

Dell PowerVault DL2000 Powered by CommVault

A Dell White Paper

CENTRALIZED DATA
MANAGEMENT FOR
REMOTE AND BRANCH
OFFICE (ROBO)
ENVIRONMENTS



CONTENTS

EXECUTIVE SUMMARY	3
RETURN ON INVESTMENT OF CENTRALIZING DATA MANAGEMENT	3
APPLICATION-AWARENESS ENSURES DATA COHERENCY	4
FLEXIBLE SOLUTION OPTIONS THAT FIT ANY REMOTE SITE	5
BRANCH OFFICE CONSOLIDATION WITH COMMVAULT SIMPANA REPLICATION	6
CONCLUSION	7

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 –POWERED BY COMMVAULT SIMPANA REPLICATION

EXECUTIVE SUMMARY

The proliferation of remote and branch office environments continues to increase pressure on IT infrastructures and teams. No matter your organization's size, you likely need to include remote site management as part of your plan: it's common for large organizations to have 100 or more remote sites; most mid-sized organizations have somewhere between 25 and 100 remote sites; even small organizations support remote systems that may be used in stand-alone, small office or home office settings.

While remote sites proliferate, the requirements for capturing, protecting and preserving remote data have also increased. Data recovery is required not only to support day-to-day business operations, but preservation of remote site data including email and end-user files is a legal requirement for organizations large and small. But managing remote data is challenging and expensive, and places a heavy management burden on already overworked IT personnel.

Most organizations have continued to rely on antiquated tape-based approaches to remote site data management. These approaches are slow, expensive, and unreliable. They also require local office personnel to perform manual tasks, such as rotation of tape cartridges and data recovery, for which the local personnel typically are not well trained or well suited. Preserving data for disaster recovery from these sites requires rotation of tapes to offsite storage locations, most often by high-cost services companies who charge to send trucks for tape pickup from every site. Sometimes tapes are rotated offsite in the trunks of employees' cars. In effect, remote data is at risk and unprotected in most cases while also being a cost drain to shrinking IT budgets.

A solution which has been available for some time but not often deployed, is to centralize data for management from remote offices across Wide Area Network (WAN) connections. Rather than attempting to manage data at remote sites where data management infrastructure is insufficient, centralization brings that data to datacenter locations where the infrastructure is robust. The cost savings are enormous: retiring backup equipment and service contracts from remote sites is an immediate benefit, while enabling local personnel to stay focused on their work rather than being distracted with backup typically raises productivity. The benefits to data management are similarly profound: datacenter locations more often stay current with modern data management infrastructure and can, therefore, more reliably protect, manage and preserve data and usually can also restore it faster even over remote network connections back to remote sites and systems.

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

However, until now, this type of approach has been too complex for many organizations to deploy. In particular, IT teams with limited resources have been reluctant to change what they have done to protect remote site data because of concerns with network costs and security of the data, and because adding replication software required to centralize data from remote sites has not been simple enough.

The Dell DL2000 powered by CommVault Simpana replication software solves the complexity issue for all types of organizations, large and small, and offers a method for centralizing data from remote sites that is simple, network-efficient, and secure. Dell has solved the complexity problem by offering a drop-in storage device with all of the capabilities required to replicate data over WAN for any end-user file, email or application data stored on the storage system. Dell also preserves the flexibility to deploy a software agent for replication from smaller sites, where adding storage may not be desirable. This solution also ensures network efficiency and data security over networks, with the use of CommVault Simpana replication which combines data compression, encryption and deduplication.

For virtually the first time, data centralization can be simple, efficient and secure for all organizations and remote sites, large and small. The Dell DL2000 powered by CommVault Simpana replication will help to accelerate company initiatives to capture, protect and preserve all data through the adoption of disk-based initiatives and tape-elimination strategies – even for data on the edge of their organizations.

COST SAVINGS OF DATA CENTRALIZATION	INCREASED PRODUCTIVITY AND SERVICE LEVELS
Remove and retire backup equipment from remote sites, saving in maintenance and upgrades	Remote and branch office systems run faster and with no backup impact, supporting higher productivity levels for business workers
Repurposes backup server systems to other workloads	Efficient use of networks through combined effect of data reduction from compression, encryption and deduplication
Eliminate the cost of tape cartridges from remote sites	Ability to enforce common policies for data management across all data, including data created in remote sites
Eliminate service contracts requiring tape collection from each remote site	Ability to report on management of all data, including centralized data, for executive oversight and validation
Eliminate or reduce human error in data management	Higher levels of reliability for preservation of data for disaster recovery and legal discovery purposes
Remove the burden of tape rotation and oversight of backup/recovery from remote sites	Higher levels of data reliability and reduced cost of remote office data protection

BRANCH OFFICE CONSOLIDATION WITH THE POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

RETURN ON INVESTMENT OF CENTRALIZING DATA MANAGEMENT

The combination of backup-to-disk, data deduplication and remote replication solves the typical problems associated with managing data at remote sites. The Dell PowerVault DL2000 powered by Commvault Simpana replication software offers high value by dramatically reducing bandwidth requirements to move data over the WAN from remote sites to a centralized data center. Utilizing deduplication, bandwidth costs are significantly reduced because less data is sent over the WAN. Data is protected much quicker than if all of the backup data was to be sent over the WAN (this could take several hours, even several days).

Because the CommVault Simpana platform comes standard with the DL2000, combining CommVault Simpana replication with the DL2000 is as simple as enabling an add-on module. All replication features are administered through the same graphical user interface. As a result, backup data can be reduced in size with deduplication so that network-based replication can be used as an alternative to tape-based off-site archival.

Unlike other disk-to-disk appliances that require their own proprietary solutions at the centralized location, CommVault Simpana replication can be used in combination with any disk hardware. Data stored on a DL2000 in a remote office can be replicated back to any centralized disk array (for example, an additional DL2000 is not required at the centralized site). Also, remote offices with alternative storage can keep that storage and use CommVault Simpana software agents to centralize data onto the DL2000. In this way, Dell offers a solution with data migration capabilities for moving data among storage devices to make it faster and easier to adopt the replication strategy with as little need to change the existing systems environment as may be desired.

Eliminate the Cost and Burden of Managing Data in Remote Environments

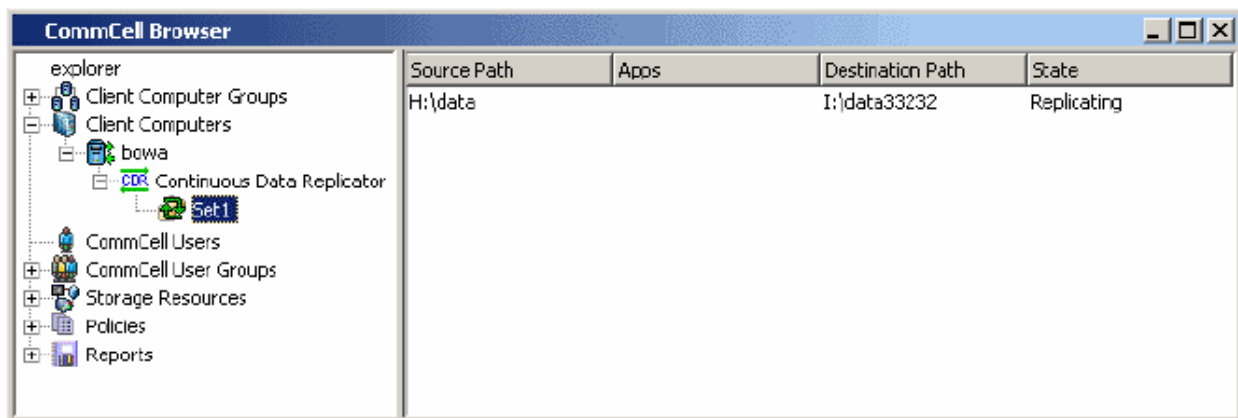
Eliminating backup infrastructure and burden from remote site environments immediately lowers costs. However, to achieve this requires comprehensive management capabilities for oversight, reporting and recovery all from a central location. CommVault Simpana software includes comprehensive remote management capabilities designed to simplify the centralized management of remote site data.

Example of an incomplete solution: Deployment of data replication to centralize management is not a sufficient replacement for remote backup. Knowledge of the state of the data is critical, so that data protected is coherent and usable at time of recovery. Knowledge of the client systems to which data belongs is critical, so that files and email can be found and efficiently returned to remote systems without having to restore all data for a branch – or for every branch. Reporting systems must also be comprehensive, so that if and when problems occur in backup then administrators will have sufficient tools to know that data is at risk and prioritize troubleshooting to find and fix the problem.

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

Example of a comprehensive solution: When sufficient knowledge of the remote clients and applications is built into the replication stream, then comprehensive management of centralized data is possible. CommVault Simpana software tracks replication to individual clients and applications, and includes these systems and data in standard backup and recovery reports generated at the central location. CommVault Simpana's unified console provides a single view of all systems, and enables IT administrators and operators to work with replicated data using the same procedures and policies that are applied to other data located at the datacenter.

Reporting, management control and the use of common backup policies are some of the key capabilities built into the CommVault Simpana solution to provide a comprehensive solution for data centralization and management.



CommVault Simpana Centralized Management Console for Replication and Comprehensive Data Management

More Reliable Data Protection and Recovery; Lower Cost Data Preservation

When comprehensive management tools are in use, then IT administrators can be assured of comprehensive protection and adherence to policies for data centralized from remote site locations. The use of datacenter infrastructure, which often includes more modern disk-based approaches to data management than what is found in remote site locations, improves the reliability and performance of data management tasks. Performing these tasks on a secondary copy of the data removed from the remote site also speeds performance of systems at the remote site which are no longer weighted down by intrusive backup operations.

Data managed using a tiered combination of disk and tape, or into an infrastructure from which tape has been eliminated completely, also lowers the cost of data preservation.

Enabling IT administrators and operators whom have been trained for data management to manage remote data is also a benefit. Organizations no longer need to wonder if backup tasks have been completed, tape cartridges rotated, or data moved to offsite locations. They can also benefit from the faster recovery that trained administrators and operators can deliver when they work with comprehensive remote management tools, such as what CommVault Simpana software delivers.

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

How does Deduplication work with Remote Replication?

CommVault Simpana software is built from a common set of re-usable services which have been optimized for modern data management. These include advanced data backup and recovery, archive, encryption, compression and deduplication. Each individual service and capability offers its own compelling advantage. In combination, the benefit for data management is compelling.

Organizations planning for data centralization from remote sites and systems are concerned with network efficiency. To address these concerns, CommVault Simpana software offers a full range of capabilities designed to ensure that only the data which must be replicated is sent across Wide Area Network connections. By combining advanced capabilities for backup and archive together with compression, encryption and deduplication, CommVault Simpana software provides better efficiencies for transmission of any type of data across WAN connections than is otherwise achievable.

The effect of efficient backup: CommVault Simpana backup software includes capabilities for selecting incremental and differential data for protection. Synthetic Full capabilities make this even more efficient, and is often suitable for protecting data across WAN connections when the rate of change on remote data is low.

The effect of archive: CommVault Simpana archive capabilities further reduce the burden on data backup by removing static and less frequently used data from primary stores. Removing this data to secondary storage not only frees up the primary disk and speeds application performance, it also reduces the amount of data that must be backed up, preserved nightly or weekly, and rotated offsite for disaster recovery.

The effect of compression: Depending on the type of data, compression algorithms can often reduce the size of the data being sent over networks – sometimes more than deduplication. CommVault Simpana software offers the ability to combine use of compression with deduplication, for maximum network efficiency for all types of data.

The effect of encryption: CommVault Simpana software includes advanced encryption algorithms including AES 256, Blowfish and other popular encryption methods. These are offered for use in combination with data compression and deduplication, to ensure that data sent across WAN connections does not present a security risk.

The effect of deduplication: CommVault Simpana offers File Deduplication capabilities for file, document and email attachment deduplication. When applied to data at remote sites, deduplication removes any duplicate data before that data is sent on WAN connections. Uniquely designed for use with compression and encryption, CommVault Simpana deduplication offers better efficiencies for transmission of any type of data across WAN connections than is otherwise achievable.

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

APPLICATION-AWARENESS ENSURES DATA COHERENCY

For supported applications, Continuous Data Replicator (CDR) will assist in configuring Replication Pair content by discovering directories that should be replicated (logs, databases, and so forth) to protect the application's data. After data for these applications has been replicated, you can browse and access the Live Copy of the data.

Windows File System Awareness

- File Allocation Table (FAT) file systems
- New Technology File Systems (NTFS)
- Access Control Lists (ACL)
- Unicode files
- Sparse files
- Single Instance Storage
- Mount Points
- Encrypted files - the key is not replicated, only the file itself
- Shared Volumes
- Compressed Data - including a compressed file contained in a compressed folder, any folder (not compressed) that contains two or more compressed files, or a compressed drive.
- Dynamic volumes are supported as Replication Pair content
- Software and Hardware RAID storage
- Migration files (stub only)

UNIX File System Awareness

- Logical volumes for AIX and Linux; only Linux supports file systems on physical disks configured as CXBF devices
- Access Control Lists (ACL)
- Unicode files; files with non-ASCII characters in their name
- Sparse files (during data replication phase only; see Note below)
- Migration files (stub only)
- Linux only:
 - Extent 2 File System (ext2)
 - Extent 3 File System (ext3)
 - Reiser File System (reiserfs)
- AIX only:
 - Enhanced Journaled File System (JFS2)

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPARA REPLICATION

CONTINUOUS DATA REPLICATOR (CDR) RECOVERY POINTS

High availability of protected data in a consistent state is also available by creating Recovery Points, through the use of snapshots created on the destination machine. These Recovery Points can be mounted, shared as a volume on the network, or recovered using Copyback. After a Recovery Point has been created, the data can also be backed up, using the appropriate File System /DataAgent. This provides additional protection for your data, as well as additional options for recovery of your data.

Consistent Recovery Points, which define a point-in-time where application data is in a consistent state on the source computer (the application server) can be created using CDR, ensuring the application data can be restored to that point-in-time. Recovery Points can also be created directly from a snapshot of the replication volume on the destination computer, but this does not preserve application data in a consistent state, and thus is more useful for file system data. Recovery Points can be created, retained, backed up, mounted, and for Windows can also be shared, and thus made available to users on the network. This allows data to be protected, accessed, and recovered in a variety of ways.

BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERSHIELD DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

FLEXIBLE SOLUTION OPTIONS THAT FIT ANY REMOTE SITE

Data Replication is the process of copying specified, file-level content from one computer, the source computer, to another, the destination computer. This is achieved through an initial transfer of the specified data, after which the replicated copy is kept updated in *nearly* real time with any changes that are made to the data on the source computer. This replicated copy on the destination computer provides on-going, nearly-real-time disaster recovery protection for the source computer, unlike most data protection solutions which require significant time to perform a complete data protection operation. In addition, data replication provides a basis for additional data protection activities, such as Recovery Points (snapshots) and backups of Recovery Points.

Replication Scenarios

Several common scenarios for data replication are illustrated below, but these by no means illustrate all of the ways in which data replication can be configured.

- Replication from one Source computer to one Destination computer
- Fan-In: Replication from multiple Source computers to a single Destination computer
- Fan-Out: Replication from one Source computer to multiple Destination computers

One-to-One Data Replication

This is the most fundamental configuration for data replication. A single computer on the LAN or WAN has its data replicated to another computer, either local or remote. This provides protection of the source computer against catastrophic failure of the computer itself.

Many-to-One Data Replication

In this configuration for data replication, referred to as "Fan-In", multiple computers on the LAN or WAN have their data replicated to a single computer, either local or remote. This provides protection of all of the source computers against catastrophic failure, while maximizing the use of resources by directing all the data to a single destination computer.

One-to-Many Data Replication

This configuration for data replication, referred to as "Fan-Out", adds significantly to the protection afforded to the source computer, because of the redundancy. A single computer on the LAN or WAN has its data replicated to multiple computers, any of which can be either local or remote. This provides protection against catastrophic failure of an entire site, as well as the source computer itself.

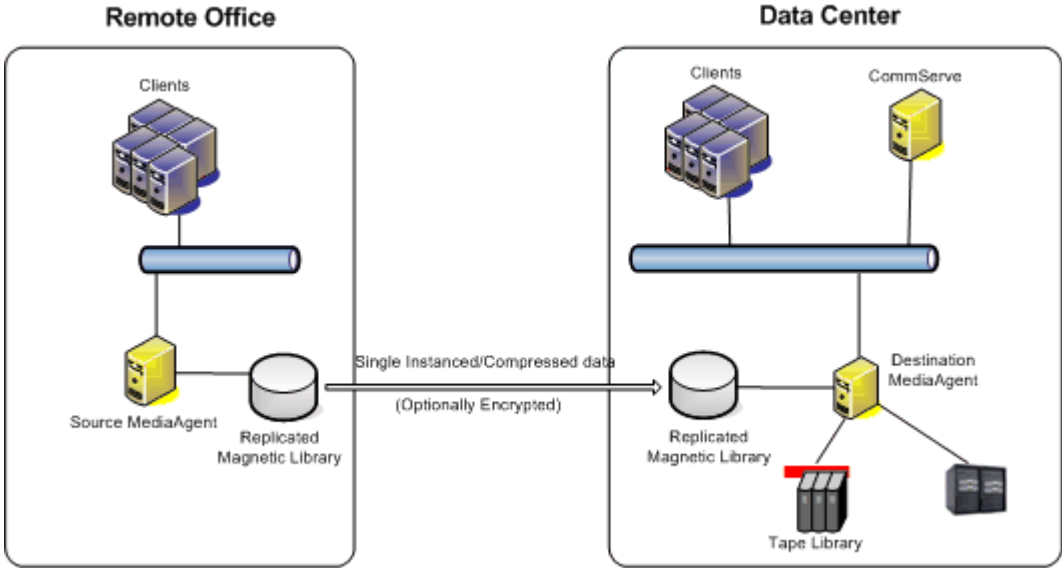
BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 – POWERED BY COMMVAULT SIMPANA REPLICATION

The Remote Office Backup solution is designed to seamlessly replicate data from remote offices to a centralized data center, where it is stored on a magnetic library (disk) for ready access. Consolidating all backup data at the centralized data center provides redundancy for disaster recovery, as well as an alternate source for normal data recovery operations. This innovative solution is easily deployed, without any disruption of current data protection operations, and with no impact on your end users at remote offices; no additional software or configuration is required on client computers, as only MediaAgents are involved. Even more importantly, the data from every iDataAgent deployed at remote offices can be replicated to the centralized data center, no matter the source of the data -- all supported file systems, databases, and data types.

Significant benefits result when this solution is combined with several other available features; de-duplication at the remote office is performed before data is replicated to the centralized data center, greatly reducing the network load. Replication Compression provides a similar benefit, compressing data at the remote office, replicating it across the network to the centralized data center, then uncompressing it on the destination MediaAgent; finally, Replication Encryption can be employed to ensure the security of data transmitted across non-secure networks, by encrypting the data before transmission and automatically decrypting it on the destination MediaAgent at the centralized data center.

Backup and recovery of company data at primary sites typically places high demands on the resources of a company's IT department. The result is that data protection is neglected at the smaller, remote sites even though information at these regional offices or branch offices may be as important as that found at the primary data center. A lack of technical resources at remote sites further complicates the problem. Even if standard backup-disk technologies are implemented at remote sites, the data continues to sit there unless a tape DR strategy is implemented because there is simply too much data to reasonably replicate back to a centralized site for DR purposes.

**BRANCH OFFICE CONSOLIDATION WITH THE DELL POWERVAULT DL2000 –
POWERED BY COMMVAULT SIMPANA REPLICATION**



A typical Remote Office Backup solution is shown in the illustration. Software is deployed on a MediaAgent located at a remote office, and a MediaAgent located at the centralized data center. Thus, there is no impact on end users and client computers, since all operations involve MediaAgents only. Deployment and configuration occur once, and the solution continues to function without any further operator intervention. Once the initial transfer of data has been completed at the time of deployment, all data replication after that point will involve only data that has changed on the MediaAgent, similar in concept to Incremental Backups. This is different from an Auxiliary Copy operation, which is similar in concept to Full Backups, copying all selected data every time it runs, and also expanding all Single Instanced data each time.

CONCLUSION

Working with CommVault, Dell has created an integrated storage solution with replication capabilities that enable better management for data from remote and branch office sites. Centralizing this data for management enables backup systems and devices to be retired from branch environments, along with maintenance contracts and tape rotation services, for immediate dramatic cost savings. Productivity is raised in the branch, because backup operations are no longer performed locally; neither local personnel nor systems are distracted nor impacted during backup. Human error is reduced because datacenter teams are better trained and managed through backup and recovery. Data is better protected, more quickly and reliably, and with usually with more cost-effective options for immediate and long-term data preservation. Policies for data management can be more consistently applied, an important consideration for corporate oversight and governance along with legal requirements in the event of discovery actions.

What makes these advantages possible is the simplicity and flexibility of the Dell PowerVault DL2000 – Powered by CommVault Simpana replication software. Offering a storage device as a single point of replication to drop into a remote site dramatically simplifies the solution. Dell also offers the flexibility of deploying just the CommVault Simpana replication agent, for smaller environments or environments in which a new storage device is not desirable. The Dell PowerVault DL2000 is also scalable to serve as the centralization point for many-to-one fan-in of replication streams coming from remote sites and systems. Dell also preserves CommVault Simpana's capability to work with virtually any disk system, for maximum flexibility and to preserve the investment in existing storage systems.

ABOUT DELL

Dell Inc. (NASDAQ: DELL) listens to customers and delivers innovative technology and services they trust and value. Uniquely enabled by its direct business model, Dell is a leading global systems and services company and No. 34 on the Fortune 500. For more information, visit www.dell.com, or to communicate directly with Dell via a variety of online channels, go to www.dell.com/conversations. To get Dell news direct, visit www.dell.com/RSS.

ABOUT COMMVAULT

A singular vision—a belief in a better way to address current and future data 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. CommVault's exclusive single-platform architecture gives companies unprecedented control over data growth, costs and risk. CommVault's Simpana software suite of products was designed to work together seamlessly from the ground up, sharing a single code and common function set, to deliver superlative data protection, 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. CommVault's corporate headquarters is located in Oceanport, New Jersey in the United States.