

Merging Backup and Disaster Recovery (DR) Processes

Abstract: For years, disaster recovery (DR) and data protection practices have consisted of two different technology islands — data protection and tape backup. The processes to implement both have usually been separate and handled by different IT departments. Enterprises have attempted to merge the two, but many still use different software products and different processes for replication and tape backup. This paper describes how continuous data protection (CDP) technology can bridge various DR methodologies to help organizations overcome the challenges of tape and achieve maximum benefits from their IT investments.

Introduction

Data protection has focused on the short term, usually for protection within the daily operating periods of the enterprise. Tape backup has been widely used for long-term data retention. Typically, data protected within storage arrays and server systems is passed to tape backup when the capability to provide protection expires.

However, when it comes to disaster recovery (DR), most disasters are not caused by hurricanes, floods, earthquakes, and other natural disasters. Certainly those types of events do occur, and a comprehensive plan to protect and recover from them is required. But the more common disasters are granular problems such as accidental or malicious data deletion, data corruption, user error, viruses, and similar events. Recovery from tape for these types of problems is slow, imprecise, and unacceptable for most business operations.

Tape Backup Challenges

Tape backup is a very reliable and affordable data protection method. Tape is durable and can be easily stored in a climate-controlled environment for long-term protection. Most tape backup implementations use backup agents on application host servers to back up data from the servers.

However, these backup agents consume application host processing time, which becomes a severe problem when virtual server technology is used. In some cases, server backup can impact the server enough to cause significant application performance degradation.

The explosive growth of digital data has significantly expanded server backup requirements to the point where completing backups on time has become increasingly difficult. In addition, online services, such as e-commerce and support service operations, have shrunk the backup window considerably, complicating tape backup and media management.

The integrity of tape backup data has also proven to be very difficult to validate, even if there were time available to conduct tape validation. This means that the reliability of DR from tape backup is questionable at best.

DR Challenges of Tape

As a DR method, tape backup can be quite risky. Data recovery from tape is time-consuming and imprecise. Recovery points — the time gaps between backups and the amount of data that can be lost as a result — can be 24 hours or longer, even with 100% valid data. The time required to restore an application to full operation from a tape backup can vary by an additional two to 24 hours, depending upon how quickly the application's data can be made consistent.

Many tape backups are never validated after backup is completed because there is not enough time to do so. Many enterprises settle for "brute-force" full tape restores of invalidated data rather than trying to find the exact information that needs to be replaced after a disaster.

Recovery from tape backup can extend from hours into days if the tapes required for recovery are stored in an offsite bunker facility. Offsite tape must first be shipped back from the bunker site before the tape restore can begin. The revenue and productivity lost during the recovery period can be devastating to an enterprise.

Many organizations have come to the conclusion that backing up data once a day does not provide enough coverage, and that a recover point objective (RPO) of 24 hours is too large. Yet they live with this risk because the cost of more aggressive replication-based DR protection is too high.

Other enterprises reduce the amount of time required for data recovery by replicating to array-based disk volume snapshots. Array-based snapshots reduce the recovery point by retaining more frequent changed data information since the last data backup.

The limitation of array-based snapshots is that most array vendors support a maximum of 8 or 16 snapshots per volume, which means that when the 9th or 17th snapshot period arrives, the first snapshot is overwritten. At best, typical array-based recovery points with snapshots are taken every one to three hours within a one-day period. These array-based snapshots are frequently connected to application server production volumes and create a negative performance impact when accessed.

Furthermore, array-based snapshot technology does not usually support a continuous journaling mode. That means that the data integrity of array-based snapshot recovery points is limited to the last good snapshot, which may be several hours old.

Improving Tape Backup with Long-term Disk-based Data Protection

The continuous data protection (CDP) concept has emerged as a solution to better protect enterprise data and to improve data recoverability. Nevertheless a CDP solution needs to unify local and remote data protection processes. It should combine replication and backup into a comprehensive data protection solution that provides (1) little or no data loss, (2) rapid data and environment recovery from failure, and (3) freedom from tape recovery as a data protection method.

The FalconStor® Continuous Data Protector (CDP) solution satisfies all of these provisions. FalconStor CDP expands disk-based data protection to significantly enhance snapshot-based DR over an extended period of time. In doing so, FalconStor CDP enhances the effectiveness of tape backup for data protection and reduces the dependence on daily tape backups for DR. With FalconStor CDP, DR policies are less dependent upon tape-based restoration, which improves the overall speed and efficiency.

By reducing or eliminating the need for tape in DR operations, FalconStor CDP enables companies to leverage tape for more suitable tasks, such as backup for long-term data retention and archiving. This enables each technology — magnetic disk and magnetic tape — to perform the role for which it is most optimal.

FalconStor CDP provides numerous DR advantages over tape backup alone:

- Instant recovery from all types of disasters
- Eliminates the tape backup window, reduces tape backups from daily to weekly
- Enhanced remote DR through WAN optimization
- Reduced costs for DR protection
- Eliminates dependence on tape backups for DR
- Physical server to virtual server (P2V) recovery
- Multiple recovery points
- Application data integrity
- Zero-impact backup
- Saves disk space through thin provisioning and space-efficient snapshots
- Non-disruptive DR rehearsals

Key Benefits

Instant Recovery

By reducing dependence on tape backup, FalconStor CDP provides fast recovery times to help organizations meet service level agreements (SLA). It allows users to recover email files and databases in five minutes, recover systems in 10 minutes, and recover entire sites in 30 minutes via remote DR.

Application-aware snapshot agents for popular email systems, databases, file system, and applications (including UNIX, Linux, and Microsoft Windows) ensure 100% transactional integrity and eliminate time-consuming file system integrity checks during system restart.

Eliminating the Tape Backup Window

One of the biggest challenges facing every IT department is the race to complete tape backups in the off-peak time available in which to do them. Even the capability to incrementally update tape backups stretches data center resources because of the reduced window of opportunity.

FalconStor CDP eliminates this dilemma by providing data protection for a much longer period of time than snapshot-supported tape backup, reducing the required frequency of tape backups. Enterprises that implement FalconStor CDP can reduce their tape backup requirements from daily to weekly, or even longer. FalconStor CDP also enables zero-impact backup, because tape backups are generated from the CDP appliances and not from the application server production disk volumes.

By moving to a weekly tape backup schedule, FalconStor CDP eliminates 313 daily tape backups per year, while maintaining instant DR and hourly snapshot protection with or without journaling. In addition, weekly tape backups can be scheduled at any time, not just during a limited backup window.

FalconStor CDP eliminates the need to perform daily backups.

Backup schedule (without FalconStor CDP)

| January | | | | | | |
|---------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| | | | | 1 | 2 | 3 |
| 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 |
| 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| 25 | 26 | 27 | 28 | 29 | 30 | 31 |

| February | | | | | | |
|----------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |
| | | | | | | |

| March | | | | | | |
|-------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |
| 29 | 30 | 31 | | | | |

Backup schedule with FalconStor CDP

| January | | | | | | |
|---------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| | | | | 1 | 2 | 3 |
| 4 | 5 | 6 | 7 | 8 | 9 | 10 |
| 11 | 12 | 13 | 14 | 15 | 16 | 17 |
| 18 | 19 | 20 | 21 | 22 | 23 | 24 |
| 25 | 26 | 27 | 28 | 29 | 30 | 31 |

| February | | | | | | |
|----------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |
| | | | | | | |

| March | | | | | | |
|-------|----|----|----|----|----|----|
| S | M | T | W | T | F | S |
| 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 15 | 16 | 17 | 18 | 19 | 20 | 21 |
| 22 | 23 | 24 | 25 | 26 | 27 | 28 |
| 29 | 30 | 31 | | | | |

= Day that backup is required

Enhanced Remote DR through WAN Optimization

Remote replication to a secondary site is a proven DR method, especially for severe production site damage. The most common method, array-based replication, has several drawbacks in that it requires a dedicated homogenous infrastructure at primary and secondary sites. Disk arrays usually do not communicate across vendor barriers. In addition, array replication is very inefficient in terms of bandwidth consumption.

Failover processed between sites can be complex and error-prone, and failback processes are often under planned and untested. For these reasons, replication is often limited to Tier 1 storage, while Tier 2 and 3 storage data goes unprotected.

FalconStor technology features WAN-optimized replication, which provides as much as a 90% network bandwidth cost savings versus array-based remote replication solutions. Patented FalconStor MicroScan™ technology reduces the amount of data sent by 64 times the amount sent by block-oriented-array replication protocols. FalconStor CDP replication is also able to switch between continuous and periodic replication modes, eliminating replication buffer overflows of array-based and host-based products that cause data inconsistencies.

FalconStor CDP provides the capability to replicate and maintain multiple versions of data at local and remote DR sites for online data retention and recovery. This enables data center operations to achieve 10-minute local full-site or database recovery and 30-minute remote full-site recovery time objectives in the event of a disaster, data corruption, or other problem.

In addition, unified tape backup and disk-based DR enables organizations to replace daily tape backup with periodic disk replication over IP or Fibre Channel (FC). This makes it possible to eliminate the tape backup window and reclaim server compute cycles with zero-impact backups scheduled on a weekly rather than daily basis. As a result, companies can reduce or eliminate the usage of large physical tape libraries in the data center and replace them with smaller units to free up floor space.

About FalconStor

FalconStor Software, Inc. (NASDAQ: FALC), the provider of TOTALLY Open™ Data Protection solutions, delivers the most comprehensive suite of products for data protection and storage virtualization. Based on the award-winning IPStor® platform, products include the industry-leading Virtual Tape Library (VTL) with deduplication, Continuous Data Protector (CDP), File-interface Deduplication System (FDS), and Network Storage Server (NSS), each enabled with WAN-optimized replication for disaster recovery and remote office protection. Our solutions are available from major OEMs and solution providers and are deployed by thousands of customers worldwide, from small businesses to Fortune 1000 enterprises.

For more information, visit www.falconstor.com or contact your local FalconStor representative.

Corporate Headquarters
USA
+1 631 777 5188
salesinfo@falconstor.com

European Headquarters
France
+33 1 39 23 95 50
salesemea@falconstor.com

Asia-Pacific Headquarters
Taiwan
+886 4 2259 1868
salesasia@falconstor.com

FalconStor
Software

Cost Savings

Embedded WAN-optimized replication keeps bandwidth costs to a minimum and accelerates data replication processes between sites for better data recovery points. Application-aware data protection features in FalconStor CDP ensure fast and accurate recovery from data center failures.

Unifying tape backup and disk-based DR processes results in the following financial benefits:

- Unified SLAs for local and remote recovery
- Up to 90% savings on bandwidth required for remote DR
- 86% tape backup operating expense (OPEX) savings by reducing tape backups from daily to weekly
- Up to 50% capital expense (CAPEX) reduction by using Tier 2 storage and eliminating all array-based software licenses
- Up to 20% OPEX reduction by eliminating backup windows and host performance impact

Summary

For years, tape backup was the only means of data protection and recovery. However, as data volumes have grown and dependence on data availability has increased, the limitations of tape media have become all-too-apparent. The FalconStor portfolio of TOTALLY Open™ data protection solutions enhance tape backup reliability, accelerate backup operations, and streamline and consolidate tape management operations. Seamless integration into any storage environment allows organizations to smoothly transition to disk-based backup with no major disruptions or changes to storage infrastructures or policies. Hardware independence provides freedom from expensive vendor lock-in and the ability to seamlessly transition between physical and virtual environments. This allows organizations to make the most of their investments while scaling backup operations, improving data protection and recovery, and reducing costs.