



Backup and Recovery by using SANWatch - Snapshot

Table of Contents

Introduction	3
Snapshot Functionality.....	3
Eliminates the backup window	3
Retrieves data from snapshot volume	3
Rollback	4
Use Case	5
How to use SANWatch/Snapshot for the EonStor arrays	6
<i>Create Snapshot Images</i>	6
<i>Back up Data to Tape or other backup space</i>	7
<i>Recover the corrupted or deleted files</i>	7
<i>Database Backup and Recovery</i>	7
<i>Cooperation with Backup Application via VSS</i>	7
<i>About VSS</i>	8
Conclusion	9
ABOUT the AUTHOR.....	9

Introduction

A snapshot is a point-in-time logical view of the data that is created by saving data differentials to a repository whenever the data in the source volume is about to be overwritten. The technique is also known as the “copy-on-write” technology that allows a snapshot to be instantaneously created while the source volume is actively serving I/Os. Essentially, the snapshot process creates an empty repository that will hold data differentials that are later changed in the source volume since the time a snapshot was taken. Unlike cloning or mirroring, snapshot does not require a similar size of disk storage when the first copy of data takes place. Snapshot images take up storage space only when the source volume receives updates that change the original data.

Using snapshot technology for data protection offers critical business values, such as:

1. Zero impact on normal I/Os,
2. The minimized or none application downtime,
3. The ability to make frequent backups (e.g., hourly) to reduce recovery time,
4. The efficiency when backing up large data volumes,
5. The reduced exposure to data loss,
6. and the ability of instant recovery.

Snapshot Functionality

Eliminates the backup window

Data backup and update can occur at the same time. Snapshot backs up the data in a RAID volume while the data is actively servicing I/Os (being changed). Compared to the traditional backup, the snapshot does not require extra downtime to perform backup and it eliminates the backup window. With the help of the automated scheduler utility, users can modify the time span between snapshots in accordance with the needs for the level of data protection.

Retrieves data from snapshot volume


An individual snapshot image can be mapped to a host ID/LUN and used as an ordinary data volume, and the volume keeps the data in exactly the state when the snapshot was taken. The snapshot volume is read-only, and users can

retrieve data files within it or the frozen records from the snapshot volume. Users can also compare the files or records frozen at different time points for special purposes, such as data mining. If the retrieved files or records were damaged by improper operation or human error, the damaged files or records can still be recovered from individual snapshot volumes.

Rollback

Using the rollback function, an administrator can easily reverse a volume to a snapshot taken previously at a specific point in time. Compared to restoration from tape, the rollback function only restores the data blocks that have been changed and access to the volume is not interrupted. This feature is useful with application testing and developing.

* Rollback will soon be available with the next revision of the SANWatch software.



Snapshot Operation Diagram

Use Case

The traditional backup is to backup data to the tape devices. If an administrator wants to backup 2TB of data to the tape library and the transfer rate of the tape library is 160MB/s, the backup process will take about 3.5 hours. Backup takes a considerably long time and the access to storage will be interrupted during the backup process.

Compared to tape backup, snapshot does not require extra downtime to perform backup. More significantly, snapshot backs up a data volume while data within the volume is still being changed. This feature keeps production volume online, and eliminates the backup window.

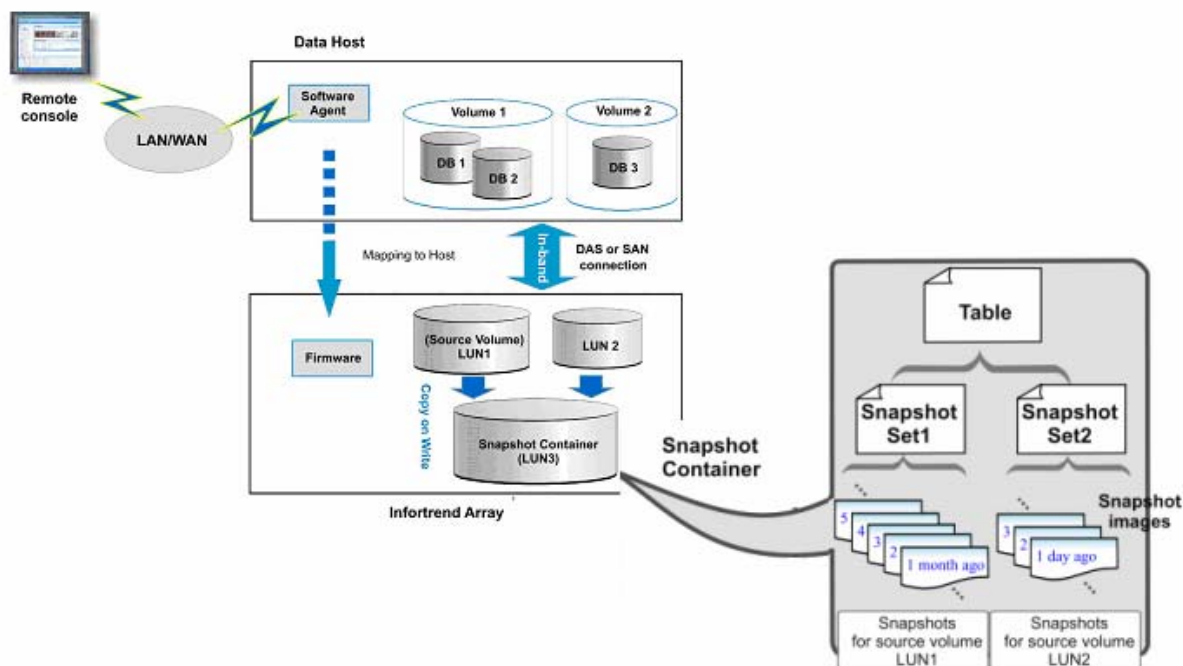
A tape library may provide the data compression function to increase its capacity. If the compression rate is 2:1, backing up 2TB of data will require 1TB of tape space. If using snapshot, the needed space is dependent on how much of data has been changed between points in time. In the worst case, a snapshot image may require a storage space as large as the source volume; but for a data volume whose data is not constantly and greatly changed, each snapshot image may only take a few megabytes.

If an administrator wants to access the backup data in tape, the backup data must be restored to a storage volume first. To restore 2TB of backup data may take about 4 hours and with the risk of high failure rate. And if a user wants to keep the currently running data volume intact, the restoration action requires another 2TB of disk space in order to retrieve the backup data.

If backing up using snapshot, the snapshot images can be mapped to the host ID/LUNs, and the mapped images behave as ordinary data volumes. Users can then access files or records in these snapshot volumes.

If users specifically want to access the snapshot volumes created at different points in time, no extra disk space is required. Access to the backup data falls on the differentials and the unchanged data in the original source volume. This feature decreases time to restore and the risk of restoration failure.

How to use SANWatch/Snapshot for the EonStor arrays



Before you start using the snapshot functionality, you will have to make sure what your target data is. For example, data can be identified as those located in files/volumes or in the form of database records. Since all cached data need to be flushed from system cache to storage to ensure usability, flush agents will be necessary on the DAS or SAN server depending on how a RAID system is attached to the application servers. Flush agents come as the SANWatch's "RAID agents," which are automatically installed along with the management software [Stand-alone (on host)].

Create Snapshot Images

1. Determine which source volumes contain the data of your interest.
2. Create one or more snapshot containers.
3. Create Snapshot Sets and associate them with source (working) volumes.
4. Install SANWatch on a DAS or SAN server using the [Stand-alone (on host)] installation option. Doing so enables an intermediate software agent that reads a source volume before taking its snapshots by flushing system cache.
5. Configure snapshot schedules and space control policies so that snapshots can take place automatically.

6. Snapshot images will then automatically generated in accordance with your schedules.

Back up Data to Tape or other backup space

1. Map the snapshot image of your interest to a host ID/LUN. The mapped volume appears as an ordinary RAID volume.
2. Perform your backup operation from this snapshot volume, i.e., if you have a backup application for tape, you can proceed with the backup procedure just like backing up from a system drive.
3. If necessary, you can keep or delete the snapshot volume through the SANWatch software.

Recover the corrupted or deleted files

1. Map the snapshot image that covers the corrupted/deleted files to a host ID/LUN. The mapped volume appears as an ordinary RAID volume.
2. Locate the files you want to recover and copy them to a preferred location.

Database Backup and Recovery

1. Install SANWatch on a DAS or SAN server using the [Stand-alone (on host)] installation option. Doing so enables an intermediate software agent that readies a source volume before taking its snapshots by flushing system cache. For example, you can install SANWatch on a server running MS SQL server or MySQL by including agents for these two applications.
2. We recommend you confine all databases within a single RAID volume because snapshot backs up the entire, specified source volume.
3. You can map the snapshot images to host ID/LUNs and open the database table to locate the data records within the mapped volumes when you want to access the historical data. You can copy the data records back to the currently running database if necessary. You can also back up your database volume to the tape library through the snapshot volumes.

Cooperation with Backup Application via VSS

If customers want to leverage the feature of RAID-based snapshot for their current storage management environment, they can

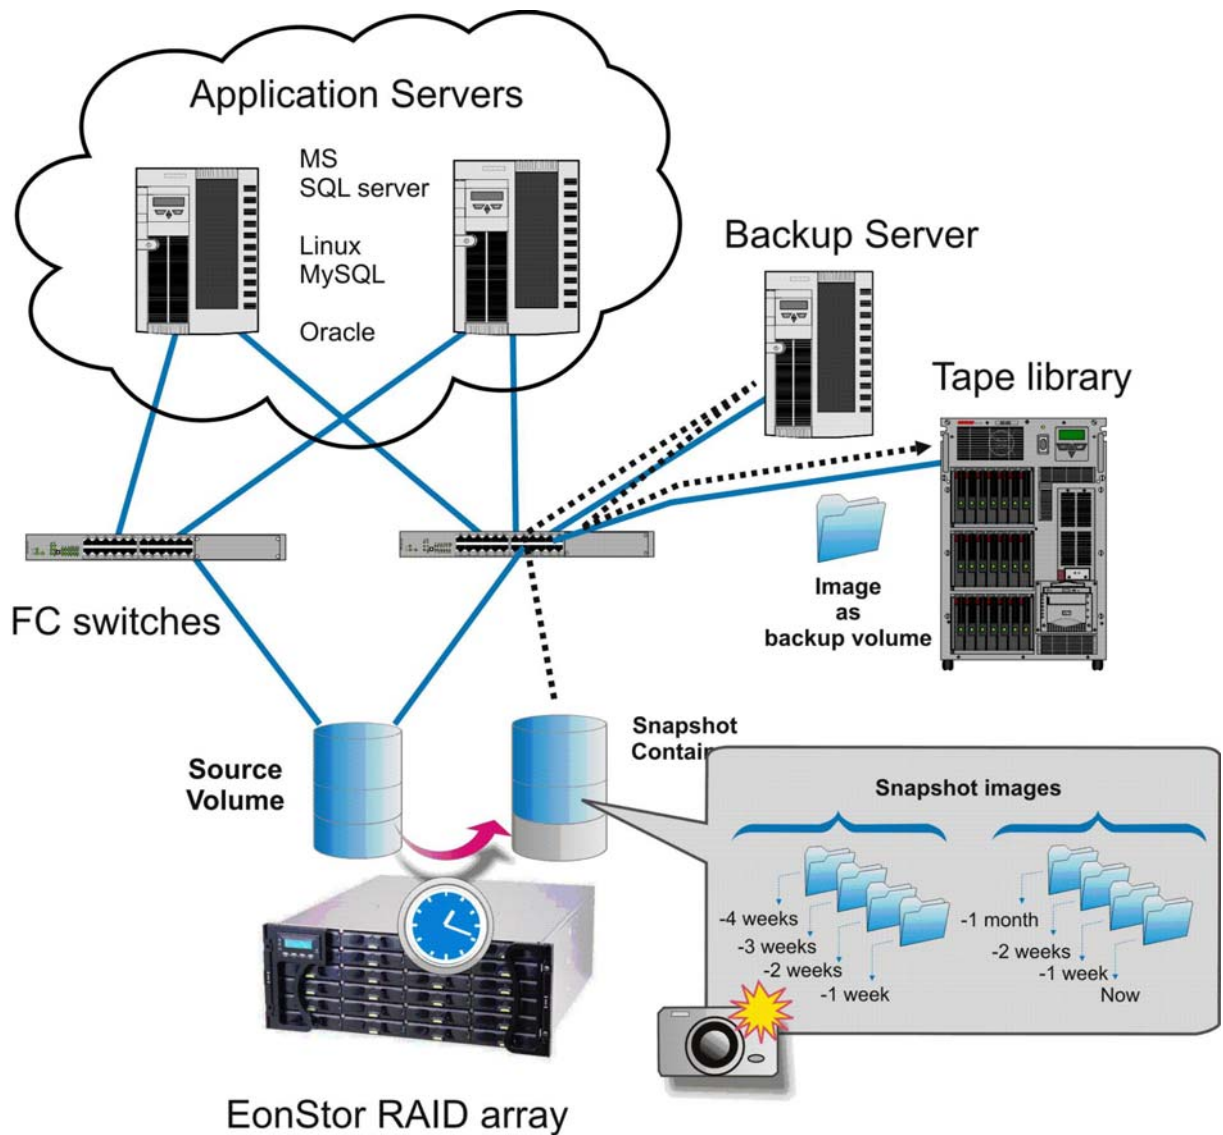
1. Install the VSS hardware provider from Infortrend onto the application server(s).
2. The VSS-aware application can trigger a snapshot through Infortrend

EonStor RAID system, and the process will be coordinated by MS VSS mechanism.

3. Specific flush agent is not necessary with VSS implementation since VSS will take care of the related synchronization process.

About VSS

The VSS support is not included in the ordinary SANWatch package, please contact your sales representatives and distributors for more information.



Conclusion

Snapshot is the first of the line of Data Service functionality that will include more exciting features such as Volume Copy and Remote Backup. In addition to the basic and well-known concepts like Copy-on-Write, Infortrend provides useful utilities such as snapshot scheduler and the purge/prune functions to increase flexibility ease of use, and to fulfill the requirements of data protection. Managed over the SANWatch interface, the snapshot implementation reduces the total cost of ownership by increasing the efficiency of the use of storage capacity and that of the backup/restore operation.

For detailed information, you can refer to the related documents for Infortrend EonStor RAID arrays and the SANWatch software.

ABOUT the AUTHOR

Mon-Chau Shiu is the Assistant Professor of Electronic Engineering Department, National Taiwan University of Science and Technology, Taipei, Taiwan. His research interests include computer architecture, embedded system, FPGA system, and image compression. Mon-Chau Shie was born in Nan-tou County, Taiwan. He received his Ph.D. degree in the electrical engineering, National Taiwan University, 2000.