

WHITE PAPER

Gaining Control of Remote Office Data Protection

Sponsored by: EMC

Laura DuBois

March 2007

EXECUTIVE SUMMARY

The exponential data growth in remote offices, the increasing importance of data to run distributed operations, and a finite pool of technical IT staff are resulting in changes to remote office data protection strategies. Vital structured and unstructured data assets distributed across different remote office geographies, server platforms, and storage technologies must be managed and protected. Historically, remote offices have been treated as standalone islands, without centralized management, policy, or visibility. Limited onsite IT resources, inadequate WAN bandwidth, and administrative-intensive and failure-prone tape hardware resulted in inadequate and infrequent backup operations, thus compromising operational recovery. For firms that have relied upon removable tape media for disaster recovery, the use of unencrypted tapes collected by third parties has placed sensitive corporate data, such as customer information or intellectual property, at risk.

These remote office challenges combined with the cost and productivity impacts of downtime are causing IT executives to take control of remote office data protection. Firms are moving to a centrally managed and controlled WAN-based remote office backup and recovery method to ensure consistent, reliable backup and guarantee fast recovery. Originally bandwidth-prohibitive backup conducted via the WAN has seen a resurgence as a result of advanced data deduplication and single-instance storage approaches. These approaches, which also eliminate the copying or storage of redundant data, are enabling firms to finally protect remote office data across existing WAN links — ensuring the preservation and recovery of remote office data.

TODAY'S REMOTE OFFICE

Offices and operations physically distributed from a firm's headquarters are commonly considered remote or branch offices. Distributed operations, such as regional banks, development sites, manufacturing facilities, sales offices, healthcare clinics, educational campuses, telecommunication branches, and home offices, all can be considered remote offices. In today's technology-pervasive business environments, remote offices are continuously connected back to a firm's main or geographic headquarters via a corporate LAN or WAN. Remote offices commonly have local IT infrastructure, such as file, print, Web, and email servers; workstations and desktops may also house distributed applications/databases. Remote offices rely upon these systems and information to support regional business functions, such as order processing, inventory management, client communications, sales activity, clinical research, and the like. As a result, data is paramount for the continued operation of these remote offices, and data must be protected and available in the event of a local failure or disaster.

While IT infrastructure exists at today's remote offices, absent are dedicated technical or IT resources to support remote office hardware, software, and networking equipment. Compounding the challenge of limited IT resources, the sheer number of remote or branch offices is very large and growing. In the United States alone, IDC estimated in 2005 that there were a total of 1,355,486 branch offices associated with mid-size and large firms (those with greater than 100 employees). Including small companies, the number of remote offices in the United States jumped to 6 million in 2005. Mid-size firms (those with between 100–999 employees) had an average of eight remote or branch offices and large and very large firms (with greater than 1,000 employees) had an average of 65 remote offices. Evidence suggests that the total number of branch offices is increasing as industry dynamics such as mergers, acquisitions, off shoring, expansion into new markets, globalization, and telecommunicating of remote office workers continue.

Combined with this growth in the number of remote offices is the growth of corporate data, both structured and unstructured. In a recent IDC survey, firms realized an average of 52% growth in disk storage capacity over the past 12 months. This was fueled by growth in email, corporate applications, documents and images, new applications coming online, disk-based data protection, and the ongoing long-term retention of corporate information due to legal and regulatory requirements. The growth of data is happening in the main datacenters as well as in remote and branch offices. Data at the edge, in remote and branch office locations, must be controlled, managed, and protected. In the event of an operating system, application, server, or site failure, data must be quickly recovered.

REMOTE OFFICE DATA PROTECTION CHALLENGES

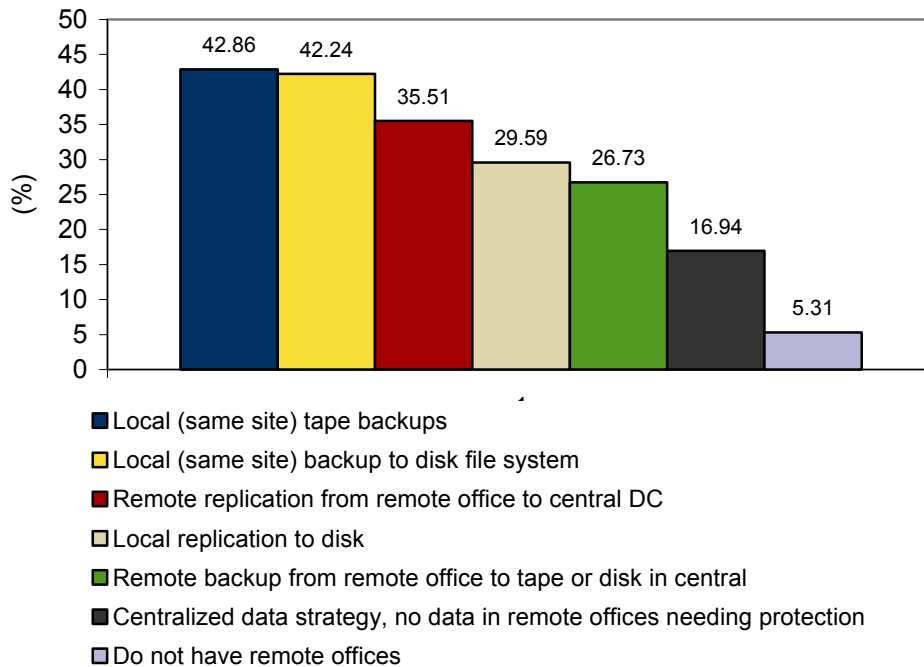
Traditional approaches to remote office data protection have been in the form of periodic and highly manual tape backups, conducted by office managers, to local tape cartridges, which were transported to offsite locations. The challenges with this traditional approach are many. This process is subject to human error when non-technical office staff manage backups, rotate tapes, or initiate recovery. The risk of data compromise due to the removable nature of tape cartridges presents risk of loss or theft of sensitive data. The high rates of both tape media and drive failure place backup in jeopardy and scarce technical resources must travel to the remote office to triage and resolve a problem. In locations where server, application, and file system configurations have not been protected, recovery requires a lengthy and complete rebuild of preexisting configurations from scratch.

However, increased focus on the value of information to the firm, legal and regulatory pressures, risk of corporate scandal, as well as limited IT resources relative to continued data growth are causing firms to consider new approaches for remote office data protection. The advent of lower-cost ATA disk storage in combination with sophisticated and centrally managed data protection and recovery software has meant firms can improve availability and recovery in remote office locations. IDC finds that firms are in the process of strategically evaluating their remote office IT strategies and looking to take control of the management and protection of remote office information technology assets. Increasingly, firms are deploying disk-based data protection and replication to address remote office protection challenges. Figure 1 outlines the methods used to protect remote and branch office data today.

FIGURE 1

Current Remote Office Backup Approaches

Q. How are you currently doing backup for remote offices? (Select all that apply.)



Note: n = 490

Source: IDC's *Disk-Based Data Protection Study*, June 2006

IDC expects firms will increasingly evaluate and deploy new disk-based data protection, backup, and replication technologies in remote offices to ensure that data can be quickly recovered and managed from a centralized point of control. In evaluating new approaches, firms must look for solutions that can address several key challenges with the protection of remote and branch office data.

Scarce IT Resources

Probably the single biggest challenge for distributed operations or remote offices is the lack of technical resources onsite to manage an administrative-intensive process such as backup. Without onsite technical resources, local office, or factory workers are left to manage backups and manually pull tapes. However, with tape backup being prone to failures, the reliability of the backup is put at risk as office workers lack the time or ability to troubleshoot problems, track tapes, or correct media or hardware failures. This leaves local data in a potentially unprotected state. When failures do occur, local office personnel lack the technical expertise to conduct a restore and ensure that data can be recovered. The lack of onsite technical resources to manage data protection processes places the firm at risk of not being able to protect data, control backup processes, or meet recovery time objectives.

Reduced Reliability

Based on recent IDC research, the most common types of failures associated with data protection are due to bad tape media, partial job completions due to the closing of the backup window, configuration errors, network timeouts, and hardware errors such as tape-drive loading, ejecting, or read/write errors. This reliability challenge is not unique to remote offices and can be found in datacenters as well. However, the reliability problem is compounded when resources are scarce. There are typically few, if any, technical resources at remote offices to monitor backups and restart them in the event of a failure. As a result, backup errors are not detected until a recovery is attempted and fails as a result of a failed backup.

Data Growth, Redundancy, and the Shrinking Backup Window

The combination of data growth year over year and shrinking or non-existent backup windows represents a significant challenge IT managers face. This problem exists not only in the main datacenter but also in distributed remote offices. When traditional local or remote backup processes are performed, data is copied to backup media. The backup includes data that has not changed since the last backup and duplicate data that exists across desktops, servers, and remote office locations. Traditional backup methods commonly require daily incremental and weekly full backups. This schedule and the backup of redundant data can result in the movement of up to 200% of primary data each week and exacerbate an already shrinking backup window. This problem is compounded over a period of months or years, as traditional backup methods inefficiently store data and unnecessarily consume valuable network bandwidth and storage capacity.

These traditional methods are not feasible for remote office backup because of the volume of data that must be transferred, the limited bandwidth connecting remote offices to datacenters, and the ever-shrinking backup windows. With more remote operations operating over a 24 x 7 period, firms are struggling to meet backup windows without impacting end-user or application productivity.

Risk of Data Compromise

As previously discussed, a common approach for remote office data protection is local backup to tape media using standalone drives. For disaster recovery purposes, removable tape media is periodically collected by third-party personnel for transportation offsite to a vault facility. The removable nature of tape media places the data at increased levels of risk for compromise, damage, or loss. Companies can face public scrutiny, fines, and damage to corporate brand if sensitive corporate or customer information is compromised. This compromise may be due to lack of physical and logical security. For firms that decide to encrypt backup data on tape cartridges, the challenge of managing encryption keys can be significant for remote offices with limited IT staff. Moreover, the loss of removable media, even if encrypted, still means the data is lost and not available for recovery.

Recent government regulations require that firms publicly disclose data loss incidents. State legislation, such as the California Security Breach Information Act (SB 1386), requires a business or state agency to disclose any security breach to customers if its data has been acquired by an unauthorized person. This regulation, while not new, now requires an agency, person, or business that conducts business in California and owns or licenses computerized "personal information" to disclose any breach of security (to any resident whose unencrypted data is believed to have been disclosed). IT executives must understand the laws governing the handling of sensitive customer data as amendments are made to state and federal legislation. Also addressing the risk of customer data loss are regulations such as the California Database Protection Act (CDPA), the California Assembly Bill 1950, U.S. Fair and Accurate Credit Transactions (FACT) Act, and the U.S. Gramm-Leach-Bliley Act (GLBA).

No Control or Standardization

Without any processes, policies, or controls in place, standardization is difficult. It is common for firms with a large number of remote offices to have a hodgepodge of different data protection approaches in place. Some remote offices have local backup to tape, others have no data protection strategy, and yet others may have some type of periodic replication or backup to disk. Also, each approach may utilize a different backup software application controlling the process. The significant variety of different hardware and software solutions in place contributes to administration overhead, increasing IT costs and risk of remote office data loss.

Lengthy Recovery Times

Remote offices, like corporate headquarters, often cannot tolerate downtime or loss of worker productivity while an operational restore is conducted. Remote offices must continue to process orders, service customers, conduct research, and communicate with suppliers. With remote office protection most commonly done with local tape backups, use of this media can significantly compromise recovery-time objectives. The problem is compounded further if tapes required for recovery are offsite or if technical resources need to travel to the remote office to initiate a restore.

Virginia DMV Centralizes Disaster Recovery Services and Eliminates Tape Backup at Customer Service Centers

IDC recently spoke with Mike DePhillip, backup administrator for the Virginia Department of Motor Vehicles, about their move to a centralized, disk-based backup approach for their remote office and datacenter environment. Currently, the Virginia DMV has 73 remote customer service centers (CSCs) where Virginia residents can replace or renew driver's licenses, create and purchase license plates, and register vehicles. Each CSC is equipped with a server that processes customer orders and sends these transactions to a main datacenter at the DMV headquarters. Each remote CSC communicates with headquarters over WAN links ranging from 56K to T1 connections. The DMV was looking for an alternative to distributed, manual, and redundant scripted tape backups for these servers running in their 73 remote offices.

In each CSC, scripts controlled backups to local tape drives directly attached to the server. This approach required DMV office workers to regularly place a tape cartridge into a tape drive and manage and rotate media. The DMV wanted to move to a centralized, enterprisewide disaster recovery approach to protect both remote office and datacenter server data. The DMV was looking for a fast and centralized backup approach that could work with their existing WAN infrastructure, provide fast recovery in the event of server failures, and centrally manage remote backup and restore processes.

While evaluating needs for the company's enterprisewide disaster recovery processes, the Virginia DMV considered several options to improve upon their older remote office backup architecture. After evaluating several products, the DMV, EMC, and CA, the DMV decided on the EMC Avamar products. "I wanted a solution that could scale and accommodate future growth as we add an additional 1,200 users in the field," DePhillip said. The Avamar data deduplication process enabled a dramatic 80% reduction in backup times and data in all 73 CSCs was able to be backed up in four to five hours. The Avamar solution will also enable significant labor savings. In its current deployment protecting 180 servers, Avamar has reduced the administration times from 35–40% to approximately 10%. As the DMV replaces its DR solution in two other server environments, the centralized, enterprisewide DR approach will enable reduction in disaster recovery administration going from an average of 35% utilization of 3-4 people to 25% utilization of a single DR administrator.

The Avamar approach not only saved in operational costs and reduced the backup window but also saved the DMV money by reducing the cost of downtime. The company had a server failure and was able to restore the entire server environment including the data in less than 45 minutes. The old tape restore method would have taken 5 to 6 hours. With an estimated \$10,000–15,000 per hour in cost of downtime, the fast recovery from the Avamar approach saved the DMV approximately \$80,000 in that instance alone. According to DePhillip, "The product has also helped reduce restore times for files as well. When a folder was inadvertently deleted by a user a restore took seconds."

"Avamar has delivered everything they promised and more," said DePhillip. "Today we use the Avamar product to not only backup our remote CSCs but also protect 180 servers in our main datacenter. In the future, the DMV plans to expand the use of the A product to replicate data from their main datacenter to a DR site for enterprisewide disaster recovery services."

Many firms would like to completely eliminate the use of tape in remote offices since it increases administration overhead, creates reliability and security risks, and is expensive to maintain. Firms are looking to other approaches, such as backup to disk and remote replication to address the branch office data protection and recovery challenge. They want a solution that can centrally manage protection and recovery of data in many to hundreds of remote offices, while reducing the capital costs for backup of remote office data. Solutions that provide centralized control of remote office backup using global data-reduction technologies can address these needs.

EMC AVAMAR'S APPROACH TO THE REMOTE OFFICE BACKUP DILEMMA

EMC Avamar addresses the challenges associated with traditional backup methods, reduces costs, and enables fast, reliable backup and recovery across the entire enterprise — including datacenters, remote offices, desktops, and laptops. Unlike traditional backup solutions, EMC Avamar uses patented data deduplication and single-instance storage technology to reduce the size of backups at the source — before it is transferred across the network. According to EMC Avamar, this reduces the required network bandwidth and backup storage by up to 300x globally, enabling remote office backup operations using existing WAN links. By leveraging existing LAN/WAN bandwidth, EMC Avamar allows enterprise organizations to centrally manage and automate remote office backup operations while significantly reducing the required backup window.

EMC Avamar's main offering is Avamar software, which is bolstered with application-specific plug-in software modules. EMC Avamar's solutions eliminate the necessity for IT organizations to rely on untrained local staff and complex manual tape rotation schemes (e.g., swapping of tapes, transport of tapes offsite, etc.) since Avamar can be automated to perform regular backups without any manual intervention. EMC Avamar solutions eliminate the need for failure-prone tape-based devices at remote offices, which can be expensive to maintain. In addition to automatic scheduled backups, Avamar verifies the recoverability of all backup data daily to ensure it is available when needed. And the security issues associated with shipping unencrypted physical tapes offsite is eliminated since Avamar encrypts data while in flight across the WAN and at rest. Data backed up using Avamar software can be quickly leveraged via powerful search engines (e.g., Google and FAST) in support of legal discovery and regulatory audits. And if tape archives are required, Avamar integrates seamlessly with existing tape environments.

EMC Avamar's data protection software is well positioned to address today's remote office challenges and can be deployed to protect the following remote office environments:

- ☒ **Small, remote offices.** For remote office laptop and desktop workers and small remote offices, Avamar's intelligent software agents can be deployed on each of the systems that require protection, and the data can be backed up directly over relatively low-speed network connections (e.g., dial-up, cable modems, DSL, etc.) to a central Avamar system deployed at the datacenter.

- ☒ **Larger, remote sites with infrastructure.** Larger remote offices or operations with server infrastructure can deploy a local Avamar server configuration. Data can first be backed up to a local Avamar system, which can then be replicated over the WAN to a larger Avamar system deployed in the central datacenter.
- ☒ **Disaster recovery site or remote datacenter.** Firms with dual datacenter strategies can protect data locally with a datacenter Avamar server configuration and then remotely replicate to a remote, yet equally active, datacenter for disaster recovery purposes. Each datacenter is responsible for the disaster recovery protection of the other datacenter.

EMC Avamar Architecture

Avamar utilizes what IDC calls a grid-based architecture to provide scalability, performance, and flexibility difficult to achieve with traditional data protection solutions.

Enabling scalability and performance, Avamar's grid-based architecture allows administrators to simply add nodes (Intel-based servers) to easily increase CPU, memory, I/O, and disk capacity for the entire grid, with no need for scheduled downtime. Data is automatically load balanced across the newly added nodes for improved performance, and nodes can consist of dissimilar hardware for added flexibility. This allows administrators to purchase only the capacity they need (what EMC calls just-in-time purchasing). And every Avamar client can connect to every storage node, which eliminates performance bottlenecks.

To ensure high availability and fault tolerance, Avamar employs what the company calls a redundant array of independent nodes (RAIN), a patented technology (US Patent No. 6,826,711) to provide failover and fault tolerance across nodes in an Avamar Server grid. Avamar can continue to provide reliable data protection and access, even if a server node fails, since data stored on any node can be reconstructed from the other nodes.

Avamar also verifies recoverability. Avamar monitors the integrity of all data stored in an Avamar Server on a daily basis to verify that all backups can be successfully restored when needed. In addition to protecting enterprise systems, Avamar also protects itself with daily internal checkpoints, consistent snapshots of the entire Avamar system that can be verified for integrity and used for system "rollbacks."

EMC's offering consists of Avamar software, along with application-specific plug-in software modules.

- ☒ **Avamar software.** Avamar software is the heart of the environment. Smaller offices can install Avamar software agents directly on the systems to be protected in order to facilitate backup across the WAN to a central datacenter. Larger remote offices and datacenters utilize Avamar software agents along with a local Avamar Server, which consists of Avamar software installed on standard Intel-based servers from Dell, HP, IBM, and Sun. A key differentiator, Avamar software agents identify and filter out redundant data segments at the source before any data is backed up across the network. Using Avamar's patented data

deduplication and global single-instance storage technologies, the required network bandwidth and backup storage is reduced by up to 300x. This ensures that only unique, new data segments are backed up across the network and stored just once globally — enabling remote offices to back up data over existing WAN bandwidth. Avamar is also the central engine for scheduling, initiating, storing, and managing retention for remote client backups. Avamar backup data can be immediately restored from an Avamar Server and leveraged by third-party applications. In addition, Avamar utilizes a grid architecture that enables the solution to easily scale from one to many Avamar Server nodes to support any size remote office.

- ☒ **Replicator software.** Replicator software enables efficient, encrypted, and asynchronous replication of heterogeneous data stored in an Avamar Server to another Avamar Server deployed in an offsite location for disaster recovery and business continuity. Avamar eliminates redundant data at the source. This allows organizations to utilize existing network infrastructure and bandwidth and avoid issues that frequently hinder multisite replication, including the cost of procuring additional network capacity and the time it takes to transmit huge quantities of data. All data is encrypted for security, eliminating the risk of data loss and compromise due to shipping tapes offsite.

Approaches to Data Deduplication


Some products with data deduplication functionality do not deduplicate data at the source. Instead, they receive traditional backup data streams, which contain a very large amount of redundant data. These backup data streams must pass through the backup server and across the network until they finally reach the data deduplication device. According to EMC, this approach can move over 200% of the primary data every week, which is very inefficient since it does nothing to reduce the amount of backup data being sent through the backup server and across the already congested LAN/WAN. Another differentiator for Avamar is that Avamar client software supports the automated protection of a broad array of operating systems and applications found in remote offices. Some alternative solutions support a limited number of operating systems and may protect only file data.

The EMC Avamar architecture provides important features that users require in a disk-based data protection solution. Capabilities such as ensured data and process integrity, easy access to data for electronic discovery, remote replication for disaster recovery purposes, data reduction support, and security capabilities are paramount to firms in selecting a backup-to-disk solution. Figure 2 outlines results from a recent survey.

FIGURE 2

Features Most Important in a Backup-to-Disk Product

Q. Please select the two features/benefits that are most important to you. (Select only two)



Note: n = 484, Currently using or planning to use a backup-to-disk product

Source: IDC's *Disk-Based Data Protection Study*, June 2006

Avamar's remote office data protection solutions provide the following key features and benefits:

- ☒ **Global data reduction and single-instance storage.** Avamar eliminates redundant backup data at the source — before data is sent across the network. This means network bandwidth savings, shortening of the backup window, and storage optimization. Avamar client software only moves new, unique data segments for backup and recovery. Only a single instance of each unique data segment is stored in a central location. The Avamar client software views primary data in terms of sub-file variable length data segments. Avamar client software assigns each data segment a unique ID, generated based upon its content, which is used to compare with data segments that have already been backed up. Only new, unique data segments are transferred during a backup operation. This approach deduplicates data at the source, and according to EMC, reduces the required network bandwidth and backup storage by up to 300x globally. The reduction of data at the source is a differentiator for Avamar and is ideal for remote office data protection over limited WAN networks.

- ☒ **Network bandwidth-optimized backup.** Avamar makes use of existing network infrastructure. It addresses the challenges associated with effectively protecting remote office data by providing automated, centrally managed, bandwidth-optimized backup across existing WANs. By reducing redundant backup data at the source (before it is transferred), the amount of required network bandwidth can be decreased significantly over traditional backup approaches. This enables companies to perform daily full backups of remote office data using the existing WAN bandwidth.
- ☒ **Grid storage architecture.** As previously mentioned, Avamar utilizes a grid-based architecture, which provides significant benefits over traditional data protection approaches. These benefits include reliability, availability, scalability, performance, and flexibility in deployment. Remote offices can make use of Avamar's distributed grid architecture by deploying many remote Avamar clients that communicate with a centralized Avamar server configuration in a main datacenter.
- ☒ **Centralized, policy-based management.** Avamar makes use of a centralized management model for controlling, configuring, scheduling, and monitoring remote office Avamar backups. Organizations commonly deploy a centralized Avamar system in the datacenter and smaller Avamar systems or just Avamar software agents at smaller remote offices. These distributed Avamar systems can be managed and monitored via the Avamar Enterprise Manager interface, and administrators can quickly drill down into a particular system to view more granular status or manage specific aspects of a remote Avamar system.
- ☒ **File system interfaces.** Avamar provides a familiar Windows or Linux file system interface to all backups stored in the Avamar server via the Avamar file system (AvFS). By using standard interfaces, Avamar can be leveraged as a strategic, enterprise data store by third-party applications. The Avamar file system interface facilitates the following:
 - ☐ **Direct to tape output.** Utilizing the Avamar file system (AvFS), a standard tape backup application can read Avamar backup data stored on disk and stream it directly to tape without the need to bring the data back through the client or application server. To standard backup applications, such as Networker, NetBackup, or TSM, the Avamar file system looks like any other server or NAS system. Once data is archived to tape via this method, data can be restored directly from tape to target clients without involving Avamar.
 - ☐ **Search and index.** Third-party search and index tools, such as Google and FAST, can be utilized with Avamar for granular search and retrieval of file content out of the Avamar file system. This search may be for regulatory audits or in response to legal discovery requests.
 - ☐ **Direct restore.** Administrators can utilize the file system view of backups to perform direct restores of individual data files or volumes (the administrator can select the level of granularity) back to the original systems or to different systems.

- ☒ **Tape integration.** Avamar provides a number of options for archiving data to tape, including direct tape output, media server tape output, and checkpoint backups. The most popular method for performing tape archival is direct tape output, which is described above.
- ☒ **One-step restore.** Unlike traditional tape storage solutions, backup data stored by Avamar is ready for immediate recovery or access on demand — eliminating the time-consuming process of restoring full and subsequent incremental backups or the need to stream thru a tape in order to locate a specific file.
- ☒ **Broad file and database support.** Avamar intelligent client agents support Unix, Windows, and Linux operating systems along with Oracle, SQL, Exchange, and DB2 databases and can also backup NetApp filers via NDMP along with VMware environments. Commonly, servers at remote offices are standard file and print servers or servers running basic networking functions. However, in remote environments where there is enough application traffic, applications and their associated storage may run locally. Avamar can support both file and application data residing at remote offices.
- ☒ **Encryption.** In today's regulatory-intensive environment, firms cannot afford to face exposure from lack of corporate controls or sound governance practices. This holds true for data across the enterprise, no matter if it resides in a datacenter, on a laptop, or in a remote office. Avamar protects backup data by encrypting it while in transit across the LAN/WAN and at rest. By encrypting backup data at rest, organizations are further protected from data theft or unauthorized access.
- ☒ **Verified recovery.** Many times a recovery may fail due to a backup that was not verified. During a recovery is not the ideal time to discover such a failure. Avamar addresses this problem through automated data integrity checking. Avamar executes integrity checks for all data stored within Avamar. According to EMC, this ensures 100% integrity so that data can be quickly and reliably restored when needed. The integrity checking application verifies that all data backups can be restored to their original state, an exacting process that would be very difficult and expensive to accomplish with removable-media solutions. If the integrity of any data is discovered to be compromised during the regular validation process (e.g., disk failures, system issues, etc.), the issue can be promptly addressed before data recovery is needed.
- ☒ **Small and large remote office support.** It's not uncommon for smaller remote offices to have limited server infrastructure but a larger quantity of desktops or workstations supporting critical business, sales, or research functions. Avamar scales up or down to meet the needs of different remote office environments. A single Avamar approach can be deployed to support remote offices of different scales.
- ☒ **Offsite replication.** Replicator software enables efficient, encrypted, and asynchronous replication of heterogeneous data stored in an Avamar Server to another Avamar Server deployed in an offsite location for disaster recovery and business continuity.

Challenges

Challenges in the development and adoption of a remote office data protection strategy include the following.

Deciding on the right approach - Remote offices can each take on very different personalities in terms of users, infrastructure and data types. Firms must inventory and evaluate assets in different remote office locations and determine if the re-centralization of remote office data to a data center while remote offices access and share files over a WAN is appropriate. This approach, known as wide-area file services (WAFS), makes files appear as if they were local to the remote office. While appropriate in some environments, a WAFS approach introduces IT management to the remote office with re-configuration of client desktops and ongoing maintenance of the WAFS system. WAFS by its very nature is file oriented and does not support non-file-based structured applications. Most importantly, WAFS requires a complete re-architecture of an enterprise's remote office IT and application management strategy, a process which may not be feasible due to organizational and management boundaries within an organization.

Legacy backup applications – Legacy backup applications may already be deployed in remote offices. These applications are entrenched in IT environments, in particular for larger firms with long or permanent data retention requirements. To de-install existing backup agents and server architecture with deployment of new software components can be a time-consuming process for a limited IT staff. Older data written in a legacy backup application format may need to be restored for legal or business reasons. This can require a firm to maintain the legacy backup product for restores while deploying the new data protection approach for backups on a go forward basis. Current best practice is for firms to take legacy tape data and restore it into new disk-based backup or archive repositories to consolidate backup data in a standard format and discovery schema.

SUMMARY

IT executives are taking control of remote office data protection. The cost and productivity impacts of downtime at the remote office can no longer be tolerated. New approaches to centrally control and manage WAN-based remote office backup and recovery ensure consistent, reliable backup and guarantee fast recovery. These new approaches offer data deduplication to reduce the size of backups sent over existing wide area networks thus making the centralized backup of remote office data a viable option.

Avamar's approach to data deduplication reduces backup data at the source. It enables enterprise organizations to efficiently backup remote office data using existing network bandwidth while reducing the required backup storage capacity throughout the organization. The solutions provided by Avamar ensure that remote offices of different sizes can reliably protect critical data, reduce overhead, and centrally manage backup operations. Avamar meets the challenges of remote office backup by addressing the people, process, and technology issues that make traditional remote data protection methods impractical. Avamar integrates with existing tape environments and enables organizations to immediately leverage their backup data in support of legal discovery or regulatory compliance objectives.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2007 IDC. Reproduction without written permission is completely forbidden.