




A Sun StorEdge™ Rapid Restore Solution for Disaster Resiliency

By Raza Hussain - Network Storage

Sun BluePrints™ OnLine - April 2000



<http://www.sun.com/blueprints>

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Part No.: 806-5396-10
Revision 01, April 2000

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303 U.S.A. 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.

Sun, Sun Microsystems, the Sun logo, Sun StorEdge Instant Image, Sun StorEdge Rapid Restore, Sun Enterprise Server, Sun BluePrints 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.

UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

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, Californie 94303 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.

Sun, Sun Microsystems, le logo Sun, Sun StorEdge Instant Image, Sun StorEdge Rapid Restore, Sun Enterprise Server, Sun BluePrints, 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.

UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

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.



Please
Recycle



Adobe PostScript

A Sun StorEdge™ Rapid Restore Solution for Disaster Resiliency

This article is derived from the Sun BluePrints™ book *Business Continuity Planning with Sun Microsystems Technologies* by Stan Stringfellow et al. The book is a practical guide for IT organizations that are tasked with implementing or revamping a business continuity plan for information systems. For more information, see <http://www.sun.com/books/blueprints.series.html>.

Scenario

Sun StorEdge™ Rapid Restore can be used as a quick and efficient method for minimizing the potential of data loss and expediting the restoration of business to normal. With Sun StorEdge Rapid Restore, customers do not wait for the lengthy tape restores to finish before bringing their business back online. Sun StorEdge Rapid Restore utilizes a point in time data copy for the recovery of data. Production is back online as soon as the recovery process is initiated, which typically takes a few seconds.

The point in time image can also be used to push the data out to tapes, which in turn makes it possible to perform tape backups without bringing the production environment down.

Current Method Used for Restoration

Tape backup has long been the tool of choice to prevent data integrity problems. Problems such as lost files, mismanaged data, and operator errors have usually been solved by tape restoration. After normal business hours, the information systems would be quiesced and tape backups would be performed nightly as batch jobs. The only requirements were that the tape backups be accurate and complete before the next business day.

The problem with tape backup today is that it does not fit into the “24 x 7” global business environment. Businesses can not afford any data unavailability. Combine the intolerance to data unavailability and the size of today’s information systems (typically from many gigabytes to terabytes and growing at enormous rates), tape backup is problematical at best as a solution.

Technology and cost reductions have offered alternative solutions to tape backup in the form of electronic vaulting and remote hot sites—both of which are considerably more costly in operation than in initial investment.

New Method Used for Sun StorEdge Rapid Restore

As leaders in storage solutions, Sun Microsystems, Inc. offers a software solution based on the Sun StorEdge™ Instant Image, Veritas NetBackup software, and Sun StorEdge Arrays and Tape Libraries. The frequency of snapshots (point-in-time copies) taken will depend upon several factors including: rate at which the data is changing, workload, and application usage. On read intensive data, which does not change too frequently, taking snapshot once or twice a day may suffice. If data changes frequently, and the work load allows for snapshots to be performed, it may be appropriate to take a snapshot every hour to two hours.

General Operational Issues

Operational procedures, automated or otherwise, will need to be established to employ the point-in-time copy technology. Deciding when the snapshots will occur, the frequency in which they will occur, as well as notifying the backup utility of when the snapshot is available are examples of operational issues that will need to be addressed.

In most cases, a point of data consistency needs to be reached before a point-in-time can be established. This is to help ensure that the snapshot image is logically consistent (i.e., you have not captured any incomplete transactions).

The general procedure used to create a point-in-time copy that can then be accessed by a backup utility is the following:

1. Ensure that the master data has reached a point of consistency and flush all data to the disk(s)
2. Create a point-in-time copy of the data
3. Resume the master data application
4. Backup the point-in-time copy

5. Resynchronize the point-in-time copy with the master data when necessary

Configuration Details

This paper discusses using the Instant Image software to make a point-in-time copy of application data for backup and restore by Veritas NetBackup. The following table illustrates the configurations used:

Applications	Interfaces	Hardware	Operating System
Oracle	File systems	Sun Enterprise™ Server models ES4x00 to ES6x00,	Solaris™ 2.6 and 2.7
Veritas NetBackup	Raw	ES10000	Operating Environment

Terminology

- **Master Volume**—A physical/logical disk device that contains the original information for a particular application or suite of applications.
- **Shadow Volume**—A physical/logical disk device that contains the changed point-in-time data or full image of the original master volume.
- **Bitmap**—A file or physical disk partition used by Instant Image software to maintain the current state of the data between the Master and Shadow Volumes.
- **Independent Copy**—A completed independent copy is an identical copy of the point-in-time Master Volume (that is, a copy of the Master Volume when the image was taken) stored on the Shadow Volume. Once it is established, an independent Instant Image operation starts a block copy operation from the Master Volume to the Shadow Volume as a background process. While copying is in progress, the Instant Image software treats accesses to the Shadow Volume as if the Shadow Volume is a *dependent* image.
- **Volume Pair**—A volume pair consists of the Master Volume and its related Shadow Volume. Each volume pair is supported by one Bitmap. You can create as many volume pairs as necessary; the Instant Image software is only limited by your amount of available storage.

▼ To Setup Sun StorEdge Rapid Restore for a File system

This section describes the steps which need to be performed to establish a point-in-time copy of the production data. Before getting started, make sure you are logged in as the root user or superuser.

1. Edit the `/etc/opt/SUNWspcv/sv.cf` file and add an entry for the disks/volumes that are required for Sun StorEdge Rapid Restore.
2. Type “`svadm -r`” so that the disks and/or volumes in the `sv.cf` become available for Sun StorEdge Rapid Restore.
3. Create a Bitmap file or logical volume to store the state of the Instant Image volume pair. If it is not correctly sized, the Instant Image initialization process will fail. The correct size is 8-Kbyte for a header plus 8-Kbytes for every 1 Gbyte of the Master Volume.

```
mkfile size_of_file /name_of_file  
mkfile 32k
```

Note – Make sure you create the bitmap file on a file system that is required for Solaris operation.

4. To insure data accuracy, put a write lock on the file system using `lockfs`. This is for the master volume.

```
/usr/sbin/lockfs -w /file_system_name
```

5. Establish the Instant Image point in time copy.

```
iiadm -e ind master shadow bitmap  
iiadm -e ind /dev/rdisk/cXtXdXsX /dev/rdisk/cXtXdXsX /bitmaps  
iiadm -e ind /dev/vx/rdisk/<namedg>/<name master vol.> /dev/vx/  
rdisk/<namedg>/<name shadow vol> /bitmaps
```

Note – Make sure that the shadow is unmounted when establishing the point in time copy. Once the point in time process has been initiated, the shadow can be mounted right away.

Note – You would only establish the Instant Image volume pair once using the `-e` option.

6. To keep the shadow copy updated with the master data, execute the following command as often as required:

```
iiadm -u s shadow_volume
```

Note – You can establish as many volume pairs as you need. The Instant Image software is only limited by your amount of available storage.

▼ To Rapidly Restore File Systems from the Shadow

In the event of a production data failure due to data corruption or disk failures, data can be rapidly restored from the point in time image. To do this:

1. Replace the failing component if its a hardware problem or create a new file system if the problem was file system corruption.

Note – The file system may be repairable with `fsck`, use your judgement to repair or create new.

2. Make sure that the master in the Instant Image volume pair is unmounted.
3. Execute the Instant Image copy command which will copy data from the shadow onto master.

```
iiadm -c m shadow_volume
```

Note – As soon as the Instant Image copy command has been executed, the master volume can be mounted and can be made accessible for production use.

▼ To Backup File Systems from Shadow to Tape

Following is the data backup procedure to a tape drive/library using Veritas NetBackup software.

1. Mount the shadow image of the file system at a mount point. For Example: /shadow.

2. Start the Veritas NetBackup GUI by typing the following command:

```
/usr/openv/netbackup/bin/xnb&
```

3. Once the Veritas NetBackup GUI comes up, click on “Backup Management” icon. This will bring up the “Netbackup Administration” window on the screen.

4. Setup a Class for the file system that you wish to backup. In our case, it is the shadow image that needs to be backed up, and is mounted at /shadow. To create a Class, from within the “Netbackup Administration” window, go to:
Actions -> New -> Class

a. Assign the class name, for example: shadow_backup, and then click on the O.K. button. This will bring up a new window on the screen with the title “Changing Class shadow_backup”. In this window setup the “attributes” appropriately and then apply them.

b. Next setup the “Clients”, the name of the host from where data needs to be backedup.

c. Next setup a schedule to automate the backup process.

d. Lastly, setup “Files”, specify the absolute path of the file system that needs to be backed up. For example: /shadow.

The above steps will setup a “Class” for the automatic backup of data according to the schedule. To initiate a backup manually, continue with the following steps:

5. In the main window of “NetBackup Administration”, select the defined Class. For example shadow_backup in our case. Click the right mouse button on the highlighted Class, which will bring up a pull down menu. Select manual backup. A window with the title “Manual Backup of Class <class name>” will come up on the screen. Choose the schedule and the client that you want to perform the backup of, and click on the O.K. button.

6. To verify the backup activity, From the main window of Veritas NetBackup, click on “Activity Monitor” icon.

This displays the activity of the current backup in progress.

▼ To Restore File Systems From Tape

In the event of a failure of both the master and the shadow images of the data, the data can be restored from the tape using the tape backup and restore software. The following is the data restore procedure from a tape drive/library using Veritas NetBackup software.

1. After replacing the failed disk/s and the creation of a new file system on disk/volume, mount the file system where data needs to be restored.
2. Launch the Veritas NetBackup GUI by typing the following command:

```
/usr/openv/netbackup/bin/xnb&
```

3. In the main menu of the Veritas NetBackup window click on “User Backup and Restore”. A window with the title “Xbp” will appear on the screen. In the Xbp window under “Directory to Search”, specify the search path for the data that needs to be restored.
4. In the Xbp window, click on “File” on the top of the screen in the drop down menu list. In the “File” menu click on “Browse Backup (restore)” to look for the backup image that you wish to restore.
5. After choosing the backup image that you wish to restore, and highlighting “Directory Structure” and “Files” of the chosen image in the Xbp window, click on “Restore” from the drop down menu, and select “Specify alternate Path”. An “Xbp_altpath” window appears.

In “Restore From” field specify the mount point that was used to backup the data. In our case it was /shadow. In the “Restore To” field specify the mount point where you wish to restore data. In our case, I would like to restore it to the master disk/volume. If the mount point for the master disk/volume was /master, then this is what I will specify in the “Restore To” field of the “Xbp_altpath” window, and click on the O.K. button at the bottom.

6. Next, in the Xbp window, from the pull down menu choices, click on “Restore”, and then click on “Restore selected files and directories”. The “Xbp_confirm” window will appear. Click on O.K. in this window to continue with the restore of data.
7. To view the restore progress, in the Xbp window, from the pull down menu choices, click on “Restore”, and select “Report progress of restore”.

▼ To Setup Sun StorEdge Rapid Restore for Oracle Application Database

1. Edit the `/etc/opt/SUNWspv/sv.cf` file and add an entry for the disks/volumes that are required for Sun StorEdge Rapid Restore.
2. Type “`svadm -r`” so that the disks and/or volumes in the `sv.cf` become available for Sun StorEdge Rapid Restore.
3. Create a Bitmap file or logical volume to store the state of the Instant Image volume pair. If it is not correctly sized, the Instant Image initialization process will fail. The correct size is 8-Kbyte for a header plus 8-Kbytes for every 1 Gbyte of the Master Volume.

```
mkfile size_of_file / name_of_file  
mkfile 32k /bitmaps
```

Note – Make sure you create the bitmap file on a file system that is required for Solaris operation.

Note – If you choose to use a raw disk for bitmaps, then assign a slice to bitmaps with appropriate size.

4. Perform the following to take a hot snapshot of Oracle.
 - a. Put the Master database in archive mode via:

```
svrmgrl <<!  
connect internal;  
shutdown immediate;  
startup mount;  
alter database archivelog;  
alter database open;
```

b. The database is then put into backup mode via:

```
svrmgrl system/manager <<!  
SELECT tablespace_name,file_name FROM sys.dba_data_files;  
!  
svrmgrl system/manager <<!  
connect internal;  
alter tablespace system begin backup;  
alter tablespace tablespacel begin backup;
```

c. Take the snapshot of the database using the following Instant Image command:

```
iiadm -e ind /dev/rdisk/cXtXdXsX /dev/rdisk/cXtXdXsX /bitmap  
iiadm -e ind /dev/vx/rdisk/cXtXdXsX /dev/vx/rdisk/cXtXdXsX /bitmap  
(How big do these need to be? Same size as master? Can they be  
larger? Can they have different disk geo?)
```

Note – Enabling of instant Image volume pair is required only the first time. Once the pair is established, you will use the Instant Image update command to keep the shadow database in synchronization with master database.

5. Use the following command to perform the update:

```
iiadm -u s shadow_volume
```

Note – Make sure that the shadow disk/volume is unmounted when the `iiadm` command is executed. After the `iiadm` command has been executed, the shadow can be mounted, immediately.

6. Take the database out of backup:

```
alter tablespace tablespacel end backup;  
alter tablespace system end backup;  
archive log list;  
!
```

All of the above will be done from within a script.

▼ To Rapidly Restore a Database from the Shadow

In the event of a production database failure due to database corruption or disk failures, data can be rapidly restored from the point in time image. To do this:

1. **Replace the failing component if its a hardware problem or create a new file system if the problem was file system corruption (When you use file systems for databases).**
2. **Make sure that the master in the Instant Image volume pair is unmounted.**
3. **Execute the Instant Image copy command which will copy data from the shadow onto master.**

```
iiadm -c m shadow_volume
```

Note – As soon as the Instant Image copy command has been executed, the master volume can be mounted and can be made accessible for production use.

▼ To Backup an Oracle Database from Shadow To Tape

Following is the procedure to backup database on to a tape drive/library using Veritas NetBackup software.

1. **Mount the shadow image of the database on a mount point, if the database is on a file system. For Example: /shadowdb. Do not start the database.**
2. **Start the Veritas NetBackup GUI by typing the following command:**

```
/usr/opensv/netbackup/bin/xb&
```

3. **Once the Veritas NetBackup GUI comes up, click on “Backup Management” icon. This will bring up the “Netbackup Administration” window on the screen.**
4. **Setup a Class for the database that you wish to backup. In our case, it is the shadow image that needs to be backedup, and the database is mounted at /shadowdb. To create a Class, from within the “Netbackup Administration” window, go to: Actions -> New -> Class**

- a. Assign the class name, for example: `shadowdb_backup`, and then click on the O.K. button. This will bring up a new window on the screen with the title “Changing Class `shadowdb_backup`”. In this window setup the “attributes” appropriately and then apply them.
- b. Next setup the “Clients”, the name of the host where the database resides and needs to be backedup from.
- c. Next setup a schedule to automate the backup process.
- d. Lastly, setup “Files”, specify the absolute path of the database that needs to be backedup. For example: `/shadowdb`.

The above steps will setup a “Class” for the automatic backup of database according to the schedule. To initiate a backup manually, continue with the following steps:

5. In the main window of “NetBackup Administration”, select/highlight the defined Class. For example `shadowdb_backup` in our case. Click the right mouse button on the highlighted Class, which will bring up a pull down menu. Select manual backup. A window with the title “Manual Backup of Class <class name>” will come up on the screen. Choose the schedule and the client that you want to perform the backup of, and then click on the O.K. button.
6. To verify the backup activity, from the main window of Veritas NetBackup, click on “Activity Monitor” icon.

This will display the activity of the current backup in progress.

▼ To Restore an Oracle Database from Tape

In the event of a failure of both the master and the shadow images of the database, the data can be restored from the tape using the tape backup and restore software. The following is the database restore procedure from a tape drive/library using Veritas NetBackup software.

1. After replacing the failed disk/s and the creation of a new file system on disk/ volume, if the database was installed on a file system, mount the file system where the database needs to be restored. If the database was installed on a raw disks, then do not create the file system.
2. Launch the Veritas NetBackup GUI by typing the following command:

```
/usr/openv/netbackup/bin/xnb&
```

3. In the main menu of the Veritas NetBackup window click on “User Backup and Restore”. A window with the title “Xbp” will appear on the screen. In the Xbp window under “Directory to Search”, specify the search path for the database that needs to be restored.
4. In the Xbp window, click on “File” on the top of the screen in the drop down menu list. In the “File” menu click on “Browse Backup (restore)” to look for the backup image that you wish to restore.
5. After choosing the backup image that you wish to restore, and highlighting “Directory Structure” and “Files” of the chosen image in the Xbp window, click on “Restore” from the drop down menu choices, and select “Specify alternate Path”. An “Xbp_altpath” window appears.

In the “Restore From” field, specify the mount point that was used to backup the database. In our case it was /shadowdb. In the “Restore To” field specify the mount point where you wish to restore data. In our case, I would like to restore it to the master disk/volume. If the mount point for the master disk/volume was /masterdb, then this is what I will specify in the “Restore To” field of the “Xbp_altpath” window, and click on the O.K. button at the bottom. Make sure that “Rename Hard Links” and “Rename Soft Links” option are checked.
6. Next, in the Xbp window, from the pull down menu choices, click on “Restore”, and then click on “Restore selected files and directories”. The “Xbp_confirm” window will appear. Click on O.K. in this window to continue with the restore of database.
7. To view the restore progress, in the Xbp window, from the pull down menu choices, click on “Restore”, and select “Report progress of restore”.

Conclusion

Sun StorEdge Rapid Restore can be used as a fast and efficient method for minimizing the potential of data loss and expediting the restoration of business to normal. Rapid restore uses a point in time data copy for the recovery of data. Production can be back online as soon as the recovery process is initiated, which can take as little as a few seconds.

The point in time image can also be used to push the data out to tapes, which in turn makes it possible to perform tape backups without bringing the production environment down.

Related Documentation

TABLE P-1 Related Documentation

Application	Title	Part Number
Man Pages	iiadm(1M) svadm(1M) lockfs(1M)	
Release Notes	Sun StorEdge Instant Image Release Notes	805-7953
Installation and User	<i>Sun StorEdge Instant Image System Administrator's Guide</i>	805-7951
White Papers	<i>Point-in-Time Copy Technologies, A White Paper Featuring Sun StorEdge Instant Image</i>	
Best Practice Guides	<i>Application/Version Testing</i>	

[Author's Bio: Raza Hussain](#)

Raza Hussain is a Product Marketing Manager/Engineer for Software products in the Network Storage division of Sun Microsystems in Newark, California. Raza has over 8 years of experience in the Network Storage Industry and has held several engineering positions at EMC/Clariion, Siemens/Pyramid Technology, and Seagate Technology Inc. Raza holds a BS in Mathematics, Physics, and Economics and has received various System Engineer awards. He was also granted an Achievement award for Sales Development and Customer Presentation.