



PS Series Groups Backup and Recovery Overview

Abstract

This Technical Report describes various methods and configurations for performing backup and recovery operations using a PS Series group.

Copyright © 2005 EqualLogic, Inc.

April 2005

EqualLogic is a registered trademark of EqualLogic, Inc.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

Possession, use, or copying of the documentation or the software described in this publication is authorized only under the license agreement specified herein.

EqualLogic, Inc. will not be held liable for technical or editorial errors or omissions contained herein. The information in this document is subject to change.

Table of Contents

Introduction	1
Backup and Recovery Operations with Different Media	2
Backup-to-Tape and Restore-from-Tape	2
Backup-to-Disk and Restore-from-Disk	2
Disk-to-Disk-to-Tape	3
Common Backup and Recovery Configurations	4
Local Backup and Recovery	5
LAN Backup and Recovery	6
SAN Backup and Recovery	6
Continuing Challenges for Backup and Recovery	7
Using Snapshots to Improve Backups	8
Local Snapshot Backup and Recovery	8
SAN Snapshot (Serverless) Backup and Recovery	10
Easily Integrating Snapshots with Backup Operations	11
Rapid Restore	12
Custom Integration of Snapshots with IT Operations	14
Summary	17
Documentation and Customer Support	17

Introduction

Traditionally, data has been backed up to tape and restored from tape. Using tape as the backup media is effective but it has limitations. This Technical Report provides an overview of backup and recovery architectures and describes various methods and configurations for performing backup and recovery operations using a PS Series group. It also describes how to use snapshot capabilities and disks as backup media to improve and simplify backup and recovery operations.

EqualLogic storage solutions represent a fundamental change in SAN implementations. By grouping together multiple PS Series storage arrays, both large and small businesses can deploy a SAN solution that easily scales to hundreds of terabytes affordably and while data remains online.

Built on ground-breaking patented architecture, PS Series storage arrays deliver enterprise-quality performance and reliability for all major operating systems. Intelligent automation and seamless virtualization of storage greatly simplify storage management. PS Series storage arrays come with a comprehensive set of high-level features—there is no extra software to purchase—and include fully redundant, hot-swappable hardware for a no-single-point-of-failure configuration.

A PS Series group is comprised of one or more PS Series storage arrays managed as a single system. A group is an excellent choice for serving primary application data and can also be used as part of a comprehensive backup and data protection solution.

Figure 1: PS Series Storage Array



Backup and recovery operations are the focus of business continuity and data protection plans and often the main source of anxiety for IT departments. Few businesses are fully satisfied with their backup and recovery solutions. Not only must data be protected from complete site failures, such as those resulting from natural disasters, data must also be protected from corruption or data loss, such as that resulting from a computer virus or human error.

An ideal backup and recovery solution:

- Maintains data integrity during backup operations to ensure that restored data is reliable.
- Retains multiple copies of data in safe locations, either local (for example, in the same building) or remote (for example, in a different geographic location).
- Ensures that backup processing has minimal impact on other IT operations.
- Allows data to be restored quickly and effectively, with minimal impact on users and applications.

A common challenge for administrators is to determine what is theoretically possible and what is practical with the backup products available today.

Backup and Recovery Operations with Different Media

The following outlines the evolution of backup media from traditional backup-to-tape to new software and hardware technologies that enable you to utilize disk as a backup media, which can improve performance and availability.

Backup-to-Tape and Restore-from-Tape

Backup-to-tape (or disk-to-tape, D2T) and restore-from-tape (tape-to-disk, T2D) procedures have evolved from configurations that use dedicated tape drives attached to servers, to centralized and shared tape libraries on a local area network (LAN) or storage area network (SAN). Tape is one of the most affordable and transportable backup media. However, tape backup media suffers from a variety of operational behaviors that make effective backup procedures difficult and inefficient. For example:

- Tape drives have unique performance characteristics. If data is not provided to the tape drives continuously at the proper rate (for example, if the data flow is uneven), performance can be significantly degraded.
- Data on tape media is not immediately accessible. Because tapes are stored separately or off-site, there are delays while the media is located and mounted for use.
- Data on tape is inherently not in the format applications can use. Data must be converted to backup format during the backup operation and then converted back to usable format during the restore operation, consuming time and resources.

Thus, using tape as the backup media imposes performance limits on backup and restore operations.

Backup-to-Disk and Restore-from-Disk

Backup-to-disk (or disk-to-disk, D2D) and restore-from-disk (disk-to-disk, D2D) procedures can be advantageous because disks do not have the performance and operational challenges of tapes. As disk capacities continue to grow and prices fall, many administrators now use disk as the media for some or all of their backup operations.

Unlike tapes, disks experience no performance problems resulting from uneven data flow during backup and restore operations. In addition, data restoration from disk does not entail long delays while the media is located and mounted.

Using disk as backup media improves data movement during backup and restore operations, significantly reducing the time needed for the operations to complete. After switching to disk, many customers experience dramatic improvements in backup and restore operations, often reducing operational times by half or more.


To achieve the benefits of disk as a backup media, EqualLogic PS Series arrays exhibit the following characteristics:

- High performance
- Scalable

- RAID protection
- Redundant, hot-swappable hardware
- Online service procedures, with no performance degradation

The Backup-to-disk operation typically works as follows:

1. The Backup server mounts a volume in the PS Series group, then, through its software, creates a backup-to-disk device pointing to it.
2. During the backup, server data is moved from its volume to the backup server's backup-to-disk volume.



There is careful planning involved when backing up only to disk. While disk drives do not share the performance and operational challenges found in tapes, as a backup media, they are not easily or reliably transportable. Because of the transportability issues, disaster protection may be limited depending on where the “backup disks” are located, relative to the primary data.

Disk as backup media also faces difficulties in long term retention and archival—most backup environments maintain backed up data for months or years. Other issues include power and cooling costs, as well as the management and servicing of disk technology (RAID protection, hardware, etc.). Although backup-to-disk has many advantages, it should not be considered a wholesale replacement for backup-to-tape, or other removable media.

Disk-to-Disk-to-Tape

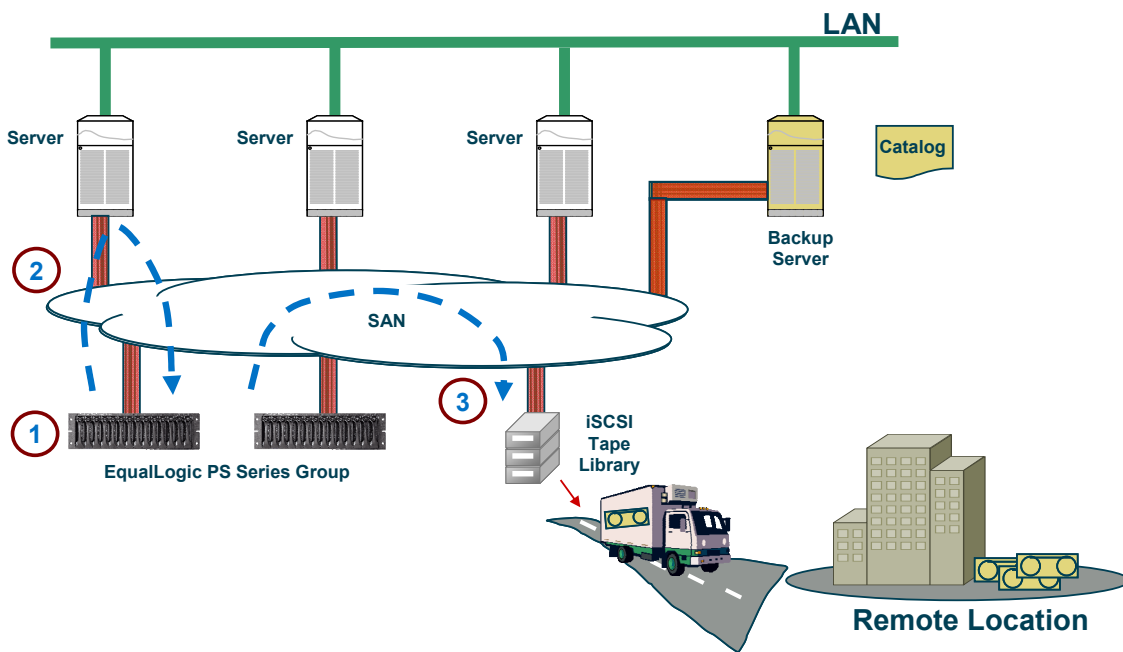
You can utilize both disk and tape as backup media with a disk-to-disk-to-tape (D2D2T) backup procedure. This common hybrid solution leverages the strengths of both disks and tapes to improve backup efficiency. Disk media is used to store data for short periods (for example, a week to a month), while tape is used for long-term data retention. This results in a solution that performs well and is cost effective over time.

Disk-to-disk-to-tape typically involves two copy operations during each backup operation: one copy from production data to disk media, and another copy from disk media to tape media. You get the performance benefits of backing up to disk, while retaining long-term copies of backed up data to meet business or regulatory requirements.

The D2D2T operation typically works as follows:

1. The backup server mounts a volume in the PS Series group and creates a backup-to-disk device pointing to it.
2. During the backup, server data is moved from its volume to the backup-to-disk volume.
3. The data that was backed up is then copied to tape.

Disk-to-Disk-to-Tape Backup Operation



Some backup applications also support running concurrent disk-to-disk backup jobs. This provides a significant performance improvement over some tape backup configurations that require each backup job to complete writing to tape before the next backup job can begin. In addition, disk-to-tape performance increases by 60% when the data is first backed up to disk.

Careful planning is required to ensure proper disaster protection during disk-to-disk-to-tape procedures. You also must have the system and storage resources to handle the second copy operation. Depending on the backup configuration, this second copy operation could be offloaded to another (backup) server in network or SAN-based backup environments.

Common Backup and Recovery Configurations

Disk can be used as backup media in a variety of backup and recovery configurations, including local, network, or SAN-based. In all cases, the backup software is configured to write to the disk backup media.

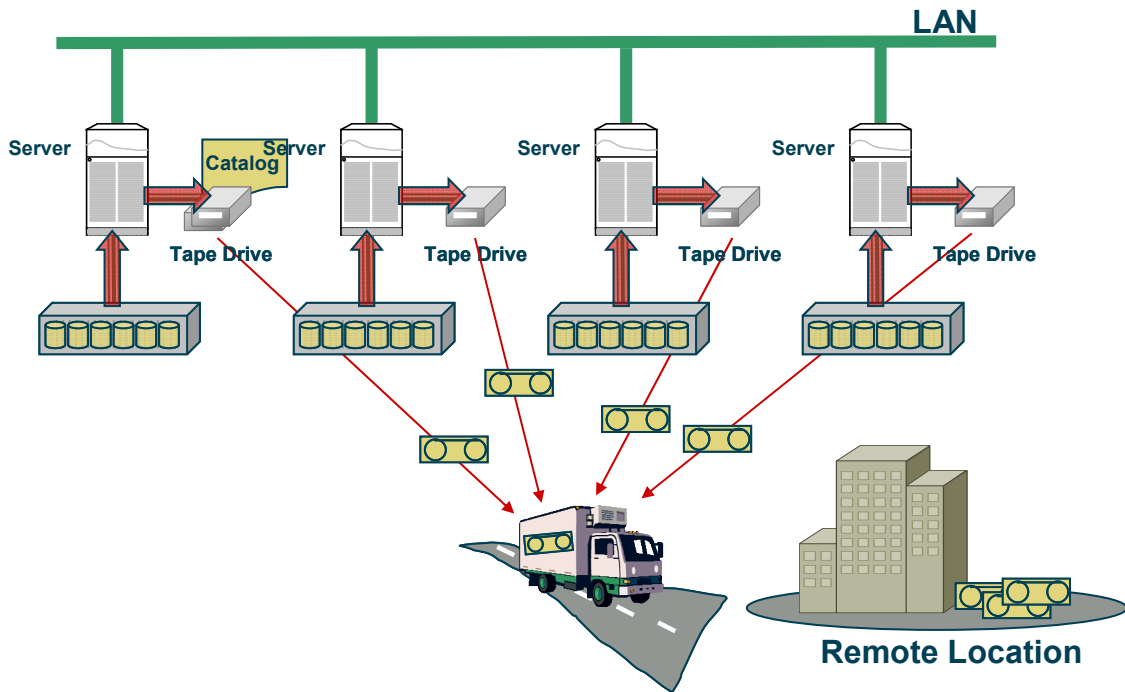
In configurations that include disk-to-disk backup, the backup software can view the disk backup media as a backup-to-disk device or as a virtual tape or tape library, depending on software and the configuration. Regardless of how the software sees the backup-to-disk media, backup and restore operations will still perform more efficiently.

The following sections describe common backup and recovery configurations. These configurations can use only tape backup media, only disk backup media, or both tape and disk as backup media.

Local Backup and Recovery

A common backup configuration is to perform backups locally on each server. With this configuration, backup operations run on the server, which copies data from disk storage to locally-attached backup media, usually tape. The tapes (or disks) can be safely stored away from the servers to protect against data loss in the event of a complete site disaster.

Backup to a Local Device




Local backup and recovery configurations have many challenges. As servers proliferate, backup operations become increasingly more difficult. Hardware and software costs increase, and there is a loss of IT productivity due to the management of multiple servers.

Recovery operations can also increase in complexity. The correct backup media must be located and mounted and the backup software must be installed on each server by experienced administrators. Because of this, recovery operations typically will take hours to complete. Note that in a complete site disaster, the backup catalogs may need to be recovered before you can recover the data, further delaying a return to operation.

LAN Backup and Recovery

As environments grow, consolidation of backup operations is key to reducing hardware and management costs. This involves using a central backup server to perform backup and restore operations. Backup agent software is installed on the application servers. The agents send data through the network (LAN) to the backup server, which then sends the data to the backup media, either tape or disk. The backup server also catalogs the data for future restore operations.



Backing up data over a LAN can lower hardware and management costs, as well as provide a more flexible, consolidated management environment. Typically, backup software licenses are needed for each server used in the backup configuration.

The challenges for LAN backup and recovery configurations are scaling and performance. These challenges are especially important when backup operations require moving a large amount of data. Because the data movement is performed by the application server using the network, it can cause performance bottlenecks for applications running on the server and also increase backup time.

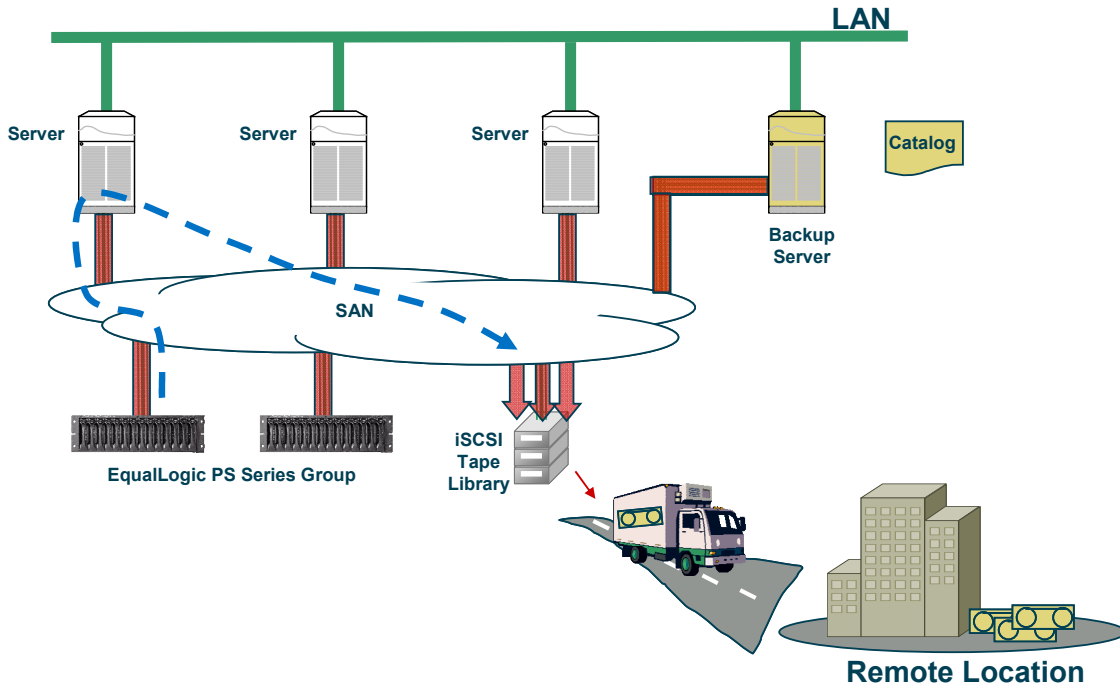
Restore operations may be more streamlined with LAN backups if a tape library or disk backup media is used. Recovery operations typically can take hours to complete when data is copied over the LAN. In disasters that involve only the application server, recovery operations can begin quickly because the backup server and catalogs are online; however, the application server must be restored with the backup agent software functioning. In larger disasters involving the loss of the backup server and the application servers, recovery requires first restoring the backup server, and then restoring the application servers.

SAN Backup and Recovery

As with LAN backup, SAN backup enables you to consolidate backup catalog management and restore operations on the backup server. Backup media can be connected directly to the network, and the backup server coordinates backup and restore activity with each server. Backing up over a

SAN can improve backup performance because all data movement is through the SAN, thus eliminating the backup traffic from the LAN. This is often called LAN-free backup.

Backup Over a SAN



Restore operations may be more streamlined when backing up over a SAN if a tape library or disk backup media is used. Recovery operations typically can take hours to complete when data is written to tape over a network. In disasters that involve only the application server, recovery operations can begin quickly because the backup server and catalogs are online; however, the application server must be restored with the backup agent software functioning. In larger disasters involving the loss of the backup server and the application servers, recovery requires first restoring the backup server, and then restoring the application servers.

Continuing Challenges for Backup and Recovery

Backup configurations (local, LAN, and SAN) share some challenges. These challenges exist regardless of backup media used (disk or tape), and are as follows:

- **Complexity in backing up running applications.** Data needs to be stable during backup copy operations; otherwise, the backed up data may not be useful. This problem is typically handled by the backup software agents on the application servers. A big part of IT overhead is complex backup environments, in which operating systems, applications, networks, and backup software must operate seamlessly.
- **Long backup times having a significant performance impact on application servers.** During backup operations, data is copied to backup media and converted from application format to backup format. IT operations can be adversely affected during backup operations. Having application servers responsible for data movement to backup media can lengthen backup windows and disrupt running applications. A backup operation can be considered an

application with a heavy workload that runs daily in your environment. This workload increases demands on servers and infrastructure and must be factored into capacity planning for each server. If demand exceeds capacity, performance may degrade and service may be disrupted.

- **Restore operations requiring backup application software to be operational and the backup catalogs to be accessible.** Because backed up data is not in usable format, the backup application must first convert it. Restore operations can take hours to convert data to a usable format so you can resume operations. Depending on the type of disaster, the backup software may need to be restored before the data is restored, further lengthening the time to recovery.

If not properly managed, these challenges can result the following situations:

- IT administrator restores backed up data, but the application fails to operate after the restore. This is a common problem that is caused by using an improper procedure to back up a running application. This problem is not related to backup media, but concerns how backup software interacts with running business applications, file systems, e-mail systems, and databases.
- Poor performance or down time for application servers during backups.
- Backups not completing within their allotted backup windows.
- Long down time after data loss due to lengthy recovery procedures.

For many businesses, backing up “live” data for a running application is an essential, but difficult, endeavor. You must consider the complexity of operating systems and applications to ensure that backups complete successfully.

Using Snapshots to Improve Backups

The challenge that administrators face today is how to set up efficient and effective backup and restore operations in complex environments, while data grows at a rapid rate.

Snapshots—point-in-time copies of data—enable you to quickly copy data at the disk (or volume) level. This stable “copy” of data can then be used as the source for your backup operations. Snapshot creation does not disrupt access to the volume. On a PS Series group, the copy is created instantly, usually in a few seconds, and maintained on disk storage in the PS Series group, providing high performance with low disk space utilization.

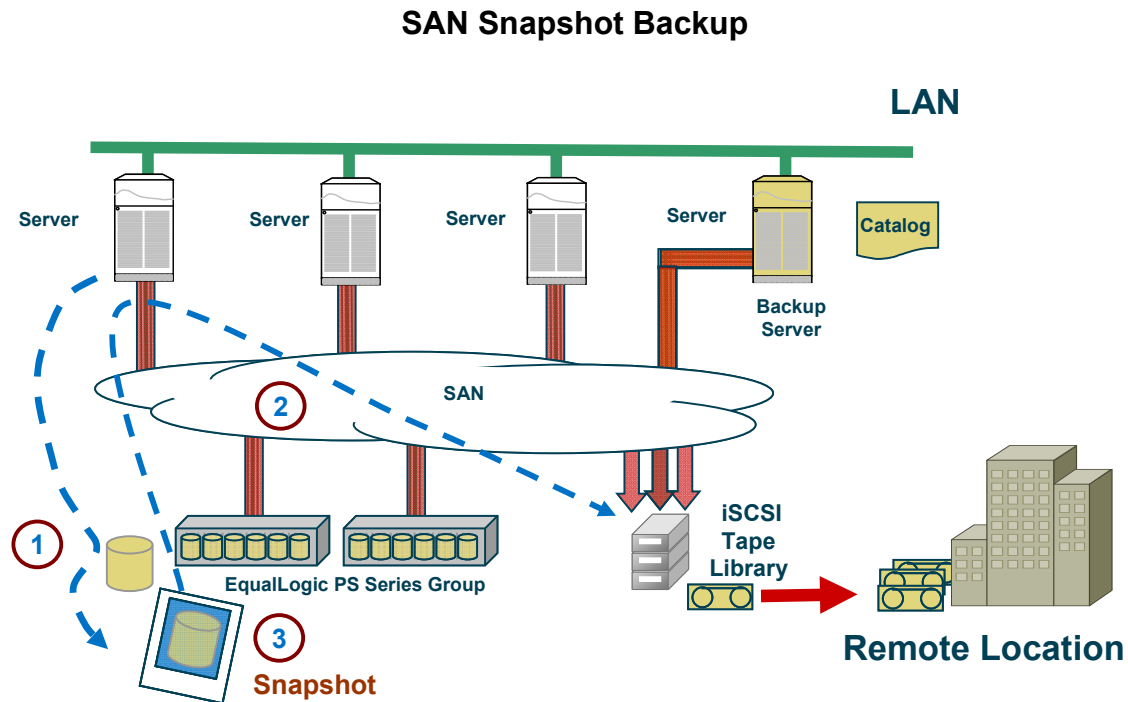
Local Snapshot Backup and Recovery

Snapshots can provide a way to back up “open files”, one of the primary challenges of IT administrators. For effective backups, you need a stable copy of data for the backup application to copy to backup media. If data is written to an “open file” while it is being backed up, the integrity of the data may be in jeopardy.

Snapshot technology enables you to back up live application data, while ensuring that the data you retrieve from the backup is reliable and useful. Snapshots can be used to stabilize application data in local, LAN, or SAN backup configurations.

The local snapshot operation typically works as follows:

1. Application server prepares for and creates a snapshot, capturing a stable copy of the data.
2. Application server mounts the snapshot and copies data from the snapshot to backup media.
3. Snapshot is deleted.



The advantage of using snapshots for backup and recovery operations is that snapshots can provide a stable copy of data for copying to backup media. The application is only briefly affected when the snapshot is created, which typically takes a few seconds. Once the snapshot is created, the application server is able to continue application operations, and the snapshot can be backed up. This may safely extend the backup window.

The application server is still involved in moving data from the snapshot to the backup media or backup server. Both full and incremental backups are available with snapshots (including brick level Exchange backups of individual mailboxes).

Restore operations are more reliable, because snapshots ensure the integrity of the backed up data. Restore procedures and timing are similar to backups without snapshots.


SAN Snapshot (Serverless) Backup and Recovery

The disadvantage of using local snapshots is that the application server still plays a role in the moving data. While the integrity of the backed up data improves, application server performance may suffer during backup operations.

Building on the configuration described in *SAN Backup and Recovery*, a serverless backup operation can create a snapshot in the SAN. Then, the backup server can mount the snapshot and copy the data to backup media. This removes the application server from the data copy process, offloading the operation to a backup server on the SAN.

The SAN snapshot (serverless) backup process typically works as follows:

1. Application server prepares for and creates a snapshot, capturing a stable copy of the data.
2. Backup server mounts the snapshot and then copies data from snapshot to backup media.
3. Snapshot is deleted.



Transportable snapshots can improve application performance in your server environment during backup operations. The snapshot provides a stable copy of data for the backup server to copy to backup media. The application is only briefly affected when the snapshot is created, which typically takes a few seconds. Once the snapshot is created, the application server is able to continue application operations, and the snapshot can be mounted and backed up by the backup server. This can dramatically and safely extend the backup window.

Using snapshots in a SAN typically requires homogenous operating system environments for the application servers and backup servers. Because the application is not running on the backup server, backup servers may require full backups of databases and e-mail systems.

This solution can improve backup data copy performance and maintain centralized backup management. SAN snapshots require shared storage devices (e.g. PS Series storage arrays) on a

SAN. Typically, backup software licenses are needed for each server used in the backup configuration.

Restore operations are more reliable, because snapshots ensure the integrity of the backed up data. Restore procedures and timing are similar to backups without snapshots.

When used in a SAN configuration, snapshots can solve all backup challenges, providing stabilized data for backups, a dramatic performance improvement for application servers, and fast restore capabilities.

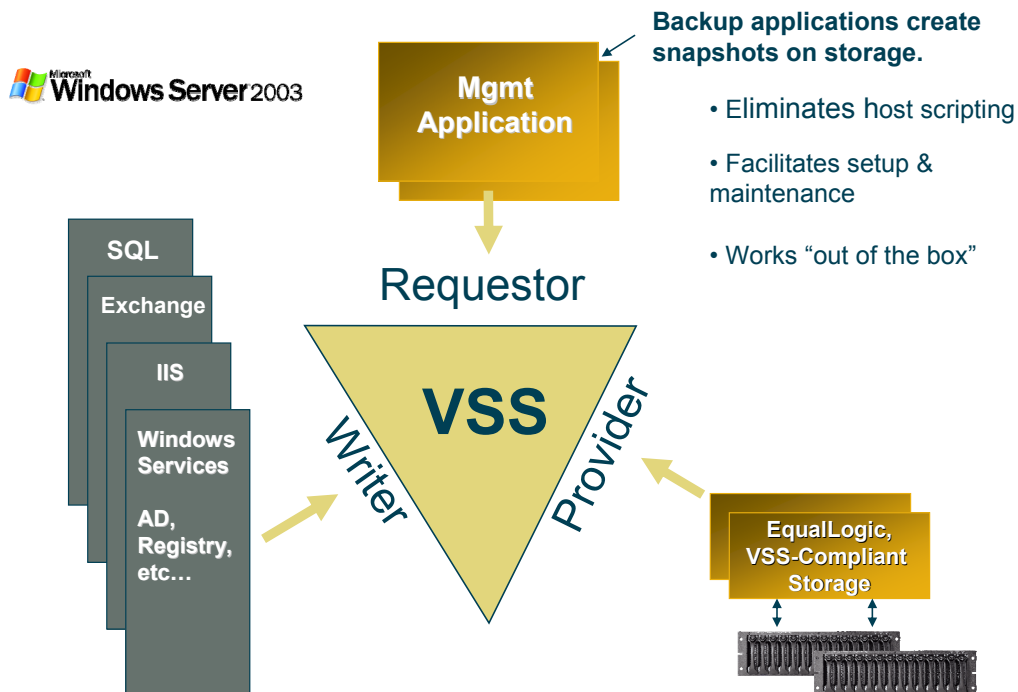
Easily Integrating Snapshots with Backup Operations

When using snapshots in backup operations, you must integrate the snapshot capability with backup applications, the applications and file systems to be backed up, and the storage devices. Historically, this integration has required using scripts, which are difficult to create and maintain for proper operation over time. These requirements have severely restricted the adoption of snapshot-based backups.

What is needed is a turn-key method that administrators can use to deploy snapshot-based backups. This can be accomplished by making operating systems and applications “snapshot aware,” shifting responsibility for snapshot integration from IT administrators to the snapshot technology providers.

Microsoft® has created a technology in Windows® Server 2003 called Volume Shadow Copy Service (VSS). VSS provides a framework that integrates VSS-aware storage hardware and applications with operating system drivers to create point-in-time copies of data (snapshots), delivering a turn-key backup solution to IT departments without the need for scripting.

VSS on Windows Server 2003



The VSS model requires the following for backing up data with snapshots:

- **Requestor** (or backup application), which “requests” the creation of snapshots, typically for backup operations.
- **Writer** (a business application such as a database, e-mail application, or file system), which prepares the application for the snapshot or restore (for example, by flushing buffers or switching logs).
- **Provider** (a storage product, such as a PS Series group), which is the mechanism that creates and maintains the snapshot.

Using this model, snapshots can be seamlessly integrated with backup operations, regardless of whether the snapshots are local or serverless. (In VSS terms, serverless snapshots are referred to as “transportable.”)

Auto-Snapshot Manager for Windows from EqualLogic can simplify and improve the performance of backup operations. Auto-Snapshot Manager is used in conjunction with VSS, backup and restore applications, and a PS Series group to create coordinated snapshots of group volumes.

When installed on a Windows 2003 Server, Auto-Snapshot Manager enables you to instantly create a snapshot—called a *shadow copy*—while the application remains online and with no impact on performance. Each shadow copy is a point-in-time copy of data, the same as a snapshot.

Backup and restore products are at various stages of VSS implementation. Backup applications that support VSS allow local snapshots at a minimum. This means the application server that created the snapshot can mount the snapshot and move the backup data. Others support transportable snapshots so that the application server is essentially eliminated from the backup process.

Note: For serverless backup operations, be sure your backup application supports the hardware provider, in addition to transportable VSS shadow copies.

Backup applications that support VSS include the following:

- VERITAS Backup Exec™
- CommVault® Galaxy™ Backup & Recovery
- CA BrightStor® ARCserve® Backup
- Bakbone® NetVault
- Legato® NetWorker

See each vendor’s backup application documentation for more information about using the application as a VSS requestor.

See EqualLogic’s Auto-Snapshot Manager for Windows *Installation and Administration* manual for more information about using that VSS provider.

Rapid Restore

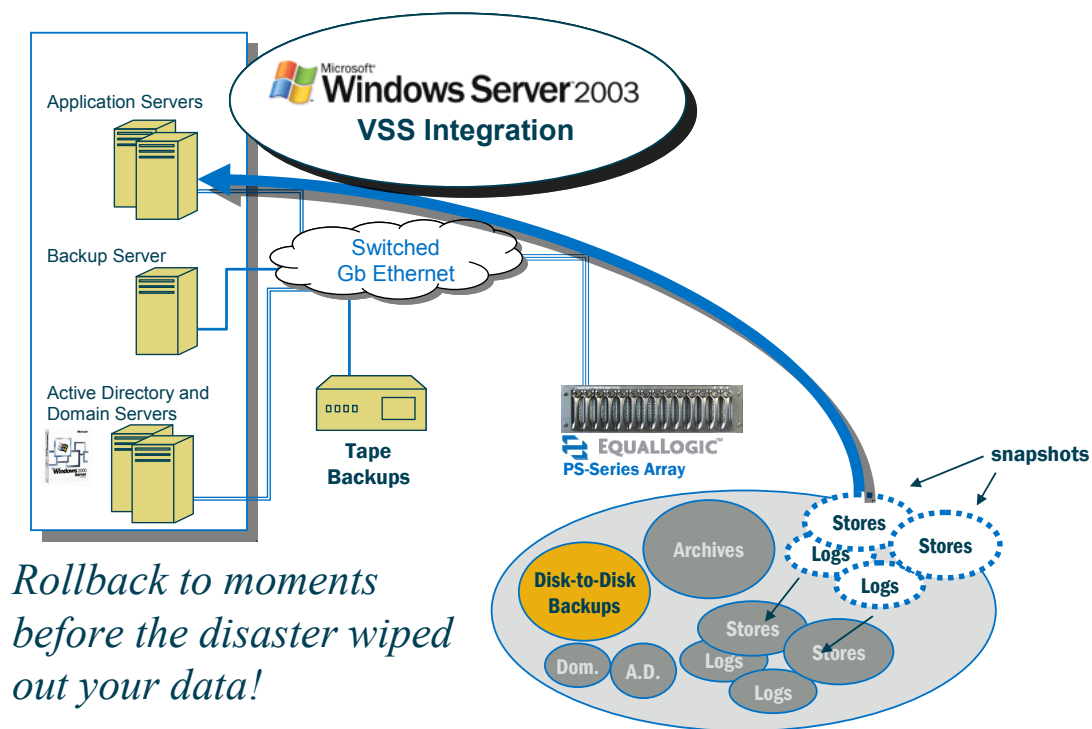
In all of the backup scenarios discussed in this Technical Report, the main focus is to improve the backup operation. Since backup operations are much more common than restore operations, it is appropriate to focus on backups.

For 24x7 businesses, recovery procedures must be fast, in addition to reliable. The problem is that data within the backup system is in backup format, not in a format that can be used by applications. Recovery involves copying and converting data from backup format to a usable format, a process that inherently takes time and increases the total recovery time.

The solution is to maintain copies of data in a usable format on disk. This will avoid the copy and convert process and provide dramatic time savings, but at a cost of additional disk storage. On-disk copies can be provided by making copies of application data available, either through scripts or quick recovery backup products. On-disk copies can also be provided by using snapshot technology and retaining the snapshot beyond the duration of the backup operation.

Some backup and restore products contain options that leverage and manage VSS, snapshot, and disk-to-disk backup (D2D) technologies to provide rapid data restore. The backup application schedules the snapshots at regular intervals, creating multiple restore points. The snapshot frequency can be scheduled hourly or more frequently, depending on application and business requirements. The backup application can directly utilize snapshots for recovery operations or can utilize their own disk copies if they are in a usable format.

Rapid Restore



By leveraging snapshot-based backup and rapid restore technologies, you can decrease backup windows and decrease restore times by as much as 90%. In addition, you gain maximum storage efficiency by leveraging snapshots and rapid restore to allow you to “rollback” to the data you had moments before the disaster wiped out your data!

All of this provides a more complete backup and recovery solution that nearly eliminates the application server’s involvement in the process. This improves restore capabilities, because snapshots can be done more frequently and can provide multiple recovery points.

Custom Integration of Snapshots with IT Operations

Many operating systems do not support native snapshot integration in the operating system, as Windows Server 2003 provides with VSS. Also, there are usage models that may not support this integration, regardless of whether the operating system provides snapshot integration.

When IT administrators use snapshots without the proper operating system integration, the snapshot (or replica, in some cases) is created with a status of “power fail or crash consistent,” “application file consistent,” or “inconsistent.” A description of each status is as follows:

- “Power fail or crash consistent” snapshots provide the volume state as if the power had been turned off to the server. If your server had a power failure, application data is not in a clean state. When you restart the server, file system or application recovery procedures are executed to make these copies usable. These procedures are typically run automatically when the server reboots.

For disaster protection, this is a very common form of data protection provided by synchronous and asynchronous replication or mirroring, and can dramatically improve recovery times, compared to typical software restore procedures. See the EqualLogic PS Series product documentation for more information regarding replication.

- “Application file consistent” snapshots provide application state recovery as if the application had been stopped just prior to the snapshot creation. This minimizes the recovery procedures needed for the snapshot data. Note that some application recovery and consistency checks are still run.
- “Inconsistent” snapshots typically involve applications that span multiple disk volumes, but the snapshots of these volumes were not coordinated (for example, created at the same time). These snapshots have a limited use, because you cannot assume any level of consistency for the data spanning the volumes.

Administrators need some form of consistent snapshots in their IT operations. Inconsistent snapshots should be avoided because they make IT operations unpredictable. The choice for snapshot type (“power fail or crash consistent” or “application file consistent”) depends on your use model, operational needs, and the level of complexity in setting up and maintaining the environment.

For sending data for long-term backup storage, the data on backup media should be in application file consistent format. This can be done in a variety of ways:

- Use backup software that ensures data is in application file consistent format.
- If an application does not provide an “end-to-end solution,” then you can:
 - Use application APIs to “flush” data to disk in application file consistent format.
 - Use application-based replication methods to create a copy at the application level and then backup this copy. This can be done with application-based replication or log shipping.
 - Stop applications before creating snapshots of the data.

- Create crash consistent snapshots and run recovery procedures before copying to backup media.

Snapshots can also be used to improve recovery operations. In this model, the snapshot is kept for some period of time, long enough so that recovery operations can use the snapshot as the source of recovered data. Data can be in either “power fail or crash consistent” or “application file consistent” format. Using snapshots results in dramatically fast restore operations. However, the tradeoffs for the two solutions are as follows:

- “Power fail or crash consistent” format is typically easier to set up and maintain. Recovery procedures will typically have an additional step or two for restoring application consistency.
- “Application file consistent” format is typically harder to set up and maintain. Recovery procedures will typically have one or two fewer steps than “power fail or crash consistent” format.

Depending on the recovery scenario, the process for restoring operations may be automatic or manual. When recovering a server from a catastrophic failure, both solutions would restore all disks using snapshot data, and the server would be restarted:

- With “power fail or crash consistent” snapshots, the server would restart as if power failed at the time the snapshot was created.
- With “application file consistent” snapshots, the server would restart as if it had been shut down at the time the snapshot was created.

In both cases, restart times are fast, and data is restored to what existed when the snapshot was created. If some of the newer data is still available (for example, if you have logs for a database or e-mail system), the IT administrator may be able to manually roll forward either of these configurations to the most current data.

If you are trying to do a partial restore (for example, to recover just a file or e-mail message):

- With “power fail or crash consistent” snapshots, you will have to manually run file system or application recovery procedures (for example, `chkdsk`, `fsck`, `eseutil`, etc.) on each mounted snapshot.
- With “application file consistent” snapshots, the mounted snapshot can be used immediately.

In both cases, the administrator must manually run file copy operations and the application stop and restart procedures to complete the restore operation.

A PS Series group provides several methods of automatically creating snapshot or scripting snapshot operations. For example, you can:

- Create schedules to automatically create snapshots (or replicas) at a specific time in the future or on a regular basis.
- Script snapshot and replica operations without using VSS.
- Use application-based backup (log shipping).
- Use quick recovery solutions (for example, DPM or backup tools).

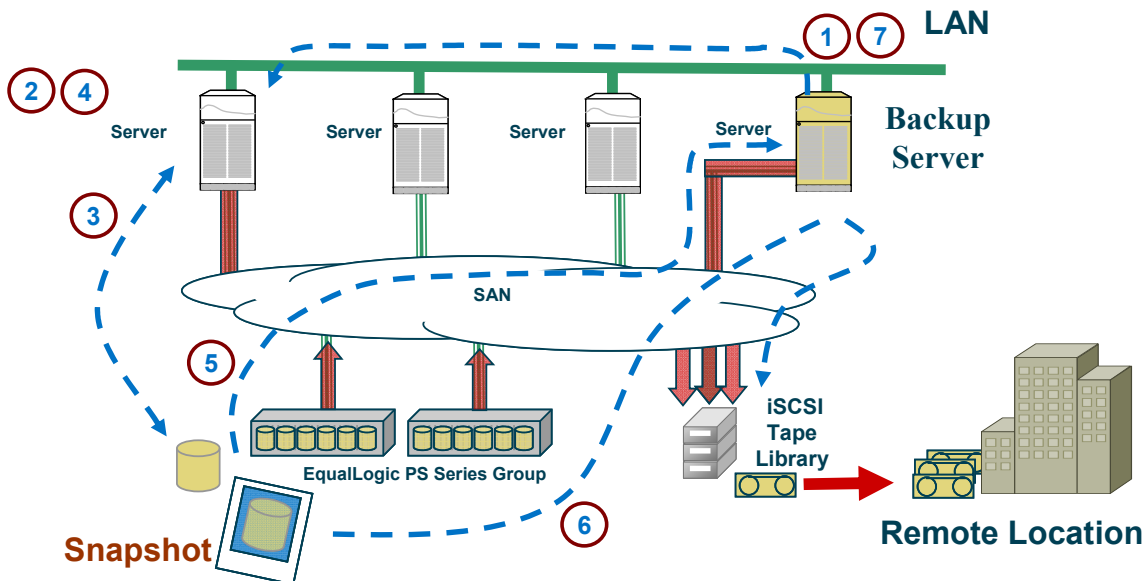
Snapshots in a PS Series group enable you to capture the contents of a volume at a specific point in time and can be used in a variety of ways to enhance recovery. Snapshot creation does not disrupt access to the volume. Like group volumes, snapshots appear on the network as iSCSI targets and can be set online and accessed by hosts with iSCSI initiators. You can also restore a volume from a snapshot or clone a snapshot and create a new volume.

When a volume is restored from a snapshot, the restored volume will contain the volume data that existed at the time the snapshot was created and will have the original volume name and iSCSI target name. The snapshot will still exist after the restore operation. Before the volume is restored, a snapshot of the current volume is automatically created. That way, more recent data is maintained if needed later in the recovery process. This snapshot is named according to the normal rules for naming snapshots; that is, the volume name plus a timestamp plus an identification number.

The scripted backup process typically works as follows:

1. Backup software or script begins the backup process based on policy schedules.
2. Using scripting, the backup server manually tells the server applications to become quiescent and then tells the file system to flush buffer cache and prepare for the snapshot.
3. Using scripting, the backup server tells the PS Series group to create a snapshot. The snapshot is created in seconds.
4. Using scripting, the backup server tells the server applications to resume.
5. Using scripting, the backup server mounts the snapshot.
6. Backup software backs up the mounted snapshot volume.
7. Using scripting, the backup server deletes the snapshot.

Scripted Proxy-Based Backup



One problem with scripting snapshots is that the procedure requires custom scripts for the overall backup processes and also to manually quiesce the application to prepare for the snapshot. Scripting may be unavoidable if operating system awareness for snapshots is not available.

Summary

This technical report provided an overview of backup and recovery architectures and methods and showed that EqualLogic PS Series storage provides an ideal storage solution for backup and recovery. Using snapshot technology and with support for Microsoft Volume Shadow Copy Services, PS Series arrays maintain data integrity during backup operations to ensure that restored data is reliable and quickly accessible with minimal impact on users and applications. It retains multiple copies of data in safe locations, either local or remote and ensures that backup processing has minimal impact on other IT operations.

Documentation and Customer Support

Visit the EqualLogic Customer Support website, where you can download the latest documentation and firmware. You can also view FAQs, the Knowledge Base, and Tech Reports and submit a service request.

EqualLogic PS Series storage array documentation includes the following:

- *Release Notes*. Provides the latest information about PS Series storage arrays and groups.
- *QuickStart*. Describes how to set up the hardware and start using a PS Series storage array.
- *Group Administration*. Describes how to use the Group Manager GUI to manage a PS Series group. This manual provides comprehensive information about product concepts and procedures.
- *CLI Reference*. Describes how to use the Group Manager command line interface to manage a group and individual arrays.
- *Hardware Maintenance*. Provides information on maintaining the PS Series storage array hardware.

To access the Customer Support website, from the EqualLogic website (www.equallogic.com), click **Support** and log in to a support account. If you do not have an account, create one by clicking the link under the login prompt. To contact customer support, send e-mail to supportnp@equallogic.com. If the issue is urgent, call 1-877-887-7337 to speak with a member of the customer support team.