

```
Testing passed
```

```
[root@hostA:/]
```

```
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /master/tmpFile  
tmpFile: Error Block 1666 Bytes 8-15: 4294901760(0xffff0000) should  
be 3067833782(0xb6db6db6)  
COMPARE ERRORS ON BLOCK 1666: 504  
Data unchanged after reread on block 1666
```

```
TESTING FAILED
```

---

**Note** – The test failure shows the data pattern on the point-in-time copy master volume is now changed, but the original data pattern is still contained on the point-in-time copy shadow volume.

---

```
[root@hostA:/]
```

```
# mount -F vxfs /dev/vx/dsk/testdg/ii_shadow /shadow
```

```

[root@hostA:/]
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /shadow/tmpFile
Pass 1, Errors 0, Elapsed time= 1:02 min.

Testing passed

[root@hostA:/]
# umount /shadow

[root@hostA:/]
# iiadm -w /dev/vx/rdisk/testdg/ii_shadow

[root@hostA:/]
# svadm -d /dev/vx/rdisk/testdg/ii_shadow

[root@hostA:/]
# mount -F vxfs -o ro /dev/vx/dsk/testdg/ii_shadow /shadow

[root@hostA:/]
# /* Run vdbench program on /shadow/tmpFile */

[root@hostA:/]
# umount /shadow

[root@hostA:/]
# svadm -e /dev/vx/rdisk/testdg/ii_shadow

[root@hostA:/]
# mount -F vxfs /dev/vx/dsk/testdg/ii_shadow /shadow
[root@hostA:/]
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /shadow/tmpFile
Pass 1, Errors 0, Elapsed time= 1 min.

Testing passed

[root@hostA:/]

```

```

# umount /shadow

[root@hostA:/]
# sndradm -un

[root@hostA:/]
# sndrstat

Type                      Vols                      s/n state
S P  vx/rdsk/testdg/ii_shadow => hostB:vx/rdsk/testdg/sndr_vol 0.00
REP

[root@hostA:/]
# sndradm -ln

[root@hostB:/]
# mount -F vxfs -o ro /dev/vx/dsk/testdg/sndr_vol /secondary

[root@hostB:/]
# dex -R 80m 1 -r -x 64k -p 4 -e 5 -P4 /secondary/tmpFile
Pass 1, Errors 0, Elapsed time= 43 sec.

Testing passed

[root@hostB:/]
# umount /secondary

/*** End Verification Phase ***/

```







**Sun Microsystems, Inc.**  
4150 Network Circle  
Santa Clara, CA 95054 USA  
650-960-1300 Fax 650-969-9131

For U.S. sales office locations, call: 800-821-4643

In other countries, call:  
Corporate headquarters: 650-960-1300  
Intercontinental sales: 650-688-9000