

**A CommVault® Business Value &  
Technology White Paper**

**Snapshot Management &  
Source-side Deduplication are  
Vital to Modern Data Protection**

## Contents

Executive Summary.....	3
New Approaches Needed to Tackle Data Growth.....	3
Snapshot Technology.....	4
CommVault® SnapProtect™ Technology.....	4
A Modern Approach to Snapshot Management.....	6
CommVault® SnapProtect™ Attributes.....	7
CommVault® SnapProtect™ Capabilities.....	7
CommVault® SnapProtect™ Benefits.....	7
Source-side Data Deduplication.....	8
Simpana® Deduplication Benefits.....	8
Simpana® Software—A Better Data Management Solution.....	10
Simpana® Software Benefits.....	10
Conclusion—Modern Data Protection.....	11
About CommVault.....	11



## Executive Summary

The amount of data managed by IT departments continues to accelerate. This data growth and the need for better availability of business-critical data are forcing IT organizations to look at new methods of backup and recovery. Companies are seeking Modern Data Protection; a single solution that enables them to efficiently capture, move, retain, find and recover data from any storage tier. Many have turned to native array-based snapshot technologies, which combine frequent recovery points and minimum operational disruption. However, snapshot management varies significantly by storage platform and often results in more complexity and risk than most organizations bargained for. Others have turned to stand-alone, deduplication approaches that are only loosely integrated with their data protection strategy. This solves an immediate need, but leaves users dealing with the consequences of limited scalability, poor manageability, and rising costs.

CommVault® Simpana® software eliminates the chaos of legacy approaches, enabling users to realize data protection and management benefits, without complexity, and that they never thought possible. CommVault SnapProtect™ technology and source-side deduplication are key capabilities that bring the concept of Modern Data Protection to reality for Simpana 9 users. With SnapProtect technology, companies reduce risk by improving data recovery times and by replacing manual scripting with an automated policy-based management approach. Furthermore, Simpana software integrates snapshots and deduplication to allow resource intensive protection activities (deduplication, encryption, compression, etc) to be offloaded from the production environment, further extending the ability to handle large datasets.

## New Approaches Needed to Tackle Data Growth

Rapid data growth, expanding virtual server deployments, and datacenter consolidation projects are creating massive amounts of data that complicate traditional backup and recovery methods. Applications, databases, e-mail stores and file systems have simply grown beyond the ability of traditional backup software to provide adequate recovery. Aggressive service levels require frequent data copies for recovery, yet traditional streaming methods cannot keep pace with critical business needs. Businesses no longer can afford to take data offline for even eight hours as traditional streaming backup processes consume valuable network and production server resources. These issues are challenging organizations of every size. New approaches are needed for recovery performance, manageability, and reliability—without adding significant cost.



Modern Data Protection meets these challenges head on, finally offering IT departments the justification they need to replace their point solutions and legacy backup and recovery software. But what is Modern Data Protection? Put simply, Modern Data Protection is a new approach to backup and recovery that leverages the integration of innovative and traditional technologies into a single solution, built from the ground up, to enable companies to efficiently capture, move, retain, find and recover data from any storage tier. Modern Data Protection eliminates the chaos of legacy backup software and point solutions and allows users to realize cost, labor, infrastructure, and performance efficiencies.

Two technologies are at the heart of Modern Data Protection. These are array-based snapshot management and source-side, global deduplication. Integrating array-based snapshot technologies with the overall data protection approach improves manageability, recovery and protection while dramatically reducing backup windows. Global, source and/or target-based deduplication is critical to Modern Data Protection because it eliminates redundant data while increasing network efficiency and reducing storage infrastructure costs. Many vendors offer these capabilities in some form, but only one company has delivered them in a way that meets the definition of Modern Data Protection. Let's look at each of these technologies in detail.

### **Snapshot Technology**

Many organizations are replacing their legacy backup and recovery software in favor of array-based snapshot technologies from vendors such as Dell, EMC, HDS, HP, IBM, Fujitsu, and NetApp. The reason often cited is that snapshots allow for more frequent copies and faster restores. However, daily operations become a challenge because snapshot scripting and management are complex. Greater numbers of discrete snapshots/backups make it difficult to enforce consistent backup policies or to quickly locate specific copies from storage tiers. Furthermore, snapshots are typically tied to the production storage array, meaning that failure of the array can lead to complete data loss.

Creating secondary data copies from snapshots can also be problematic when the links between the second copy and the production system break, thus losing application awareness. As a result, recovery from the secondary copy requires complex, manual processes before the data is available to the production system. However, a Modern Data Protection approach ensures the creation of more recovery points and reduces the complexity of application data recovery.

### **CommVault® SnapProtect™ Technology**

CommVault offers SnapProtect technology as the foundation of Modern Data Protection for its customers. SnapProtect technology dramatically improves the ability to meet broader service level commitments, without sacrificing the ability to meet recovery point objectives (RPO) and recovery time objectives (RTO). SnapProtect features make it feasible to increase the number of recovery points kept in any storage tier. This helps increase disk-based retention periods while decreasing restore and recovery times.

SnapProtect technology is designed to work with native storage array snapshot capabilities for even the largest databases and applications, like VMware, Hyper-V, Oracle, SQL, DB2, SAP, Exchange, Sharepoint, etc . It leverages the native snapshot engine built into the storage array to create and manage application-consistent, validated, point-in-time recovery copies.

CommVault's unified software platform allows customers to use a single management console to create and manage snapshot copies in a heterogeneous mix of leading storage arrays; without customization or scripting. This seamless approach facilitates multiple storage tiers using disk, tape, or even the cloud. Data in all storage tiers is application-consistent and can be recovered directly and quickly to the production system with minimal administrator intervention. SnapProtect technology allows users to easily create secondary data copies, migrate them to lower-cost storage tiers including tape and cloud, perform data validation activities, and more importantly, obtain multiple recovery points from heterogeneous storage tiers.

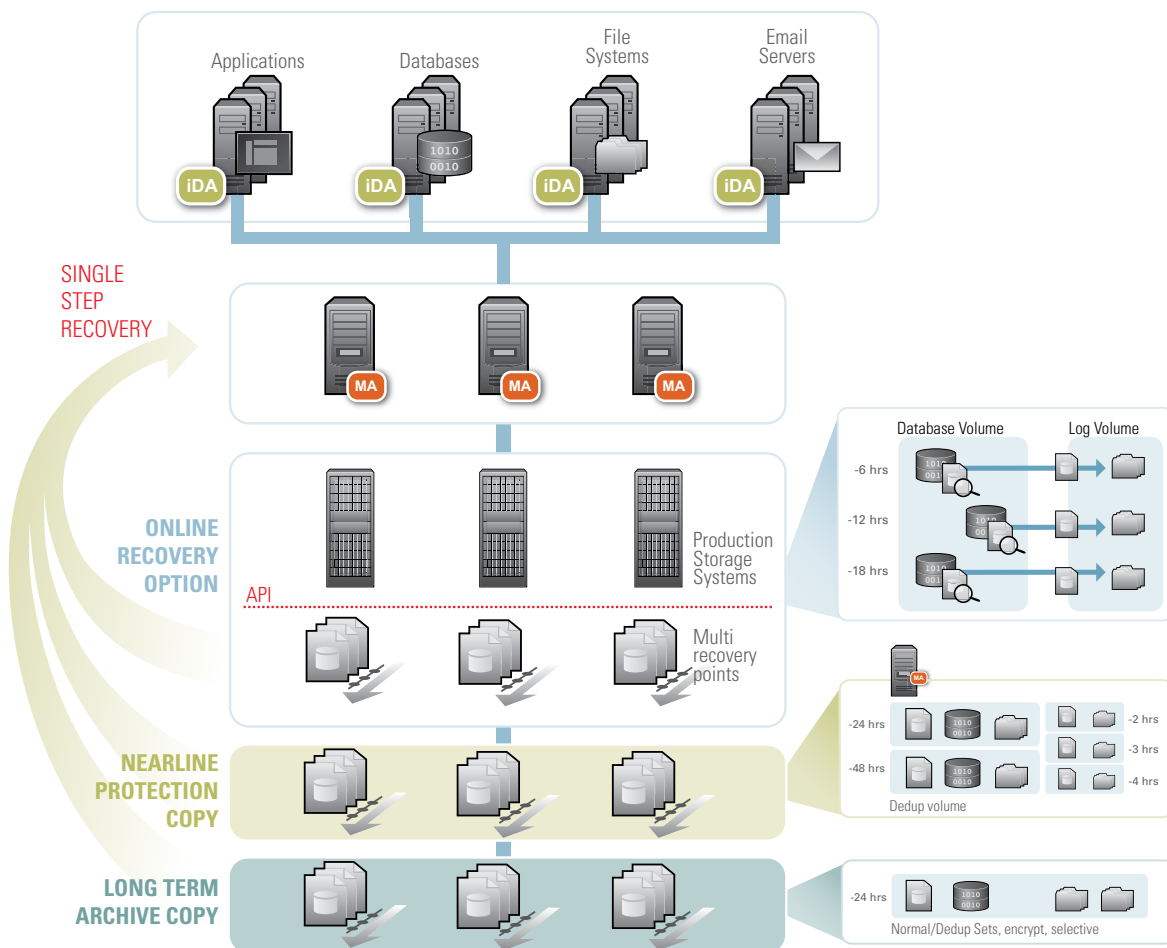


Figure 1: Overview of the Migration and Recovery Process.



## A Modern Approach to Snapshot Management

Modern snapshot integration unlocks the power of native hardware snapshots that customers have already invested in while improving recovery and protection and dramatically improving backup windows. SnapProtect technology lets users meet RPOs with more frequent recovery points and delivers faster and more granular data recovery performance than traditional backup and recovery approaches.

SnapProtect technology leverages native, array-based snapshots using the next-generation data management capabilities in Simpana 9. Unlike traditional streaming backup methods, SnapProtect technology quickly quiesces the application or file system, triggers the snapshot and returns the system to a fully operational state within a few minutes. These snapshots can be mounted on a proxy server and validated as part of a secondary process to ensure the snapshot is an application-consistent copy and readily available for recovery.

Organizations can replace selective, infrequent backups with multiple recovery points each day to reduce both RPO and RTO. This enables you to minimize the impact on production systems, while realizing faster full-file and application recovery. Depending on your configuration, automatic rules for maintenance tasks (for example, proxy validation, log management or snapshot aging) and intelligent storage policies further simplify management. Administrators can select from full application recovery or granular object recovery, all without the manual steps traditionally needed to make the data available to the application, database or file system.

Snapshots can be migrated with one of two approaches. The first uses a proxy server and a secondary copy process to create an independent recovery copy. The second is provided by application-consistent integration with storage-array replication to efficiently move full or incremental block changes from snapshot copies. Both approaches significantly reduce the backup window while offloading server and network resources. They also ensure that data is directly recoverable to the production system or application even if the production storage array is not available. Furthermore, multiple data copies can be created across different storage tiers. Each copy will be recovery-ready, which means that data can be restored directly from any recovery copy to the production system.

By leveraging array-based snapshots and its leading data management and deduplication technologies, Simpana software offers faster full-file or application recovery than legacy backup solutions, minimizes the impact on production systems, and simplifies the task of data management.



### CommVault® SnapProtect™ Attributes

- Use snapshot copies to create multiple tiers of recovery copies with different retention settings, based on SLA and business requirements.
- Integration with leading array-based snapshot technologies to create application-consistent recovery copies with minimal production impact. Integration with Microsoft Volume Shadow Copy Service (VSS) on Microsoft Windows Server platforms to manage protection of Microsoft Exchange Server and SQL Server, as part of a seamless snapshot-based process.
- Support for leading applications such as VMware, Hyper-V, Oracle, SQL, DB2, SAP, Exchange, Sharepoint, etc.
- Support for file systems along with leading operating systems, including Windows, UNIX and Linux.

### CommVault® SnapProtect™ Capabilities

- **Centralized, Unified Management Control**—Works in tandem with leading array-based snapshot technologies to unify protection and recovery operations across systems, applications and locations.
- **Eliminates Scripting**—Utilizes Simpana software's automated policy-based approach to eliminate manual scripting for creating, migrating and aging of application-aware snapshots across storage tiers.
- **Heterogeneous Flexibility**—Enables migration of snapshots across storage tiers and platforms, and allows for direct heterogeneous restores to any system.
- **Application Awareness**—Supports application-aware integration across storage tiers to ensure coherent data handling for leading messaging, database and application systems.

### CommVault® SnapProtect™ Benefits

- Quickly restore entire file systems and databases within moments of failure, without administrator intervention.
- Protect growing mission-critical data stores with minimal impact on production systems, allowing for the creation of more recovery points for reduced RPO/RTO.
- Improve data availability by offloading resource intensive protection activities (deduplication, encryption, compression, etc) from the production hosts.
- Meet recovery SLAs, as well as create multiple recovery tiers, including near-line recovery copies and long-term archive copies. Leverage fully integrated source-side deduplication to improve backup storage and network utilization and store more recovery copies on disk for longer periods.
- Use offline data mining tools for granular recovery of files and database objects from backups.



## Source-side Data Deduplication

Working in tandem with modern snapshot management are the technologies from CommVault that make up Simpana 9's source-side, global deduplication functionality. Source-side deduplication dramatically reduces the resources required to move data as part of the protection and management processes within an organization. Each client that is associated with the same Global Dedupe Database will deduplicate itself against the stored collection so even redundant FULL backups will only move the unique changed blocks to the store.

The new Simpana 9 Dedupe Accelerated Synthetic Hash (DASH) full copy option dramatically changes the operations profile for creating a synthetic full copy set on deduplicated data. DASH full copies allow users to create a synthetic full set by merging all incremental backups to create a new full backup. Since all the data is already sitting on the backup storage device, DASH copy simply updates the appropriate reference counts for the digital hashes associated with deduplicated data. As a result the synthetic full process takes just minutes.

A combination of content-aware and segment-based deduplication with a Global Deduplication Database delivers better results compared to other appliance-based approaches. Using the Deduplication Database, the digital signature of each segment is compared against all the other segments of data in the global storage policy, maximizing deduplication across all backup and archive data. Combining client-based generation of hash signatures with integrated data compression and data encryption capabilities minimizes network bandwidth requirements and maximizes end-to-end network speed and security that goes all the way from clients to disk, tape, and cloud-based storage tiers.

With Simpana software, you can expand the use of data deduplication across all tiers of storage, including tape, without the need to "rehydrate" the data back into full form. Using an automated, policy-based approach, Simpana software users can implement a hands-free approach to migrate and manage data quickly and efficiently. Furthermore, migration is invisible to end-users and ensures that the deduplicated data is brought back seamlessly from any storage tier. CommVault deduplication also combines a shared index/catalog and unique deduplication to tape so users can easily track and retrieve only the data required, without the need for lengthy restores and searches through full data volumes.

## Simpana® Deduplication Benefits

- **Reduce Storage Consumption**—Reduce storage needs by up to 90 percent. Eliminate redundant backup and archive data by starting at the source. Each source, whether server, laptop, or storage pool, uses a two-part process to eliminate duplicate data from its own data set, then from all other sources using a global policy.

- **Optimize Network Infrastructure**—Source-side deduplication typically moves only one-tenth the data of target-based approaches. This helps enable practical protection of remote sites and a mobile workforce.
- **Scale Out Easily**—Protect PBs of backup and archive data using a modular, extendable approach that eliminates the need for re-hydration and re-deduplication.
- **Deduplicate Globally, Retain Locally**—Deduplicate across many locations and multiple storage policies. Unique copies with independent retention policies can be created on disk or tape and stored onsite, at a central location, or even in the cloud.
- **Begin Quickly**—With just a few simple clicks, Simpana 9 users can begin to realize the benefits of deduplication.

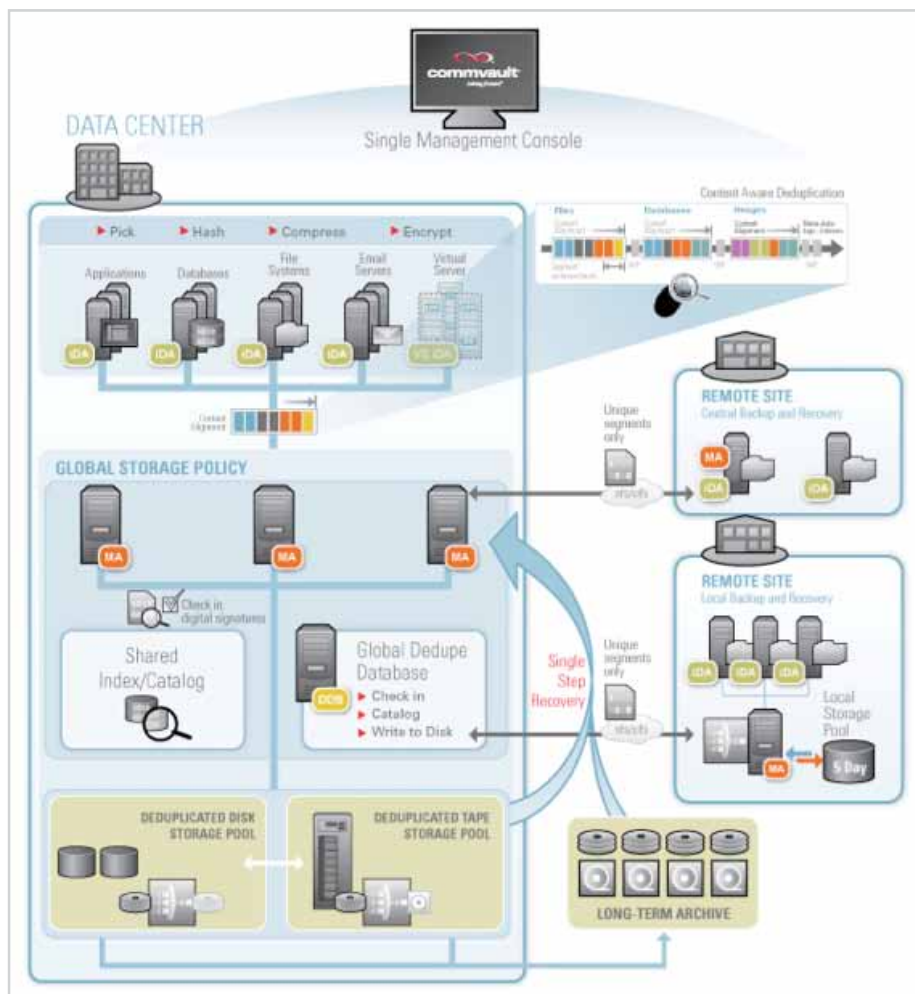


Figure 2 Embedded Global Data Deduplication

## Simpana® Software—A Better Data Management Solution

CommVault Simpana software is a single, efficient, scalable platform that supports a complete range of data management functions. Offering an expandable range of application modules, Simpana software provides multiple backup and recovery options; a flexible IT infrastructure for data protection across physical and virtual servers; and fast, cost-effective disaster recovery to ensure business continuity. CommVault Simpana software includes data archiving, backup and recovery, data reporting and resource management and replication and snapshot management. When used together, these capabilities deliver the unique and compelling benefits of Modern Data Protection, and it's easy because of the single platform and single console architecture.

### Simpana® Software Benefits

Simpana software provides a single console from which to easily manage all aspects of data protection, archive management, recovery management and replication of physical and virtual servers. Benefits of Simpana software include:

- Simplifies and expands data protection recovery options
- Facilitates implementation of tiered storage
- Increases application and data availability
- Enables cost-effective, consolidated backup of virtual servers
- Improves backup efficiency and reliability
- Enables hot backups of virtual machines
- Shares storage for backup and archiving across physical and virtual servers
- Provides one interface to learn for data management




Figure 3—CommVault® Simpana® Software's Singular Information Management Architecture.



## Conclusion—Modern Data Protection

As key elements of a Modern Data Protection strategy defined by CommVault, SnapProtect technology and source-side, global deduplication work together to enhance data backup and recovery capabilities by simplifying and seamlessly automating snapshot management while, at the same time, reducing the resources needed to move and store data. With the downtime of business-critical applications kept to a minimum, organizations increase productivity and reap the benefits of rapid backup and recovery. Management complexity is reduced while less time and fewer resources are required to oversee protection activities or to manage existing storage systems. This, in turn, reduces or delays the need for additional infrastructure, enables headcount to be re-deployed to other value-added activities, and shrinks the impact of data management on IT budgets.

## About CommVault

A singular vision—a belief in a better way to address current and future data and information management needs—guides CommVault in the development of Singular Information Management® solutions for high performance data protection, universal availability and simplified management of data on complex storage networks. The CommVault exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk. CommVault Simpana software modules were designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative Backup & Recovery, Archive, Replication, Search and Resource Management capabilities. More companies every day join those who have discovered the unparalleled efficiency, performance, reliability and control only CommVault can offer. Information about CommVault is available at [www.commvault.com](http://www.commvault.com). CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States.

For More Information on Modern Data Protection:

<http://www.commvault.com/simpana.html>

<http://www.commvault.com/solutions-snapprotect.html>

<http://www.commvault.com/solutions-deduplication.html>



[www.commvault.com](http://www.commvault.com) ■ 888.746.3849 ■ E-mail: [info@commvault.com](mailto:info@commvault.com)

CommVault Worldwide Headquarters ■ 2 Crescent Place ■ Oceanport, NJ 07757 ■ 888-746-3849 ■ Fax: 732-870-4525

CommVault Regional Offices: United States ■ Europe ■ Middle East & Africa ■ Asia-Pacific ■ Latin America & Caribbean ■ Canada ■ India ■ Oceania

©1999-2011 CommVault Systems, Inc. All rights reserved. CommVault, CommVault and logo, the "CV" logo, CommVault Systems, Solving Forward, SIM, Singular Information Management, Simpana, CommVault Galaxy, Unified Data Management, QiNetix, Quick Recovery, QR, CommNet, GridStor, Vault Tracker, InnerVault, QuickSnap, QSnap, SnapProtect, Recovery Director, CommServe, CommCell, ROMS and CommValue are trademarks or registered trademarks of CommVault Systems, Inc. All other third party brands, products, service names, trademarks, or registered service marks are the property of and used to identify the products or services of their respective owners. All specifications are subject to change without notice.