



Sun StorEdge™ Availability Suite Software Point-in-Time Copy Software – Maximizing Backup Performance

A Best Practice

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

Part No. 816-7183-10
June 2002, Revision A

Send comments about this document to: 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, 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 relatant 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, 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 Snapshots 2

Issues With Snapshots 2

Maximizing Backup Performance 3

▼ To Maximize Backup Performance 3

Sun StorEdge Availability Suite Software Point-in-Time Copy Software – Maximizing Backup Performance

The Sun StorEdge™ Availability Suite Software point-in-time copy software is a snapshot facility that runs in the Solaris™ Operating Environment (Solaris OE). A point-in-time snapshot is an instantly available, time-fixed, replicated view of a momentarily quiesced storage volume. The use of snapshots for backup purposes helps to increase data availability by helping to reduce the backup window to just seconds. However, use of a snapshot can add I/O processing that impacts the performance of the backup procedure. This best practice discusses a method for maximizing backup performance using the export feature of the point-in-time copy software.

Point-in-Time Copy Snapshots

The point-in-time copy software creates a volume set in which an application volume is copied to a snapshot volume. The snapshot volume can then be used for backup purposes without accessing the application volume. The point-in-time copy software's export feature enables the snapshot to be removed from the volume set in which it was established for the duration of the backup procedure. When the snapshot is exported, the application server no longer maintains the snapshot volume, which improves backup performance. When the backup is complete, a fast resynchronization occurs during which the snapshot is updated with only the changes to the application volume, rather than a full volume copy.

The backup can be run on the application server, or the snapshot can be imported by a backup server to move backup processing off the application server. For a backup server to import the snapshot requires that the snapshot reside on a dual-ported storage device. Whether you are using a single server or a dual-server configuration, the major advantages of using the point-in-time copy software and its export functionality are as follows:

- Helps improve backup performance because the I/O overhead of reading from a snapshot is eliminated
- Helps improve application performance because the backups are performed on separate disk devices
- Helps increase application availability because the backup is done with a snapshot instead of waiting for the backup to be written to tape, which helps to reduce the backup window to seconds
- Helps protect investment, because existing backup software can still be utilized

Issues With Snapshots

Operational procedures, automated or otherwise, must be established to employ the point-in-time copy technology. Examples of operational issues that must be addressed are:

- When to do the backup snapshots
- How often to do the backup snapshots
- How to notify the backup software when the snapshot is available

In most cases, a point of data consistency must be reached before a point-in-time copy can be established so that the snapshot image is logically consistent.

Maximizing Backup Performance

This procedure explains how to make a point-in-time copy and how to export it for backup purposes. Before starting, make sure you are logged in as the root user. The following steps assume that the point-in-time copy software is installed on each server that has access to the point-in-time copy.

▼ To Maximize Backup Performance

1. Determine the volumes required for your point-in-time copy volume sets.

The point-in-time copy software uses shadow volume sets to manage point-in-time copies. A shadow volume set consists of a master volume, a shadow volume, and a bitmap volume. The master volume is the volume accessed by the application. The shadow volume is the point-in-time copy of the master. Use raw volumes for bitmaps. The bitmap in the volume set is where changes to the master are tracked while the shadow is exported.

The shadow volume set must be enabled as an independent shadow volume set so that a full physical copy of the master is performed.

Define an additional bitmap volume to track changes to the shadow volume while the shadow is exported. If the exported shadow volume is a backup, no changes are written to the shadow. However, the bitmap is still required for consistency checking when joining the shadow volume with its associated volume set on the application server.

To make the shadow and secondary bitmap volumes accessible to another server, place them in a separate disk group from the master and primary bitmap volumes. This enables the disks associated with the shadow volumes and their respective bitmap volumes to be deported in a single operation.

Example volumes are as follows:

- Master volume (belongs to the disk group `production`):
`/dev/vx/rdisk/production/vol01`
- Shadow volume (belongs to the disk group `backup`):
`/dev/vx/rdisk/backup/vol01`
- Bitmap volume associated with master (belongs to the disk group `production`):
`/dev/vx/rdisk/production/vol01.bmp`
- Bitmap volume associated with shadow during export (belongs to the disk group `backup`):
`/dev/vx/rdisk/backup/vol01.bmp`

2. Size the volumes.

An independent shadow volume requires the same amount of storage space as its master volume. The master volumes can be any RAID level supported by the Solaris OE and the volume manager software. Master, shadow, and bitmap volumes do not need to be the same RAID level.

The bitmap volumes are sized according to the size of the master volume. For each 1 Gbyte of raw disk space of the master volume, 8 Kbyte are required for the bitmap, plus an additional 24 Kbyte for overhead. So an 8 Gbyte master volume requires a bitmap of

$(8 * 8 \text{ Kbyte}) + 24 \text{ Kbyte} = 88 \text{ Kbyte}$.

3. Momentarily quiesce the application.

Flush to disk any data that is in the server's page cache. For `ufs`, do this with `lockfs`. Refer to your database software user manuals to determine the proper procedure for flushing data to disk.

4. Enable the point-in-time volume sets.

For manageability, you may group volumes being exported together into an I/O group. You can create an I/O group called `Backup1` using the `iiadm -g` command, so that the volumes can be managed as a single set.

```
# iiadm -g Backup1 -e ind/dev/vx/rdisk/production/vol01 /dev/vx/rdisk/backup/vol01 \
dev/vx/rdisk/production/vol01.bmp
# iiadm -g Backup1 -e ind/dev/vx/rdisk/production/vol02 /dev/vx/rdisk/backup/vol02 \
dev/vx/rdisk/production/vol02.bmp
# iiadm -g Backup1 -e ind/dev/vx/rdisk/production/vol03 /dev/vx/rdisk/backup/vol03 \
dev/vx/rdisk/production/vol03.bmp
```

5. Resume the application.

The application can be resumed while the copy to the independent shadow is in progress.

6. For I/O grouping, create a single point-in-time copy for the group.

Although the volume sets have all been enabled and inserted into the I/O group `Backup1`, they were set at different points in time. To have a single point-in-time copy for the group, invoke an update operation for the group. The `-p` option for `iiadm` locks the volume so that no other `iiadm` command can be run until the update completes.

```
# iiadm -g Backup1 -p -u s
```

7. When the independent copies are complete, export the shadow volumes.

For the creation of the initial point-in-time copy, this can take considerable time; subsequent updates are much quicker with fast resync capability.

- For an I/O group, export the group.

```
# iiadm -g Backup1 -E
```

- If there is no I/O group, export the shadow and bitmap volumes one pair at a time.

```
# iiadm -E /dev/vx/rdisk/backup/vol01
# iiadm -E /dev/vx/rdisk/backup/vol02
# iiadm -E /dev/vx/rdisk/backup/vol03
```

- If the shadow and bitmap volumes are not being imported on another server, skip to [Step 11](#).

8. Deport the disk group that hosts the shadow volume and import it to the backup server.

For example:

```
# vxdg -g backup stopall
# vxdg -g backup deport
```

9. Import the disk group on the backup server.

```
# vxdg -g backup import
# vxdg -g backup startall
```

10. Import the shadow volumes and bitmaps on the backup server.

```
# iiadm -I /dev/vx/rdisk/backup/vol01 /dev/vx/rdisk/backup/vol01.bmp
# iiadm -I /dev/vx/rdisk/backup/vol02 /dev/vx/rdisk/backup/vol02.bmp
# iiadm -I /dev/vx/rdisk/backup/vol03 /dev/vx/rdisk/backup/vol03.bmp
```

11. Perform the backup.

If the backup is performed on the application server, skip to [Step 14](#).

12. Disable the shadows and bitmaps from the point-in-time copy software on the backup server.

```
# iiadm -d /dev/vx/rdisk/backup/vol01
# iiadm -d /dev/vx/rdisk/backup/vol02
# iiadm -d /dev/vx/rdisk/backup/vol03
```

13. Reestablish the disk groups on the application server.

- From the backup server, deport the disk groups.

```
# vxdg -g backup stopall
# vxdg -g backup deport
```

- At the application server, import the disk groups.

```
# vxdg -g backup import
# vxdg -g backup startall
```

14. On the application server, join the shadows to their respective masters.

The bitmap associated with the shadow is merged with the original bitmap volume to reconstruct the original volume set consisting of the master, shadow, and bitmap volumes.

```
# iiadm -J /dev/vx/rdisk/backup/vol01 /dev/vx/rdisk/backup/vol01.bmp
# iiadm -J /dev/vx/rdisk/backup/vol02 /dev/vx/rdisk/backup/vol02.bmp
# iiadm -J /dev/vx/rdisk/backup/vol03 /dev/vx/rdisk/backup/vol03.bmp
```

15. Quiesce the application.

16. Perform a fast resynchronization to update the shadow with changes from the master.

For an I/O group:

```
# iiadm -g Backup1 -u -p s
```

Or, for each shadow volume:

```
# iiadm -u -p s /dev/vx/rdisk/backup/vol01  
# iiadm -u -p s /dev/vx/rdisk/backup/vol02  
# iiadm -u -p s /dev/vx/rdisk/backup/vol03
```

17. Resume the application.

The application can be resumed immediately after issuing the update commands.

18. Repeat [Step 7](#) through [Step 17](#) to continually update and back up the shadow volumes.



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