Business Continuity Solutions for Microsoft Exchange using HP Network Attached Storage (NAS) built on Microsoft Windows® Storage Server 2003 and CommVault® QiNetix™ Software

Microsoft Windows Storage Server 2003 | CommVault Systems | hp invent

Abstract

The combination of HP StorageWorks network attached storage (NAS) products, built on Microsoft® Windows® Storage Server 2003, and including QiNetix software from CommVault Systems, offers a commanding set of storage management tools to assist in the delivery of business continuity solutions for Microsoft Exchange environments. The solution provides the protection of consolidated Exchange Server 2003 data and file server data hosted on HP NAS, utilization of HP NAS as a target for Disk-to-Disk-to-Tape (D2D2T) backup for existing Exchange servers, and deployment of HP NAS as a target for Exchange e-mail, delivering on the HP Information Lifecycle Management (ILM) strategy for active, time-based management of data, from creation to deletion, based on application and business needs. These solutions leverage the strength of industry-leading storage products from HP, the reliability of HP ProLiant-based technology, and CommVault's comprehensive data management software brought together by Microsoft Windows Storage Server operating system to meet the needs of today's businesses.

Introduction

Windows-powered HP NAS makes an ideal platform for both primary and secondary (backup/archive) storage for Microsoft Exchange e-mail data. HP StorageWorks NAS devices provide low acquisition and ownership cost and high scalability. These benefits, coupled with the HP StorageWorks high-speed storage for primary Exchange data stores and low-cost bulk storage for secondary Exchange backup/archive storage, provide a compelling platform for Exchange data management solutions. Combining this HP NAS platform with the "best-of-breed" data and storage management software from CommVault Systems creates a total solution for the management of Exchange e-mail and the environment it is hosted in. Why focus on e-mail? Consider the following:

- The United States generates more than 9 billion e-mail messages daily.
- IDC projects 1.2 billion global e-mail mailboxes by 2005.
- E-mail contains 60% of a business's critical information.

However, most e-mail is archived on end-users' hard drives, and e-mail content management ranges from primitive to non-existent.

E-mail and databases house the vast majority of critical business information. The loss of just a few important messages, or the discovery of incriminating ones, can impact a business' viability, operations, employee communications, customer relations, and e-commerce. So given the importance of e-mail, how do you manage it? How do you archive and restore e-mail messages? And how do you manage content?

CommVault's QiNetix software suite for HP NAS is composed of three primary solution components—Galaxy Backup & Recovery, QiNetix DataMigrator, and QiNetix DataArchiver—that answer these questions and focus on the protection, availability, and accessibility of a company's e-mail resources.

Protecting consolidated Exchange and file server data on NAS

HP StorageWorks NAS products enable you to store both file server data and Exchange 2003 data stores on the same NAS device. This ability is facilitated by Microsoft Windows Storage Server 2003 Feature Pack for Exchange storage, which delivers no-cost functionality for HP NAS customers running the Windows Storage Server operating system. With business-critical data being consolidated on HP NAS, effective and reliable data protection solutions are essential. CommVault Galaxy backup software provides the kind of backup and restore capabilities required for businesses running Microsoft Exchange.

Disk-to-NAS-to-Tape protection of Exchange servers

Continuous access to business-critical data stored in Microsoft Exchange servers is a well-recognized need. In the event of data loss, quick and immediate recovery and access to data are essential. Solutions that enable this are mandatory in today's business world. The use of disk as a backup medium for data protection has obvious benefits. With the availability of HP StorageWorks Serial ATA (SATA) disks supported by HP NAS, this option is now a reality. The low total cost of ownership (TCO), ease of deployment, and scalability of HP NAS make it an ideal platform for disk-based data storage solutions. Coupled with CommVault Galaxy software optimized for disk storage, HP NAS provides an ideal solution for D2D2T backup—or in this case, Disk-to-NAS-to-Tape (D2N2T) deployments.

NAS as an Exchange e-mail information lifecycle target

This solution builds on the concepts and values presented in the previous, "Disk-to-NAS-to-Tape for Exchange Servers" solution. It leverages the same strengths of the Windows-powered HP NAS platform and CommVault's QiNetix software. The result is a complete information lifecycle management solution, providing capabilities ranging from basic e-mail lifecycle management to regulatory compliance solutions.

The HP award-winning NAS products built on Windows Storage Server, coupled with CommVault's own award-winning software, create solutions that solve the complex challenges facing businesses today. The strength of the partnership is clear when you consider that CommVault is a Microsoft Gold Certified Partner, and their Galaxy software is the Microsoft System Architecture (MSA) recommended backup and recovery solution for Microsoft-based operating systems.

Solution: protecting Exchange and file server data consolidation on NAS

With the introduction of a new Windows Storage Server 2003 Feature Pack, it is now possible to host Microsoft Exchange 2003 data stores on HP NAS devices, built on Microsoft Windows Storage Server 2003. The use of NAS storage for Exchange data has benefits over alternative storage area network (SAN) deployments in the areas of flexibility, cost-effectiveness, and ease of deployment. These benefits are easily realized in small and medium business environments, where server consolidation and branch office business require solutions with low TCO. It is imperative that such valuable data is protected from accidental loss and that a comprehensive plan for data continuity is called for. HP has partnered with CommVault Systems to provide a robust data protection and business continuity solution in these environments.

Before looking at the data protection needs of this scenario, a more in-depth look at the benefits for such deployments is helpful. Deploying HP NAS as an Exchange storage platform delivers three key benefits when compared to traditional SAN deployments:

- **Flexibility:** HP NAS allows you to pool storage for application, file, and Exchange use without the complexity of SAN solutions. Consolidating storage from server-based direct attached storage (DAS) onto Windows Storage Server hardware makes it easier to allocate and reallocate storage as business requirements dictate, something that is difficult with DAS systems. With sufficient storage, deploying Windows Storage Server 2003 allows consolidation of Windows NT® 4.0 and Windows 2000 file and print services with Exchange 5.5 and Exchange 2000 databases when migrated onto one set of storage hardware.
- **Cost-effectiveness:** Because the cost of HP NAS running on Windows Storage Server can be amortized across Exchange and file services, economies of scale and consolidation of hardware combine to make Windows Storage Server 2003 extremely price competitive on a per-megabyte basis. Windows Storage Server 2003 is cost-effective enough to deploy in places like branch offices that could not normally justify their own SANs or DAS arrays.
- **Ease of deployment:** HP NAS products provide seamless integration with Exchange and Active Directory. They can be managed and updated with the rest of your Windows infrastructure; they use familiar Windows tools and utilities, and they are supported both by Microsoft and its server hardware partners.

For additional details of how the Windows Storage Server 2003 Feature Pack delivers these capabilities for Exchange while still maintaining excellent performance, robustness, and cost effectiveness, see:

http://www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/exchange_fp.mspx

Possible deployment

The following diagram (Figure 1) shows a possible deployment of hardware and software components. Backup resources such as tape drives and libraries may be connected directly to the NAS device or to a backup server elsewhere in the environment.

Figure 1.

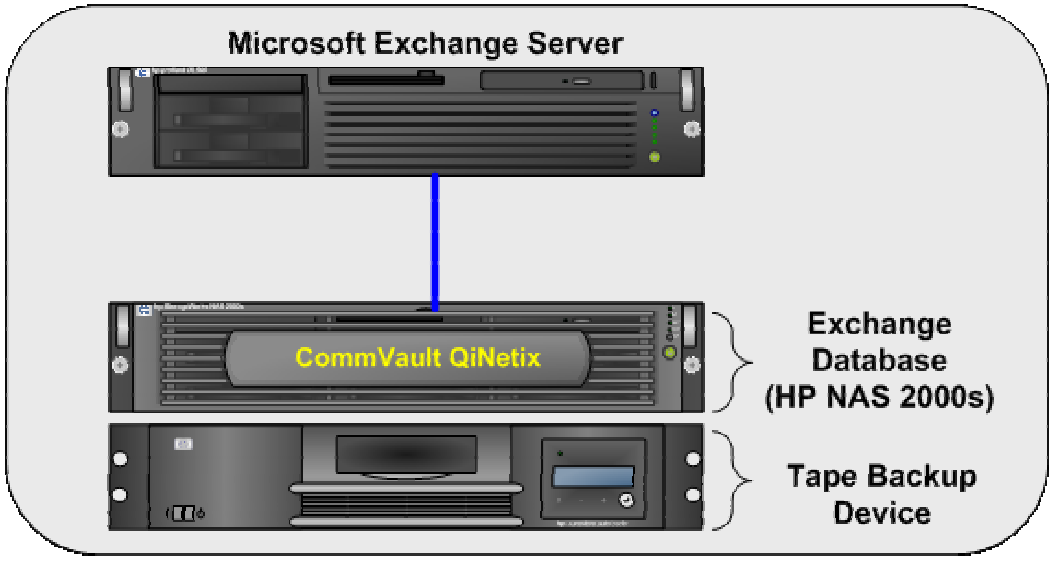

Incorporate Exchange business continuance into your NAS storage management strategy

Using HP NAS to consolidate Exchange data with file server data reduces "server spawl." With data previously hosted on numerous servers migrated to more powerful and inexpensive HP NAS devices, it is essential that this business-critical data be well protected. CommVault® Galaxy™ Backup and Recovery software (part of the CommVault QiNetix unified data management suite), coupled with HP NAS on Microsoft Windows Storage Server 2003, provides a compelling answer to the question of business continuance.

The goal of business continuity from an IT perspective is to restore access to critical data and e-mail as quickly as possible to restart business operations. In formulating a business continuity strategy, one must consider how long the business can afford to have this critical information unavailable—whether it is minutes, hours, or days. There should be a corresponding investment in business continuity products and services. Downtime results in lost revenue, declines in customer and shareholder confidence, damage to company image, and potential loss of market share. Arriving at an estimated cost per hour of downtime is one useful method for determining an appropriate investment in disaster recovery planning.

The key element is what is referred to as time-to-data. How fast can the data that is absolutely required to restart business operations be accessed? Should it be replicated, mirrored, restored, or recovered? Several factors determine time-to-data, including the granularity of the software used to perform the recovery. Can the recovery software retrieve an individual e-mail or file? Or is a time-consuming, full recovery all that is required to get a single piece of information? Was the program designed to minimize time-to-data? CommVault System's Galaxy Backup and Recovery software, coupled with HP NAS devices, provides answers for all these questions and meets today's needs.

How CommVault Galaxy Backup and Recovery software improves data access and business continuity for HP NAS deployments

CommVault's Galaxy Backup and Recovery software provides a complete management infrastructure for managing, safeguarding, and recovering your mission-critical data stored within HP NAS environments. Key advantages include:

- Flexibility, scalability, reliability, and ease of use. Galaxy software provides a series of default settings, automatically created at installation, so that all connected data can be managed and immediately protected using these default configurations.
- A logical view versus a physical view of all storage resources in a single unified browser-based management console, enabling fast, easy restores without having to know where the data resides.
- Policy-driven or user-directed data management at various levels within the solution. The level of flexibility is driven by site-specific needs and unique data requirements.
- Management of a virtual storage pool consisting of disk, tape, and optical. This includes the management of robotic libraries, tape devices, magnetic disk backup pools, and tape media lifecycle management.
- An automated backup and restore process, flexible enough to provide protection to single system NAS devices, and robust enough to scale up to meet widely distributed backup needs.
- Magnetic disk as backup media allows customers to add magnetic disk to their backup hierarchy providing an additional media choice for fast backup and restore.
- Auxiliary Copy functions allow duplicate copies of backup data to automatically be generated for remote disaster recovery use or off-site vaulting.
- Auto Discovery of tape libraries and drives. Galaxy software automatically discovers available tape library and tape drive resources, significantly reducing initial configuration time and effort. Additional resources added at a later date are also auto-discovered and can be added to existing policies easily.
- Support for a wide range of tape library units (TLUs), stacker/loader tape devices, and stand-alone tape drives giving a comprehensive choice in media performance and capacity to meet a site's multi-faceted backup storage needs.
- Intelligent Dynamic Drive Sharing (iDDS) enables tape devices to be shared between multiple applications ensuring best utilization of tape resources and maximum return on investment (ROI).

Ideal Exchange e-mail backup and recovery capabilities

A good e-mail backup and recovery operation should be able to:

- Handle single-instance message storage. If 1,000 people have received the same message, the backup and restore software recognizes the duplicates and backs up just one instance of the file—and uses pointers to the other 999. This function saves considerably on tape and disk space, and significantly shortens backup windows. This enables CommVault's Exchange solution to be ready to manage an entire enterprise's Exchange server.
- Provide backup/recovery of all configuration and state data that is stored in the Active Directory server.
- Restore at granular levels. Restores should work from the server and volume level, through user mailboxes, folders, subfolders, and individual messages and attachments. Restore operations should also work without having to take the databases down.
- Restore by message characteristics and properties. Restore operations should search through a variety of message attributes, including attachments, message properties, and date/time stamps. IT should be able to restore the message to selected targets, such as individual mailboxes or restricted databases for discovery searches. IT can now effectively restore all aspects of e-mail, protecting historical and hidden properties.
- Restore by content. Content-based searches should identify, sort, and search through messages by sender, recipient, date, and subject line. This ability is vital to restore individual messages, to group messages relating to business functions, to prove compliance, or to satisfy litigation requirements.
- Provide remote restore capabilities. Remote restore operations are a critical function for distributed messaging systems. In this case, the restore application receives a recovery request and alerts a network administrator. The administrator should immediately be able to restore the requested message to any other messaging server on the network, wherever it may be.

CommVault Galaxy Backup and Recovery software encompasses all of these capabilities—and more. For complete details on additional Galaxy backup and restore features, visit: http://www.commvault.com.

Conclusion

With HP StorageWorks NAS supporting the consolidation of file server data with Microsoft Exchange data stores, business continuity solutions become a necessity for such deployments. The goal of business continuity is to provide a rational plan to assure access to data on a priority basis in the event of loss—but it should not be relegated only to severe or significant data loss scenarios. A good storage management strategy takes into account all possible scenarios. A good storage management implementation needs support for advanced tools, such as application data views, data copy management, media virtualization, sophisticated job scheduling, insightful reporting, and hardware discovery. With policy-based storage management, complex data management practices can be easily and systematically implemented, and then, most importantly, those practices can be automated. With CommVault QiNetix software, users can easily build a collection of storage management policies that provide a robust business continuance strategy for data stored on HP StorageWorks NAS servers.

Solution: Disk-to-NAS-to-Tape for Exchange servers

Over the past few years, we have seen disk being used as a target for backup. Clearly disk is being used for more than just fast backup. With the introduction of SATA drives, the speed, reliability, and capacity of these low-cost disks is continuing to expand its storage niche. IT is now starting to look at disk as both a local and remote layer in its storage hierarchy. Disk is being used for snapshot copies of data to deliver even faster application recovery. Effectively utilizing disk is the key to controlling TCO. Maximizing multiple levels, locations, and types of disk and tape storage is the secret to positively impacting your corporate bottom line.

With Windows-based HP StorageWorks NAS products supporting other HP StorageWorks products that utilize SATA disks, a platform for inexpensive and effective disk-based data protection solutions now exists. The low TCO costs, ease of deployment, and scalability of HP NAS make it an ideal platform for disk-based data storage solutions. This platform, coupled with CommVault Galaxy software optimized for disk storage, provides an ideal solution in this arena.

Advantages of Disk as Secondary Storage

The advent of lower cost disk solutions has created disk to disk (D2D), a category of secondary storage. Using disk for secondary storage offers the following advantages versus tape:

- Faster restore of data.
- Disks are generally faster than tapes, especially when mount and seek times on tapes are factored in.
- Tape failure issues are eliminated.
- Disks are random access and are optimized for lookups—tapes are sequential and can be much slower for random file recalls.
- Multiple hosts can access disk simultaneously—disks have multiple read/write heads.
- Using disk eliminates human errors in tape handling.

Gartner Group has said the most costly storage administration task is backup/data protection. Tape costs have not reduced in cost on par with the reduction in disk prices. In addition, backing up to tape has become more, not less, complex as data volumes continue to exponentially grow. Backup windows have ceased to exist. Critical applications such as Exchange must run almost continuously. This is essential when you consider that as much as 60% of businesses critical information is stored in e-mail. Downtime is not an option, whether planned or unplanned. Numerous studies have placed downtime costs at anywhere from $80,000 to $2 million dollars an hour, depending on the industry. Add to this the scarcity of skilled IT personnel, the shrinking capital budgets, the economic uncertainties, and the need for each department in every company to reduce costs **and** add to the company revenue, it is clear the average IT manager is in quite a dilemma.

With SATA technology, the cost of disk is diminishing quickly, especially with regard to the cost of tape. Is this good? Yes. Is it as good as the hype? Yes and no. D2D does not eliminate tape, but it does offer an additional method for taming your data management challenges.

D2D is more than just faster throughput versus current tape solutions. It is still about the data. It is not about backup to disk or backup to tape; it is really about maximizing effective use of all storage types in all locations. Each storage solution is meant for a particular phase in the lifecycle of a piece of critical information.

Possible deployment

Figure 2 illustrates a possible deployment of this solution. Notice that in addition to Exchange servers, the solution supports Disk-to-NAS backup of other applications.

Figure 2.

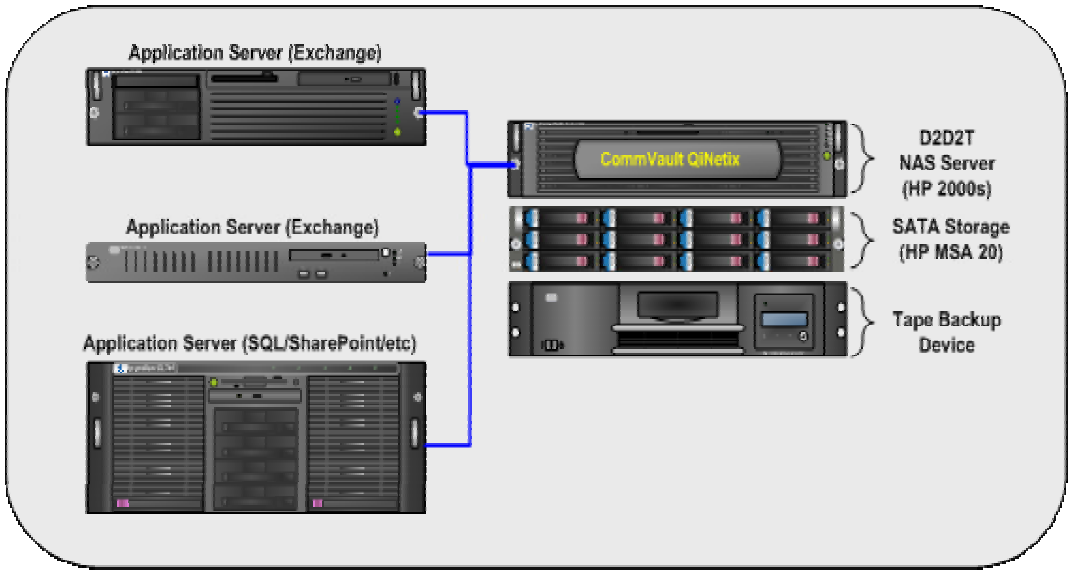

D2D options

Traditional backup

Traditionally data has been stored on disk and backed up to removable media—most commonly magnetic tape cartridges. This has become a tried and true method, perfected by independent software vendors over the last 15 years. Copies of data have been backed up to tape for both local and remote storage (Figure 3).

Figure 3. Traditional backup to tape, local and remote/vault copies on tape

As data growth has grown through the years, tape manufacturers have responded with greater cartridge capacity and faster throughput. However, in the last few years, data volume has grown

exponentially and tape has been unable to provide capacity and speed enough to enable customers to protect their data in their shrinking backup windows.

D2D backup and D2D appliances

So what should I do now? Backup as much as you can in your backup window? Don't backup? These are not viable options, especially considering the ever-increasing government regulations regarding maintaining copies of data for future compliance needs. To meet the need for faster and higher capacity backup, low cost ATA and SATA disk devices have surged to the forefront. These D2D disk products offer lower speed than primary disk, but preserve the fast access and provide faster throughput than tape. D2D prices offer much lower price points, now enabling customers to add the D2D layer to their backup storage hierarchy.

With opportunity comes hype. As you have now experienced, there are many types of disk-based backup devices or "appliances." These typically offer proprietary solutions that may or may not use standard operating systems or file systems, and most likely include embedded software to perform data protection functions. This is contrary to the approach taken with HP NAS, which highly leverages industry-standard platforms and software. It is also contrary to CommVault, which supports the standards-based NAS products from HP, in addition to HP StorageWorks SATA, tape, and optical products.

When looking at D2D devices or solutions, you must understand the complete implications of any solution. For instance, does the solution have the ability to make copies of data from the D2D device to local or remote tape? What important features provided in your backup and recovery software are, or are not, available in your D2D device? How about media management, retention, and tracking? Is the data indexed so that single files can be recalled quickly from the tape copies? Does the solution solve this by claiming you do not need tape at all? Better call your insurance carrier if this is the case. To quote Jon William Toigo, industry expert, "My message to you is simple: stop and think before you act."

D2D with CommVault QiNetix

Some other key points to scrutinize include finding out whether your data management software takes full advantage of the properties of disk. Does your data management software write to disk as a random access device or does it write to disk as if it were a tape in a single-threaded sequential data stream? Unlike tape, disk has many read/write heads, allowing multiple jobs and multiple read and write commands to be executed simultaneously. Data management software that takes full advantage of this can more effectively maximize your investment. CommVault has been writing to disk as a random access media for 15 years. CommVault is designed to take full advantage of every type of media and its specific characteristics (Figure 4).

Figure 4. With CommVault QiNetix copy data from any device to any device at any location automatically by way of policy management

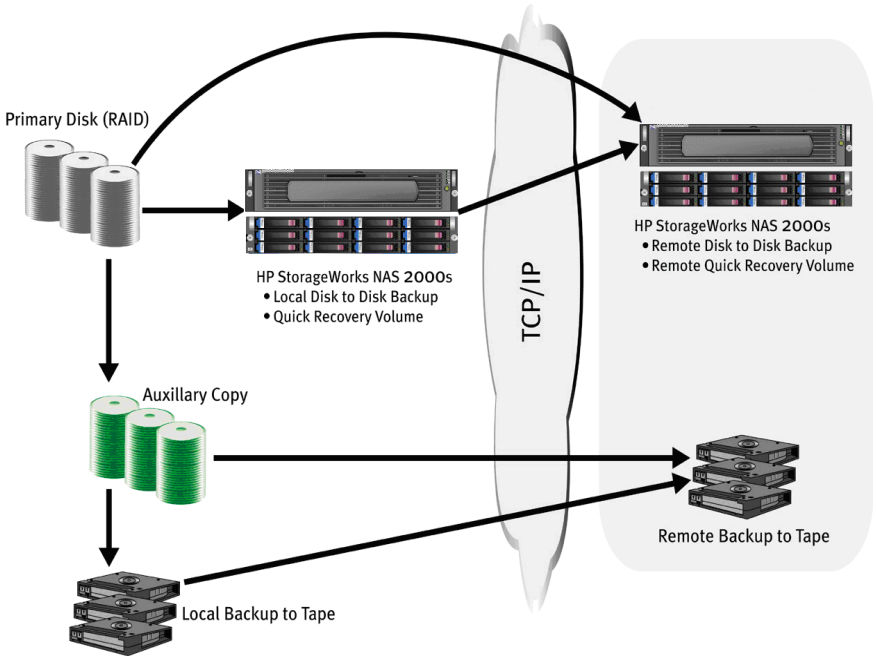

Another key CommVault feature is the ability to make copies of data from any media type to any other media type in any location. Additionally, CommVault policy-based management provides the ability to automatically make these copies from one media type to the other. Many solutions do not support copies of data from media type to an entirely different type, such as from disk to tape. With this support, CommVault customers can implement a D2D layer for fast backup and restore, then migrate copies of that data to tape after a time period (that is, 30 days). Automatic data migration frees up the D2D storage to hold more short-term copies of data. This hierarchical management of the storage and automatic migration of data from storage layer to storage layer truly empowers you to micromanage your data storage and maximize the investment in storage devices to help the bottom line of the business.

With optional product offerings including data migration, compliance archiving, and snapshot management, CommVault provides a complete set of data and storage management tools designed specifically to provide total data protection. The following sections describe the CommVault approach and briefly describe some of the core benefits of CommVault QiNetix software products. CommVault has designed products to provide superior ease of use, management, and technical superiority in assisting customers to manage their data protection requirements.

CommVault Galaxy D2D2T features

Some other key features with CommVault Galaxy software include:

- Backup data spans several mount points seamlessly.
- Backup volumes can be shared with other applications.
- Backup volumes can be moved between media agents without affecting operations and without any loss of backed up data.
- Dynamic mount path sharing enables a pool of SAN-attached disks to serve as targets for several Galaxy media agents. Access to the disks is coordinated by the CommServe resource manager.

Understanding the nature of random access devices and implementing a solution that takes advantage of the specific capabilities of a storage resource enables CommVault to help IT staff optimize their storage resources. Taking advantage of the native file systems and the two-part indexing scheme, CommVault can provide superior performance versus competing solutions that treat disk as a sequential "tape" device. This ability to recognize and optimize for disk as a backup target is especially important when data cannot be delivered to the backup media fast enough:

- Slow WAN links
- Many small files
- Proprietary Application APIs that feed data slowly
- Disk writers can support slow or intermittent writes without the stop/go/position problem of tape (shoe shining)

In tests with disk backup devices, CommVault has consistently demonstrated the ability to provide maximum throughput to disk used for backup and consistently higher backup and restore versus the most advanced tape devices available.

Media spanning on disk

Similar to media spanning on tape, Galaxy also allows backup to disk jobs to span from one file system to another without user intervention or job disruption. This is critical in reducing management complexity by not having to constantly monitor and manage available disk space before backup job execution.

Synthetic Full

The Synthetic Full feature allows for the creation of a new "full" backup image by combining a previous full backup with the associated incremental backups. The advantage of the Synthetic Full includes the ability to do this without touching the source data again. This removes the impact on the application server CPU and allows the customer to create a new set of full backup tapes for other uses such as vaulting or setting up a new site or test system.

Synthetic Fulls can be created from different media types. For instance, a weekly full could occur to tape, while daily incrementals are stored to disk for faster backup and more importantly, faster restore. A Synthetic Full copy can be created using the "full on tape" and the "incrementals on disk" to create the "new" full for offsite storage or disaster recovery. This flexibility empowers IT staff to tailor recovery plans to provide maximum restore performance and efficient use of all layers of the storage hierarchy.

Auxiliary Copy

With Auxiliary Copy, Galaxy provides the ability to create/migrate copies of data between different media types, different locations (Figure 5), or both. This "storage HSM" is completely policy based and occurs in the background without user intervention. Copies of data are created based on these policy settings and copies are retained and pruned based on retention policies. In this manner, copies of data can be moved throughout the storage continuum and expired or created on media as the access pattern requirements change over time.

Figure 5. Disk to Disk movement with Auxiliary copy

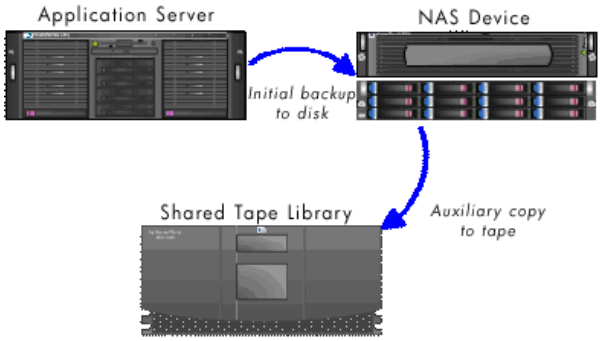

Auxiliary Copy is key in managing the cost of storing data. Auxiliary copies are used to make copies of data to less expensive media over time, allowing the more expensive media to be freed up for mission-critical data with faster access requirements. Using Auxiliary Copy effectively can significantly impact both the ROI and TCO of a company.

Storage Policies

Storage Policies (SP) are where key data management decisions are defined. The SP is a logical method of representing physical source and target locations. The SP consists of the following configurable items:

- Location of data copies
- Storage target types
- Optical disk
- Magnetic disk
- Magnetic tape
- Drive pool
- Library
- Capacity
- Data path
- Number of copies of data to create
- Retention periods per copy
- Recycle thresholds

When the SPs are defined, data is "assigned" to an SP for management following those guidelines. It is easy, point and click, to assign data to a different SP, changing the way the data is managed. There is no reconfiguration of hardware, any re-cabling or networking to accomplish these changes. It is all handled within Galaxy software and is based off the definitions of the SP. This feature reduces the cost and complexity faced in many IT shops in setting up and maintaining the infrastructure and management schema for safeguarding the data. With Galaxy SPs, much of this complexity is simplified, and the corresponding cost of managing the data storage environment is greatly reduced.

Auto discovery of storage devices

Galaxy software uses auto discovery techniques to provide a list of available storage devices and locations making these automatically available to the administrator when they begin to set up the SP. By presenting to the administrator all the possible targets for backup copies (disk, tape, and optical), the administrator has a clearer picture of the environment and can set up the corresponding data protection schema utilizing all available resources. This reduces the potential to overlook possible storage devices or more importantly to misconfigure the devices by making a mistake on a path entry or similar detail.

Application level integration

One of CommVault's fundamental goals with the delivery of Galaxy has been to provide data protection and management tailored to the specific needs of how clients use their data. This is directly opposite from other vendors who build their solutions around the physical storage devices and their idiosyncrasies. This paper's focus is Microsoft Exchange. However, with tight integrations for applications such as SQL Server, Active Directory, and SharePoint, Galaxy software can present an application view of data both from a backup and a restore perspective. The goal is to enable clients to back up the data in a manner that makes sense regarding how they might have to restore lost data. In the case of Exchange for instance, having to restore an entire database to get back a single lost message is both inconvenient and costly to the business.

VSS Integration—"60 second" recovery

CommVault also supports Microsoft's Volume Shadow Copy (VSS) framework. This means that backups of applications such as Exchange can use the VSS framework for near instantaneous backups. CommVault provides simple, transparent integration with Microsoft's VSS. A simple checkbox within the backup configuration screen enables VSS for all VSS-enabled applications. CommVault also makes the use of VSS hardware shadows transparent as the Galaxy and QR software will automatically default to a valid, registered VSS hardware provider, if one is available from the hardware server in question. This includes hardware "transportable shadows."

Granular restores

As mentioned previously, Galaxy provides the ability to backup the data and index it in such a way that clients can, for certain applications, restore individual objects instead of being forced to restore entire systems. For Exchange, this means the ability to restore that single piece of mail directly back into the mailbox of the user without causing an inconvenience or data loss to the rest of the users. The important item to remember for this is that the restore is coming directly from the backup media to the user. With built-in granularity for applications, Galaxy reduces the downtime and potential cost of not having the data available to the user when it is needed. The simplicity with which these restores occur also significantly reduces the personnel time and cost of having to do multi-stage restores and search and discovery of data.

Conclusion

The combination of HP StorageWorks NAS, built on Windows Storage Server 2003 and using CommVault QiNetix Software, provides the ability to combine technologies intelligently and to apply these effectively in a real-world scenario. By understanding the data, establishing data prioritization, and optimizing movement and storage of that information asset through the Storage Continuum, a business can improve the ability to:

- Safeguard Microsoft Exchange data
- Assure fastest access to Exchange data throughout its lifecycle
- Manage costs
- Effectively utilize personnel
- Maximize ROI, thereby improving company revenue

Knowing how and when to implement the current and future advanced technologies, like D2D, for managing critical data assets are prime factors in changing IT from a cost center into a tool for improving profitability. With the functionality contained in the CommVault QiNetix suite, businesses now can take control of managing their data in a more effective manner and positively impacting their company profits. Technology sprawl can now be contained.

Solution: NAS as an Exchange e-mail ILM target

This solution builds on the concepts and ideas discussed in the previous "Disk-to-NAS-to-Tape for Exchange Servers" solution. It leverages the same strengths of the HP Windows-based NAS platform and CommVault's QiNetix Software previously detailed. The resulting solution helps deliver on the HP ILM strategy, and provides capabilities ranging from basic e-mail lifecycle management to regulatory compliance solutions.

As noted previously, e-mail and databases house the vast majority of critical business information. But while databases are methodically protected and routinely queried, messaging systems go begging—their performance is ragged, most messages are never archived, and they lack the tools for comprehensive searches. This is true even though the loss of just a few important messages, or the discovery of incriminating ones, can impact a business' viability, operations, employee communications, customer relations, and e-commerce. So given e-mail's importance, how do you manage it? How do you archive and restore e-mail messages? And how do you manage their content? The remainder of this white paper will answer these questions.

Possible deployment

Figure 6 illustrates a possible deployment of the hardware and software components for this solution. Note that numerous tape and optical devices are supported. The devices can be attached directly to the NAS device or exist elsewhere in the environment.

Figure 6.

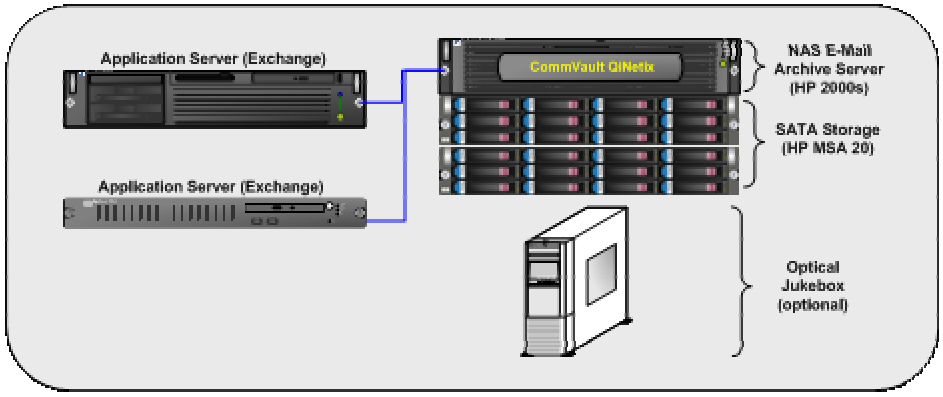

HP Information Lifecycle Management

It is worth taking a moment to set forth a simple description of the HP ILM process (Figure 7) and how this solution delivers on this strategy.

Figure 7.

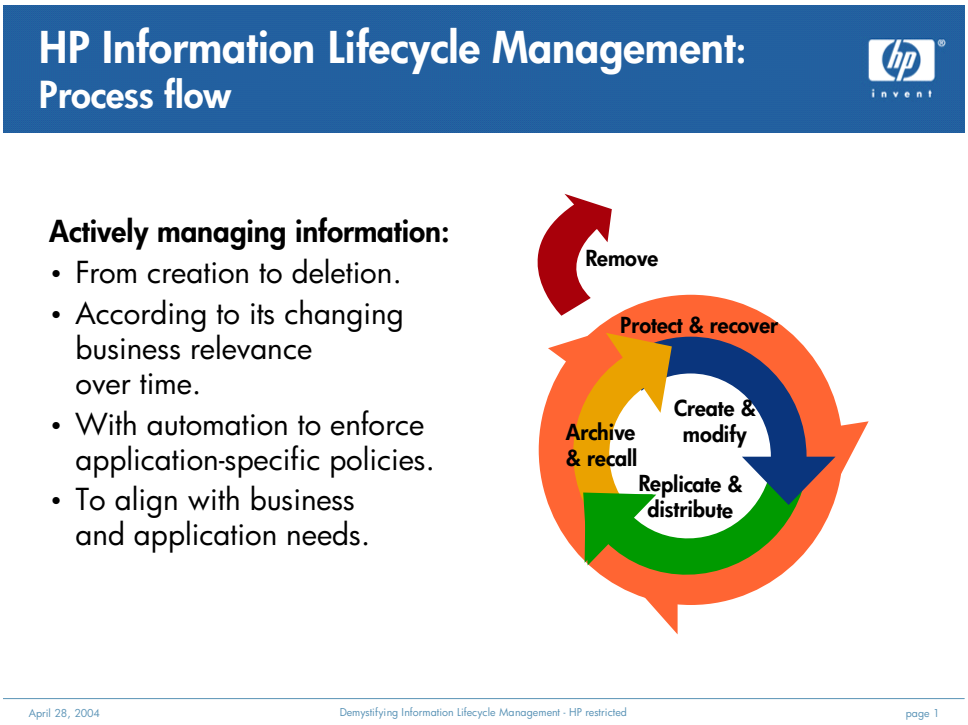

Delivery on the HP ILM strategy by this solution can be thought of as three parts:

- **Part one** is a set of storage devices of different capacity, throughput, connectivity, and most importantly price. An example is an Exchange server with SCSI disks, a NAS device with SATA disks with an attached SDLT autoloader, or an magneto-optical WORM drive.
- **Part two** is an engine that can move mission-critical data to and from each of these storage layers, from disk to disk, to tape or other mediums such as magneto-optical, without restriction on whether the source or target is local or remotely attached. The CommVault QiNetix suite is an example of such an engine.
- **Part three** cannot be purchased. It is having an understanding of one's data and being able to ascribe proper value to each piece of data at the various points in the data's lifecycle.

Information (e-mail for the purpose of this paper) must be recognized as having a life of its own—it is dynamic. It is created and as it ages, it has periods of more or less use during its existence. The management and protection of the "data" through this lifecycle is critical to maximizing the effective use of IT dollars, personnel, and time. ILM, as this is called, is the combination of breaking data down into categories of priority. These categories can be determined by common factor, some of which follow:

- How long can you survive without access to the data (minutes, hours, days)?
- How long must you keep or have access to the data (days, months, years, forever)?
- What is the required response time for access to the data (seconds, minutes, hours, days)?
- Are there regulatory issues regarding location and retention of copies of the data?

Understanding access patterns is a key to providing the proper data management. Only by understanding these patterns can IT make the proper technical decisions to meet the data retention, recovery, and access needs.

Today the IT administrator has many disparate technology tools, each of which has compelling benefits to improve any number of IT issues. The problem is the immense investment of personnel time and training required to become familiar with each technology. This prevents IT from reducing the cost of utilizing these advanced technologies and diminishes its ability to positively impact the company bottom line. What is needed is a single ILM tool providing IT the ability to use and manage these technologies, based on the data's characteristics and data protection needs. CommVault and HP have joined forces to solve this issue. CommVault's QiNetix suite of data and storage management products provide an engine for ILM in the Exchange e-mail space while the Windows-based HP NAS products provide a flexible storage environment to host the various tiers of required storage.

At the end of your deployment of any ILM solution, the following questions must be answered by a response of "yes:"

- Does the solution pay for itself in terms of business value and improved operation efficiency?
- Using polices to select and move data to more cost-efficient secondary storage?
- Ensuring transparency for users and applications?
- Making it easier to administer, preserve, and access that data?
- Scaling across ramping growth in users, servers, and data volumes?
- Providing a solution with inherent business continuity and disaster recovery?

CommVault's QiNetix software addresses the requirements by unifying the traditionally separate functions of data movement with data management; the entire storage stack, from application to device, can be managed as a cohesive whole to provide Exchange data access and availability in an automated manner.

Data migration or archiving: what do you really need?

There is renewed interest in data migration and archiving. This is especially true for Microsoft Exchange as a solution to control or manage growth of online data. But, migration and archiving are not the same, in spite of the interchangeable way vendors use the terms; so, short definitions are offered to help clarify the benefits of each.

E-mail migration: A data management strategy designed to improve IT operations (server performance, backup, recovery, availability, and management) by moving a portion of online data to cheaper secondary media. Users access migrated messages from their familiar Microsoft Outlook interface. By cutting administrator time spent juggling storage resources, reducing the time end-users squander organizing mailboxes and deleting e-mail, plus postponing hardware upgrades, data migration yields a real ROI.

E-mail archiving: A data management strategy intended to provide a permanent, searchable record of e-mail to meet regulatory, governmental, or organization needs. A separate, duplicate store of e-mail is created. It is accessed by the archiving tool's interface, not Microsoft Outlook. Many archiving products provide sophisticated search parameters, interactive monitoring, and fully auditable transactions. Thus, the driving rationale for archiving is compliance, not expense reduction.

An integrated approach is needed: Managing the complete lifecycle of e-mail covers four phases: creation, storage, and maintenance, plus expiration. Today's point level solutions, from backup to migration and archiving, offer individual pieces to the data management puzzle but lack the integration necessary to deliver the optimum management and storage efficiency. For example, a second e-mail set must be created, stored, and maintained with the current e-mail archiving solutions. The backup software makes another copy of every e-mail. Add user .pst files and administrators are managing up to four copies of every message.

To streamline the e-mail data lifecycle process and minimize costs, e-mail archiving capabilities should be integrated into existing Data Protection and Data Migration tools. That strategy would offer a multi-tiered, cost-effective data management strategy with the added benefit of reducing the number of copies of each message by at least two thirds. It also ensures that valuable information can be retrieved quickly and efficiently.

CommVault® QiNetix™ DataMigrator™

The QiNetix DataMigrator product migrates older Exchange data from primary storage (for example, SCSI storage on an Exchange Server) to secondary, less expensive storage (SATA storage on an HP NAS device). Using DataMigrator software shrinks backup windows, reduces primary disk storage, eases IT administrative burden, and improves primary application or file system performance—all reducing customers' TCO. Because DataMigrator is integrated with Microsoft Exchange, end-users do not have to access a separate interface to recall migrated data. They simply use the familiar Microsoft Outlook e-mail interface.

The QiNetix DataMigrator product supports Exchange 5.5/2000/2003 messages or attachments migration to less expensive secondary storage while letting users transparently view migrated data through Microsoft Outlook interfaces. With QiNetix DataMigrator, IT organizations have an automated way to control users' storage on network shares and other corporate storage resources. In addition, end-users now have a seamless way to extend the capacity of their mailboxes without resorting to local, unmanaged .pst files. By centralizing storage of all file and Exchange messaging information, administrators can be assured all data is properly protected and catalogued, or purged when the retention period is met.

As part of the QiNetix solution, DataMigrator allows the clients to set policies to automatically manage the size of these key database applications while still allowing for access and tracking of information that has been migrated to less expensive storage (Figure 8).

Figure 8. QiNetix DataMigrator for Exchange migration policy settings screen

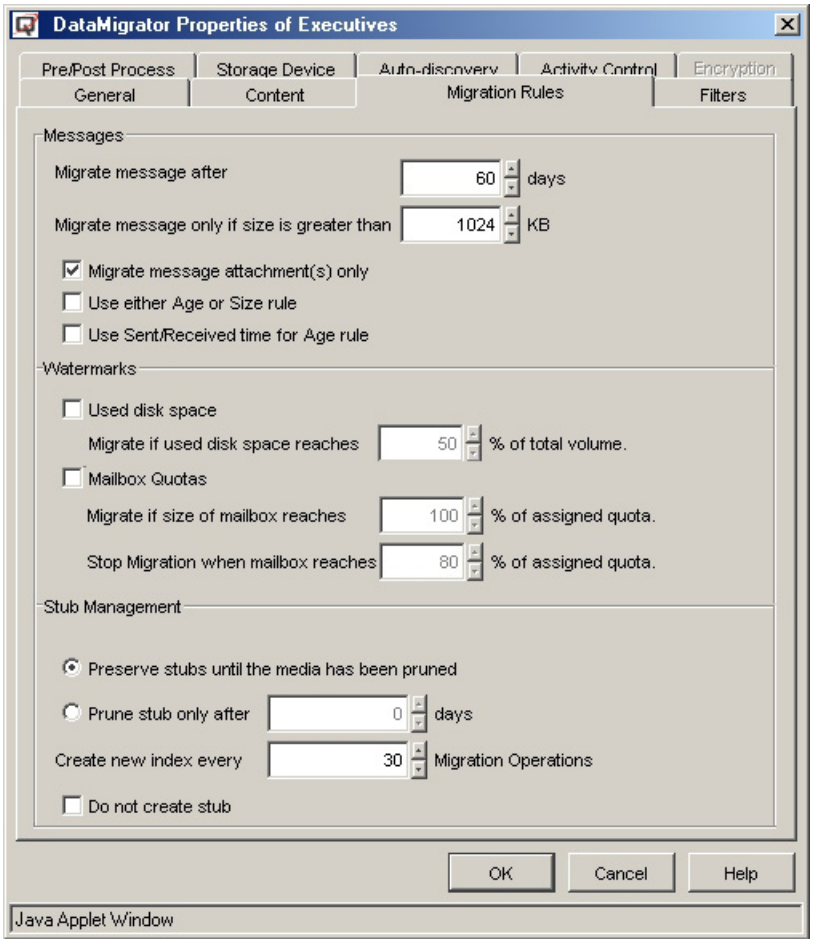

Through effective use of the DataMigrator tool, clients can keep the cost of primary disk under control and more effectively utilize all aspects of the storage management infrastructure.

As with the other QiNetix solutions, the DataMigrator features are all accessible from the same QiNetix interface, providing the single, unified console view for management of all QiNetix products. DataMigrator also shares the same D2D and disk optimization features discussed previously.

CommVault® QiNetix™ DataArchiver™

The QiNetix DataArchiver product provides a method for capturing and archiving e-mail data for compliance with various regulatory, corporate, or legal retention requirements. Additional capabilities such as content indexing, keyword search, and Auxiliary Copy provide flexibility and electronic tracking and search for easy satisfaction of legal search requests. The software's secure and audited retrieval meets the need for legal or regulatory compliance. When combined with HP NAS utilizing SATA drives, HP StorageWorks tape drives/libraries and HP MO/UDO WORM drives provide a complete compliant solution that is within reach of the small and medium business customer.

Does your company archive e-mail or will it soon? A recent survey by *Storage Magazine* cited 70% of customers polled as replying yes. That number is sure to rise as businesses are facing new regulatory enforcement of legal, federal (state and local), and industry rules regarding retention of unalterable copies of data. Each company must soon develop compliance management practices for managing their risk regarding data storage and more importantly data access if records are subpoenaed. What are the main concerns when researching compliance archive solutions?

Customer needs include:

- Meet government/industry requirement to maintain and e-mail archive
- Complete collection of e-mail of a select group of users
- Be based on corporate or industry requirements
- Ensure e-mail accessibility for retention period
- Easily retrieve and package e-mail for audit purposes
- Address technology obsolescence and redundancy for long-term preservation
- Enable deep search and retrieval with audit controls
- Need full text, sampling and retrieval for risk management
- Have indexes that scale and remain available for entire retention lifecycle

Meeting these needs can be daunting. E-mail archiving very quickly grows into a data management challenge. For instance, take the following scenario, easily faced in most medium-sized business. "I need to archive data from 500 mailboxes and ensure that we can capture and search all of the messages."

> 60 messages a day at 5kb, 10% attachments at 500kb, full text indexing at 2:1 compression, assume 20% traffic/size growth rate. 3-year retention

That means in year one you have 7.5 million messages archived, in year two it increases to 16.5 messages, and by year three you have 27 million messages archived—to the tune of 1 TB of raw storage—just for your compliance archive! This is not a "do" or "don't do" event. The challenge now is ensuring that the initial rollout of your compliance product can scale for the long haul. Nothing could be worse than deploying once and then discovering that your compliance solution cannot grow as your company and regulatory compliance needs grow.

Figure 9. QiNetix DataArchiver—Compliance Archive process flow

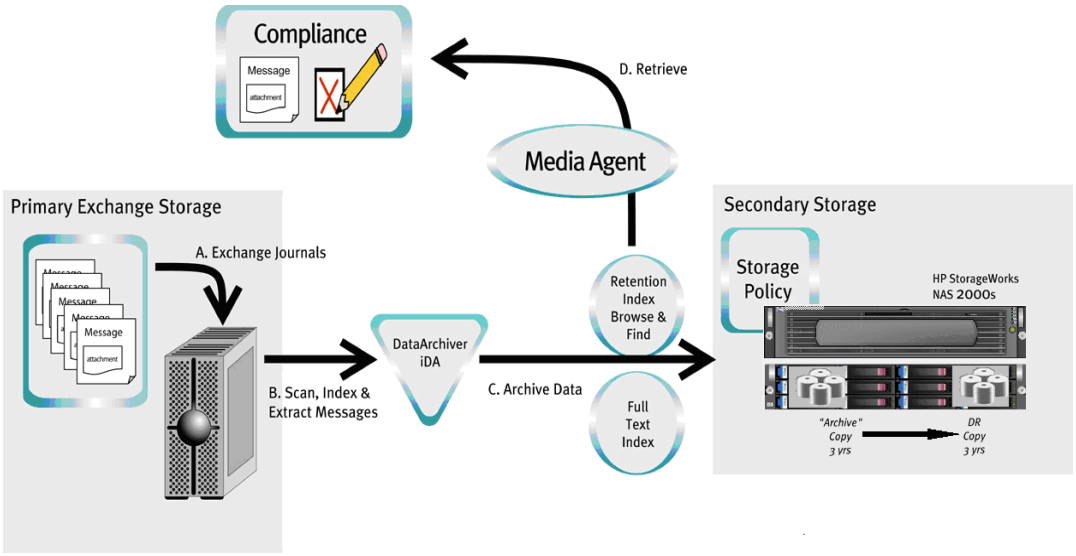

When looking for a compliance solution, keep the following requirements in mind and ask your vendor to show it can do all of the following.

Create and maintain an e-mail archive for compliance purposes with the following features:

- Capture journal, index e-mail contents, and provide user-automated retrieval
- Ensure all required controls (audit, tamper proof, and so on) and safeguards are in place
- Enable proactive surveillance and risk management to ensure compliance is working
- Lower TCO in archive management
- Is designed in data protection to make sure the compliance archive is backed up
- Provide common management tools for both backup and compliance
- Reduce the retrieval and discovery costs by automating all search and retrieval effort—remove the need for manual search of the archives
- Provide an integrated platform to perform compliance and data protection
- Ensure one system is both highly scalable to grow as the archive grows, and also compatible between different disciplines of ILM
- Provide built-in redundancy and protection to ensure the survivability of the archive in the event of disaster or human error leading to data loss

Compliance archiving is not as simple as making a copy of your e-mail or files and putting them in a back room. It requires significant planning to understand the obligation and the risk to the company in not being able to deliver on a request for specific data.

QiNetix DataArchiver, like the rest of the QiNetix suite, is a stand-alone solution designed to allow companies to comply with various regulatory agency requirements for information storage and access. Utilizing the Common Technology Engine provides DataArchiver with significant value versus the competition. By inherently sharing storage, having a single user interface, unified browser, and a consistent operational approach with the Galaxy and DataMigrator products, users can reap the benefits of lower ongoing operational costs and easier, more effective data management. This is true with DataArchiver deployed in a stand-alone mode and exponentially so when deployed in concert with the rest of the QiNetix suite.

Conclusion

Knowing how and when to implement advanced technologies for managing critical data assets are prime factors in improving a company's bottom line. With the CommVault QiNetix ILM functionality and Windows-based HP NAS, businesses can take control of managing their data in a more effective manner and to positively import their company profits while at the same proactively implementing strategic compliance and regulatory initiatives.

For more information

Visit the following links for additional information and resources for the solutions mentioned in this paper:

- Managing Exchange Storage with the Windows Storage Server 2003 Feature Pack: www.microsoft.com/windowsserversystem/wss2003/techinfo/plandeploy/exchange_fp.mspx
- www.hp.com/go/nas
- www.commvault.com
- www.microsoft.com

5982-6622EN, 06/2004