LAN-Free Backups Using the Sun StorEdge™ Instant Image 3.0 Software

Art Licht, Sun Microsystems, Inc.

Sun BluePrints™ OnLine–June 2002

http://www.sun.com/blueprints

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

Part No.: 816-4527-10
Revision A, 05/10/02
Edition: June 2002

Please
Recycle

LAN-Free Backups Using the Sun StorEdge™ Instant Image 3.0 Software

As data grows in size and backup windows shrink, performing backups across the LAN is no longer the ideal method. LAN-free backups should be used when the amount and importance of the data increases. This method enables the data and applications to remain available, while supporting shorter backup windows.

This document contains overviews of LAN and SAN backup practices and includes procedures for performing LAN-free backups.

LAN Backups

Currently, administrators perform backups across a network. They use products such as VERITAS NetBackup or Legato NetWorker™, but as the product names imply, they are designed to back up data across the network. These products are robust, and procedures for network backups are already in place.

Performing backups across the LAN is fine for small amounts of data, but as the data reaches the multi-terabyte range, the network is no longer the most effective method for large enterprise-wide backups. The SAN backup method takes advantage of these existing products and procedures to move the data across the SAN instead of the LAN.

SAN Backups

The Sun StorEdge™ Instant Image 3.0 software enables administrators to use their existing backup software procedures and benefit from the newer SAN technology. Most backup software can perform local backups, backups across the LAN, or backups across the SAN.

Traditionally, the data was only accessible over the LAN. With its deport shadow capability, the Instant Image 3.0 software enables administrators to make a snapshot of their application data, regardless of the format, and to export it to the backup server. The backup server can then back up the data directly through the SAN. This method has the following advantages:

- Higher performance

  The backup can be performed faster because the tape drives can be connected through several 100-Mbyte/second links through the SAN.

- Better application performance

  Application performance improves because the backups can now be done off a second set of spindles, in separate arrays, and down separate host bus adapters.

- Better application availability

  The availability of applications increases because the hot backup can now be done by using a snapshot, instead of waiting for the backup to be written to tape, which reduces the backup window to seconds.

- Uses existing backup software

  The existing backup software can be used to perform LAN-free backups.

Terminology

The following phrases are used in the document.

- Application server

The server that runs the application whose data is going to be backed up.

- Bitmap volume

The bitmap volume is a volume or physical disk partition used by the Instant Image software to maintain the current state of the data between the master and shadow volumes.

- Deporting the shadow volume

The shadow volume is deported by using the export, import, and join functions. This is the process by which an Instant Image software point-in-time (PIT) can be moved around the SAN. The procedure involves exporting the snapshot from the application server, importing it to the second server, then merging the master and shadow bitmaps together. This is called a join function.

- Independent copy

A completed independent copy is an identical copy of the master volume (that is, a copy of the master volume when the snapshot was taken) stored on the shadow volume. After it is initiated, an independent Instant Image software operation starts a block copy operation from the master volume to the shadow volume as a background process. After the process has completed, the independent PIT can be exported to another Solaris™ Operating Environment (Solaris OE) server.

- Master volume

The master volume is a physical disk device or logical volume that contains the original information for a particular application or suite of applications.

- Shadow volume

The shadow volume is a physical disk device or logical volume that contains the changed point-in-time data or full image of the original master volume.

- Volume pair

A volume pair consists of the master volume and its related shadow volume. Each volume pair is supported by one bitmap volume. You can create as many volume pairs as necessary. The Instant Image software is limited only by the amount of available storage.

LAN-Free Backup Procedure

This section contains general operational issues and a procedure for performing a LAN-free backup.

General Operational Issues

You will need to establish general operational procedures, automated or otherwise, to employ the point-in-time copy technology. You will also need to decide when the backup snapshots occur, how often they occur, and how the backup software is notified when the snapshot is available.

In most cases, a point-of-data-consistency needs to be reached before a point-in-time can be established. This ensures that the snapshot image is logically consistent (that is, the application can recover from any incomplete transactions).

The following list contains the general steps for creating a point-in-time copy that can then be accessed by a test application.

1. Make sure that the master application has reached a point of consistency (that is, it has flushed all data to the disk).
2. Create a PIT copy of the data.
3. Export the PIT copy to a backup server by using the deport shadow capability.
4. Initiate tape backup.
5. When required, resynchronize the PIT copy with the master volume data by using the Rejoin and the Fast Resynchronize functions.

▼ To Perform a LAN-Free Backup

This procedure describes how to make a PIT copy and export that copy over the SAN to a backup server. The following steps assume that the Instant Image 3.0 software is installed on all of the computers that have access to either the master or snapshot image.

1. **Log in to the system as superuser.**
2. **Determine which volumes are the master volumes and which volumes you will use for the independent shadow volumes.**

   It is key that the shadow volumes be on SAN-attached or dual-ported storage so that other servers can access these volumes.

   This step typically requires some planning. The idea is that the master volumes and original bitmaps remain available to the application server at all times. The PIT copy and its newly created secondary bitmaps should be available to the second server or backup server. Therefore, put the master and bitmap volumes in a diskgroup separate from the shadow and secondary bitmap volumes. This enables the deported shadow and its associated secondary bitmap volume to be deported and imported as a single diskgroup.

3. **Calculate the maximum amount of raw volume to be duplicated.**

   For an independent copy, you need a duplicate amount of usable storage to store the point-in-time copy. Take your existing diskgroup and volumes and duplicate them to another diskgroup. The online or master volumes can be any RAID level supported

by the Solaris OE or by a volume manager. The PIT, or shadow volumes, does not have to be the same RAID level as the master. The independent shadow volumes must be at least as large as the master volumes.

4. **Make both the master and shadow volumes available for the Instant Image software.**

   a. **Create two diskgroups, one for the online production volumes and another for the snapshots.**

      Any volume that is available to the Solaris OE is eligible to be used with the Instant Image software. If the shadow volume is to be exported to a second server, the shadow volumes must be in their own diskgroup. This enables the shadow disk group to be exported as a single entity using VxVM.

   b. **Create a volume for the bitmaps that can be export and imported along with the snapshot.**

      The easiest way to do this is to create a volume within the diskgroup.

5. **Use VxVM or the `format`(1M) command to allocate space for the Instant Image software to keep track of the differences between the master and the shadow volumes.**

   This is the bitmap associated with the master volume. The Instant Image software requires approximately 8 Kbytes for each gigabyte of raw disk space, plus space for header information (24 Kbytes).

   One bitmap is required for each Instant Image volume group (a group contains one master and one shadow volume). It is recommend that raw volumes be used for the bitmap volumes.

   The following is an example of the bitmap volumes used in this document:

```
/dev/vx/rdsk/production/vol01.bmp
/dev/vx/rdsk/production/vol02.bmp
/dev/vx/rdsk/production/vol03.bmp
/dev/vx/rdsk/production/vol04.bmp
/dev/vx/rdsk/production/vol05.bmp
```

   Additionally, the same bitmap volumes should be created in the backup diskgroup. It is these bitmaps that are used on the secondary server.

6. **Prepare the application and/or file system to be backed up.**

   Any data that is in the server page cache should be flushed to disk. For the UNIX™ file system (UFS), use the `lockfs`(1M) command. Databases have different procedures for flushing the page cache. For example, The Oracle® database software

uses the SQL commands `BEGIN` and `END BACKUP`. The Informix database software uses its `onmode` command. No matter which method you use, get all of the data to be backed up to the disk.

7. **Enable the point-in-time volume sets.**

   For manageability, all of the volumes that belong to this server-application should be grouped in an Instant Image software I/O group.

   This enables you to manage this backup as a single set.

```
# iiadm -g Backup1 -e ind /dev/vx/rdsk/production/vol01
/dev/vx/rdsk/backup/vol01 /dev/vx/rdsk/production/vol01.bmp
# iiadm -g Backup1 -e ind /dev/vx/rdsk/production/vol02
/dev/vx/rdsk/backup/vol02 /dev/vx/rdsk/production/vol02.bmp
# iiadm -g Backup1 -e ind /dev/vx/rdsk/production/vol03
/dev/vx/rdsk/backup/vol03 /dev/vx/rdsk/production/vol03.bmp
# iiadm -g Backup1 -e ind /dev/vx/rdsk/production/vol04
/dev/vx/rdsk/backup/vol04 /dev/vx/rdsk/production/vol04.bmp
# iiadm -g Backup1 -e ind /dev/vx/rdsk/production/vol05
/dev/vx/rdsk/backup/vol05 /dev/vx/rdsk/production/vol05.bmp
```

   The above commands create an Instant Image software I/O group called Backup1 and inserts each master and shadow volume set into that group. You can now manage the snapshot as a single entity. Although, you have enabled all of the volume sets into the same I/O group, at this time, they all have different PITs. To set a single PIT, you must invoke an Instant Image software update operation.

   The following is an example of an update operation.

```
# iiadm -g Backup1 -u
```

   In versions 1.x and 2.x of the Instant Image software, it was necessary to register the volumes into the storage volume software. This is no longer necessary. The `iiadm -e` command performs all of the needed storage volume registrations.

8. **Bring the application back online, if necessary.**

   In most cases, this step is not necessary because the application was never taken offline. For databases, you must take the database out of backup mode, clear any `onmode` command, and unlock any file systems, if previously locked.

9. **Export the shadow to another server.**

After the initial copy has been completed (only done once), you can export the shadow volumes to a secondary server. You can monitor the enable progress by using the `iiadm -i` command. If you are using scripts, the `iiadm -g Backup1 -w` command enables the script to wait until the copy is completely established.

The following is an example of how to export the shadow copy:

```
# iiadm -g Backup1 -E

or
# iiadm -E shadow_volume (for each volume)
# svadm -d shadow_volumes
# vxdg -g backup stopall
# vxdg -g backup deport (only necessary if using VxVM)
```

Note – It is necessary to disable the volumes from the storage volume software when they are going to be deported. This is volume manager specific. If no volume manager is used, removing the volumes from storage volume may not be necessary, but it is a good practice to follow. Logical volume managers require that volumes to be deported have no open processes.

---

10. Import the shadow disk group and snapshot volumes to the backup server.

Mount the file systems, if necessary. The shadow disk group and snapshot are imported by using the following commands:

```
# vxdg -g backup import
# vxdg -g backup startall
# iiadm -I /dev/vx/rdsk/backup/vol01 /dev/vx/rdsk/backup/
vol01.bmp
# iiadm -I /dev/vx/rdsk/backup/vol02 /dev/vx/rdsk/backup/
vol02.bmp
# iiadm -I /dev/vx/rdsk/backup/vol03 /dev/vx/rdsk/backup/
vol03.bmp
# iiadm -I /dev/vx/rdsk/backup/vol04 /dev/vx/rdsk/backup/
vol01.bmp
# iiadm -I /dev/vx/rdsk/backup/vol05 /dev/vx/rdsk/backup/
vol05.bmp
# svadm -e shadow_volumes
```

The next step is to mount the file systems, if any, that reside on the above volumes. The snapshot volumes can be put back into the Instant Image software I/O group, if desired, by using the following command:

```
# iiadm -g Backup1 -m shadow-1,shadow-2, &.. shadow-n
```

It is necessary to insert the volumes into the storage volume software to track all of the changes that occur to the shadow volumes while they are being used by the secondary server.

11. Start the backup process.

The only difference from a normal backup is that the data is now local, instead of out on the network. After the backups have successfully completed, the PIT can be removed from the backup server by using the following commands:

```
# iiadm -g Backup1 -d
# svadm -d shadow_volumes
# vxdg -g backup stopall
# vxdg -g backup deport startall
```

12. Rejoin the master and shadow volumes.

The Instant Image software logs the changes that are occurring on the master volumes as well as the shadow volume.

a. When you are ready to update the snapshot, export the bitmap from the second server back to the where the master database resides.

The Instant Image software determines which blocks need to be updated. This is referred to as a *join*. You do not need to recopy the entire volumes to update the shadow volumes. Only the modified data needs to be copied.

Next, ready the snapshot to be brought back to the master.

b. On the secondary server, stop the application, remove the volumes from storage volume, and deport the volumes by using the volume manger, as in Step 9.

To reattach the snapshot to the master, use the following commands on the master:

```
# vxdg -g backup import
# vxdg -g backup startall
# iiadm -J /dev/vx/rdsk/backup/vol01 /dev/vx/rdsk/backup/
vol01.bmp
# iiadm -J /dev/vx/rdsk/backup/vol02 /dev/vx/rdsk/backup/
vol02.bmp
# iiadm -J /dev/vx/rdsk/backup/vol03 /dev/vx/rdsk/backup/
vol03.bmp
# iiadm -J /dev/vx/rdsk/backup/vol04 /dev/vx/rdsk/backup/
vol04.bmp
# iiadm -J /dev/vx/rdsk/backup/vol05 /dev/vx/rdsk/backup/
vol05.bmp
# svadm -e shadow_volumes
```

The above commands take the changes on the secondary server, recorded in the bitmap volumes, and merge them with the changes done on the master-production server.

13. Resynchronize the master and shadow volumes.

a. Put the application into a known state (see Step 6).

b. Execute the following commands to update the shadow volumes:

```
# iiadm -g Backup1 -u s (for an atomic case)

or

# iiadm -u s /dev/vx/rdsk/backup/vol01
# iiadm -u s /dev/vx/rdsk/backup/vol02
# iiadm -u s /dev/vx/rdsk/backup/vol03
# iiadm -u s /dev/vx/rdsk/backup/vol04
# iiadm -u s /dev/vx/rdsk/backup/vol05
```

If this procedure executes from a script file, the Instant Image software can wait for the updates to complete before allowing any other commands. Use the `iiadm -w` command in the script.

You can apply the same scheme to applications that are on file systems; the Instant Image software works at the raw device level. Therefore, if the application resides on a file system, follow the these steps:

1. Quiesce the application (put the application into a known state).

2. Flush the page cache, if necessary.

   For example, use `lockfs /`*mount_point* on UFS filesystems

3. Establish or update the PIT copies, as described above.

4. Unlock the file systems, if previously locked.

5. Use `lockfs -u /`*mount_point*.

6. Mount and export the backup snapshot.

If you add more volumes and partitions in the future, follow the steps described in this document to enable the Instant Image software. Adding table spaces and data files to volumes that are already being used with the Instant Image software does not require changes to the Instant Image software.

You must modify the preceding for each particular installation. For repetition and ease-of-use, the sample commands used in this document should be put into a script. It is highly recommended that error checking be inserted as well. If an error should occur, the script should stop, and the error should be addressed.

Restoring Data From the Backup

You can restore the data from the backup by performing a Rapid Restore from the disk snapshot or by copying the data from a backup tape.

The fastest and most efficient way to restore data is with a Rapid Restore of the data. This is the same data that was last saved to tape. The snapshot can be mounted directly on the production server, and data or individual files can be copied directly from the snapshot to the production file systems. Additionally, the Instant Image software can, if desired, restore the entire set, group, or individual volumes by copying only the data that was modified since the last backup or snapshot. In this instance, the snapshot is the source volume, and the master is the destination volume during the fast resynchronization.

The slower, but more widely practiced method, is to restore the data from tape.

About the Author

Art Licht is a Storage Technology Engineer for Sun Microsystems, Inc. Art works with customers to build storage and SAN architectures, as well as architecting storage solutions such as disaster recovery, data-sharing applications, and backup and restore. Previously, Art worked in engineering on various storage hardware and software products.

Related Documentation

The following documents are available at the `http://docs.sun.com` site:

- Solaris 8 Reference Manual Collection:
  - `lockfs`(1M)
  - *Sun StorEdge Instant Image Release Notes*
  - *Sun StorEdge Instant Image Installation Guide*
  - *Sun StorEdge Instant Image Configuration Guide*
  - *System Administrators User Guide VERITAS Volume Manager User Manual* 806-7677-10

The following man page is part of the Sun StorEdge Instant Image software:

- `iiadm`(1M)