



WHITE PAPER

Iron Mountain Delivers File Archiving Service

**By Brian Babineau
With Lauren Whitehouse**

February, 2009

Table of Contents

Table of Contents	i
Introduction	1
Why Archive File Content?	2
Control Storage Costs without Compromising Access.....	2
Enable Electronic Records Management Programs.....	2
Understanding Iron Mountain’s VFS Service	3
The Basics.....	3
Deploying VFS.....	3
VFS Service Retention Management Capabilities.....	4
Data Integrity and Preservation.....	4
Large-scale Data Movement.....	5
VFS Facilitates Compliance with SEC 17a-4	5
Why WORM?.....	5
SEC 17a-4 is the Standard.....	5
VFS for Broker/Dealers.....	6
Conclusion	8

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482-0188. This ESG White Paper was developed with the assistance and funding of Iron Mountain.

Introduction

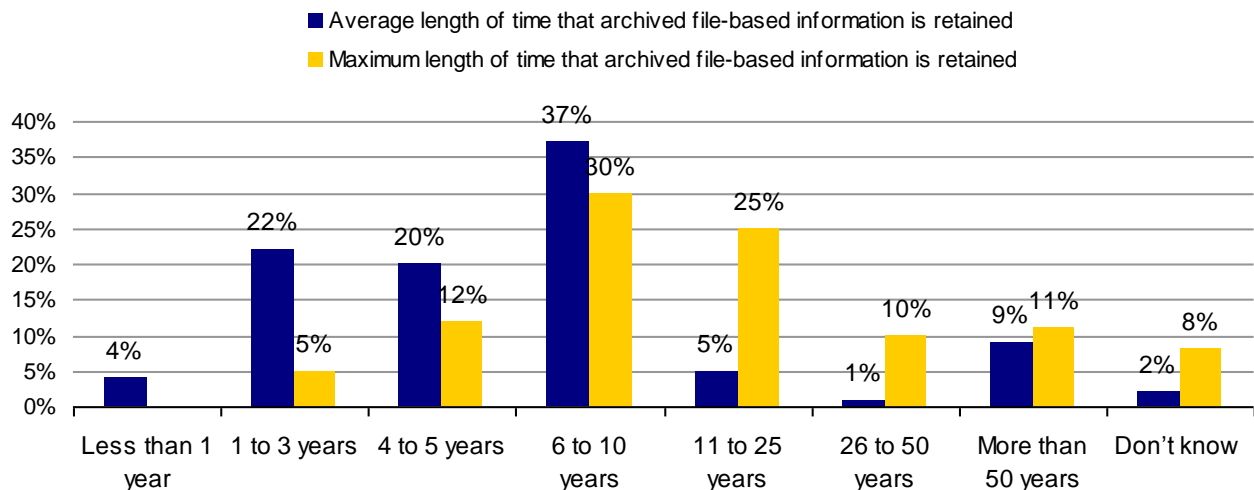
ESG estimates that file-based information will increase at a nearly 60% compounded growth rate over the next four years.¹ Although some content is not frequently accessed, this unstructured data consumes valuable storage resources and complicates data protection processes. Another subset of the file information consists of business records subject to industry regulations. Yet another portion of a company's unstructured data might fall under the scope of an ongoing electronic discovery request.

These challenges create a conundrum for IT: balancing the need to save more file data—in an accessible format—for longer periods of time with the need to control growing storage costs. To help strike this balance, companies have turned to purpose-built file archive solutions. These solutions identify a subset of information to be retained for a specified period of time and move this data to a secondary environment—sometimes located offsite—with the goals of storing it cost effectively and maintaining accessibility.

The downside of most file archive solutions is the implementation of additional technology infrastructure—particularly storage—that IT must manage. ESG estimates that organizations will archive nearly 80,000 petabytes of file content over the next three years,² with a majority of retention policies ranging between four and six years (see Figure 1). Running a large file archive adds another burden to IT's workload, in addition to managing primary file storage.

FIGURE 1. AVERAGE AND MAXIMUM RETENTION PERIODS FOR ARCHIVED FILE-BASED CONTENT

To the best of your knowledge, what would you say is the length of time for which your organization retains archived file-based information? (Percent of respondents, N = 93)



Source: ESG Research Report, 2007 File Archiving Survey, December 2007

Companies can avoid the drawbacks of onsite file archive storage management by investing in an offsite information archiving service for long term retention. One such offering—the Virtual File Store (VFS) service available from Iron Mountain's Digital Division—boasts a unique architecture that simplifies file archiving processes as it removes the need for IT's involvement in the file archive infrastructure. The VFS service allows companies to move information off of primary file servers to a lower cost environment inside a secure Iron Mountain data center, set and enforce retention policies, and save information in Write Once Read Many (WORM) format.

¹ Source: ESG Research Report, *File Storage Market Sizing*, August 2008.

² Source: ESG Research Report, *2007 File Archiving Survey*, December 2007.

Why Archive File Content?

Control Storage Costs without Compromising Access

Traditional file archive processes involve moving old data to tape and saving it for extended periods of time. The problem with this approach is that the information is not easily accessible. Employees cannot use the information and if compliance officers or corporate counsel need to find files, it can take a while.

Tape was a logical choice for file archives because it costs less than disk. However, new accessibility needs—driven by compliance and electronic discovery requirements—and increasing availability of denser, less expensive storage solutions created a new opportunity to keep more information accessible without spending a disproportionate amount of money on storage.

File archive solutions enable customers to deploy various tiers of storage—each with their own cost and performance characteristics. Customers can archive (move) data to the appropriate tier, helping to align accessibility needs with the cost of saving information. Over 40% of archive customers surveyed by ESG store their file archives on disk for short term retention and then move the data to tape for longer retention periods.³ This keeps newer data more accessible while historical information is stored on less costly media.

Because archive solutions allow files to move from one storage system to another, customers can free up valuable primary file server capacity for newer data. With less information residing on the primary file server, backups run faster. IT can complete data protection operations within the allotted backup window. Of those organizations currently using file archive solution, 70% experienced improved primary file server performance.⁴

Enable Electronic Records Management Programs

Almost every organization creates business records that must be kept for compliance or business reference reasons, and most have traditionally saved these records in paper format. With more business being conducted electronically, records are now being born digital—printing them for retention purposes is no longer feasible.

Purpose-built archive solutions enable electronic records management programs because they can identify specific content that needs to be saved, then set and enforce a retention period—all while ensuring that the information is accessible throughout the retention timeline.

A subset of business records, file-based data, can include documents such as budgets, employment offers, customer account statements, and sales contracts that can be managed easily by a purpose-built archive solution. With file archives enabling an electronic records management program, companies can comply with record retention regulations such as those cited in the Securities and Exchange Act of 1934, the Financial Service Authority (UK) rules, the HIPAA Security Rule, and the Family Educational Rights and Privacy Act. There are also government-specific requirements, with examples being ‘Sunshine Laws’ enforced at the state level in the US and New South Wales, Australia’s Policy on Digital Records Preservation.⁵

Companies may be forced to retain information as part of a legal matter where electronically stored information falls under the scope of a discovery request. ESG research indicates that 60% of electronic discovery events involve general office productivity and application files.⁶ As a result, companies must be able to identify, retain, and preserve (prevent deletion or modification) relevant unstructured content until the legal matter or regulatory investigation is resolved.

³ Source: ESG Research Report, *2007 File Archiving Survey*, December 2007.

⁴ Source: ESG Research Report, *2007 File Archiving Survey*, December 2007.

⁵ <http://www.records.nsw.gov.au/recordkeeping/>

⁶ Source: ESG Research Report, *Electronic Discovery Requirements Escalate*, November 2007.

Understanding Iron Mountain's VFS Service

The Basics

Iron Mountain's VFS service is a file archive offering that enables companies to remove data from primary file servers and retain it on storage residing at Iron Mountain data centers. Customers also have the option to establish retention policies for specific files to meet regulatory mandates, records management requirements, and legal hold instructions.

There are two major components of the VFS Service: the VFS Appliance and the VFS repository which resides on Iron Mountain's data center infrastructure. The VFS Appliance serves three roles, first as a target for data archival, second as a content mover, and lastly as the interface for retrieval. As soon as the VFS Appliance receives archived data, it moves the content to the remote repository. The VFS Appliance integrates with Microsoft Active Directory and manages all of the archive data (location, access permissions, etc.). In addition, the VFS Appliance tracks everything that occurs within it, including what data has been archived, where it came from, and who has accessed it. The VFS repository receives the archived data from the VFS Appliance and retains it for the appropriate period of time.

Deploying VFS

The VFS service is managed and operated by Iron Mountain. Iron Mountain coordinates the installation of a VFS Appliance on the customer's IP network, similar to a traditional file server. The VFS Appliance, which can support one hundred file shares and petabytes of data, is a stateless appliance with a large amount of cache storage designed to handle sequential I/O traffic—ideal for moving and retrieving archive data. There are three VFS Appliance options with different cache levels: G1 (250 GB), G3 (2 TB), and G5 (4 TB). Iron Mountain also configures the VFS repository inside its data center, which will receive data from the VFS Appliance.

The VFS Appliance exports CIFS and NFS file shares, which serve as targets for the archived data. Customers can then move data to the VFS Appliance by:

- Manual (or scripted) data movement from primary file shares, which enables primary applications or individual employees to send data directly to the VFS Appliance.
- Using file archive or backup solutions that typically run an agent on the primary file server tasked with determining what data is moved. In this scenario, the VFS Appliance file shares become targets for the data movement software. Iron Mountain has established partnerships with a few companies that can assist with these migrations—customers should check with Iron Mountain for an updated list of those solutions that are certified with the VFS Service. Customers may also choose to utilize software packages such as RoboCopy and SafeCopy for file movement to the VFS Appliance.
- Leveraging existing file system virtualization and management solutions, which serve as a global namespace for multiple file shares. The VFS Appliance file shares are managed under the global namespace, becoming the archive location in that schema. Iron Mountain has developed a partnership with F5's ARX file virtualization line, allowing customers to automate movement between primary file shares and VFS Appliance file shares under a single ARX namespace.

As it receives the archive data, the VFS Appliance sends it via a VPN to the designated VFS repository at an Iron Mountain data center, where it will be stored and then replicated to a second Iron Mountain site for disaster recovery purposes. The VFS service is committed to maintaining two geographically separate copies of archived data at all times.

Customers can retrieve data by going directly to the file share exported by VFS or via the backup/data mobility/file virtualization solution that removed the content from the primary file server in the first place. Iron Mountain saves all of the data online (on disk) facilitating swift, seamless access—a sharp contrast when compared to tape-based archives.

Although the VFS Appliance resides inside a customer's data center, VFS is a hardened appliance that is managed and monitored remotely by Iron Mountain. It integrates with a customer's existing Microsoft Active Directory implementation, ensuring that file access permissions are extended from the primary environment to the archive. The G5 also has multiple NIC cards that can be bonded for additional aggregate throughput and incremental NFS security controls.

VFS Service Retention Management Capabilities

When authorized administrators configure NFS and CIFS exports on the VFS Appliance, they can also establish a retention period for each share. By setting the retention period on a file share, the administrator makes the file system 'read-only.' The 'read-only' status means the data will be stored in WORM format as access is still permitted, but no action (i.e., deletion or modification) can be taken. Each file is kept in WORM format for the retention period based on when it is archived to VFS. For example, if a file share has a 30-day retention period and a file is moved to it on January 1, it will be saved as 'read-only' until January 30; a file archived to VFS on January 2 will be retained as 'read-only' until January 31.

Because each file share can have its own retention policy, customers can address a multitude of records management requirements. One file share may be set up for finance with a retention period of 90 days while human resources may leverage its own file share with a retention period of two years to meet specific employment laws.

The minimum VFS retention period per file share is 30 days, but it can be limitless (data cannot be deleted, ever). Currently, customers cannot reduce the retention policy; however, it can be extended, which is extremely helpful if information is being requested as part of an ongoing legal or regulatory discovery process. Very often, regulatory investigations last longer than a business record's retention requirement and, therefore, the retention period must be extended.

Customers can implement VFS to enforce legal holds in two different ways. The first is to set up a share with a long or indefinite retention period and copy over potentially relevant content from a primary file share. There, it can be preserved until matter is resolved. Alternatively, customers may choose to enforce legal holds on existing archived data. In this situation, existing retention periods on specific file shares may need to be extended until the matter resolved. Organizations may even choose to use a combination of both approaches, depending on what data falls within the scope of an electronic discovery request. Regardless of the option, corporate counsels and IT can easily place files on legal hold with the VFS.

Data Integrity and Preservation

If an organization wants to deploy archiving to meet record retention and legal hold requirements, a solution must be secure, complete automated data integrity checks, and audit all activity within the solution. VFS has several capabilities that address all three requirements:

Security

- The VFS Appliance is a Linux server with most services, such as Apache Web server, turned off. Additionally, Iron Mountain has removed potential vulnerabilities from the SUSE Linux Kernel to prevent the risk of operating system-level breaches. Certain network services and user accounts within the operating system are also disabled on the VFS Appliance.
- Only Iron Mountain can maintain the VFS Appliance; customers only have permission to access data, and view reports.
- The VFS Appliance sends data to the remote repository using a VPN. Public key encryption is used for mutual authentication between the VFS Appliance and the repository, and the data is transmitted and stored using 128-bit AES encryption. The latter prevents the archived data from being read while it is being sent from the VFS Appliance to the repository and it also eliminates the possibility that an Iron Mountain employee can access a customer's data once it is stored.
- When the data reaches the Iron Mountain data center, it passes through a firewall and is stored within a US-based facility hundreds of feet below ground. These facilities contain several physical and technological information protection features, including armed guard details, fire suppression systems,

redundant ISP connections, and emergency generators that can run the data center for up to seven days.

Integrity

- Every file is stored with a unique signature, which serves as its fingerprint. This fingerprint is checked periodically during the content's retention period, including when it is replicated to the second data center and accessed. If compliance officers or corporate counsel need to prove the authenticity of a file, they can run a report comparing the signatures of the file when it was stored and when it is retrieved from the archive.
- All data is encrypted as it is transmitted and stored, also mitigating the risk of unauthorized access.
- Initial file controls are established by the VFS Appliance—the G1 and G3 appliances use POSIX ACLs and the G5 uses NFS V4 ACLs.

Auditability

- Any attempts to access the VFS repository, as well as component audit messages, are tracked. Each of the VFS Appliance's hardware components generates alerts on system activity. All of this data is viewable (not alterable) by an employee with administrative network view credentials.

Large-scale Data Movement

When an organization first implements a file archive, a large amount of data may need to be moved initially. The same may hold true if a company needs to retrieve a significant volume of content in response to a regulatory or legal inquiry. When dealing with a service, a large bulk data transfer could take time and consume valuable network bandwidth.

The Iron Mountain VFS Data Shuttle option enables customers to physically ship their information to Iron Mountain's data center. The Data Shuttle service will encrypt the data before it is sent, create a checksum algorithm so that the data can be authenticated when it arrives at the data center, provide tamper proof containers to transport the physical storage media, and track the actual shipment via a trusted third party carrier. When the data arrives at the Iron Mountain facility, the operations staff will run the aforementioned checksums to ensure that all the sent data has arrived. The process can then be reversed if the customer requests data from Iron Mountain, with integrity checks being completed at the customer site.

VFS Facilitates Compliance with SEC 17a-4

Why WORM?

ESG has yet to encounter a regulation that requires organizations to store business records in WORM format. However, saving data in WORM format helps customers preserve the integrity and authenticity of business records because it guarantees that content is not altered or deleted during the assigned retention period. Not only does this help companies ensure that business records are properly saved, but it also assists corporate counsels in enforcing legal hold mandates during electronic discovery.

SEC 17a-4 is the Standard

17 CFR 240.17a-4—more commonly known as SEC Rule 17a-4—establishes standards for the proper retention and preservation of business records. The rule, which applies to registered broker/dealers, contains several storage media requirements that must be adhered to, including 17a-4(f)(2)(ii)(A), which mandates that the storage media *'preserves the records exclusively in a non-rewriteable, non-erasable format.'*

There have been several interpretations of the 'non-rewriteable, non-erasable' storage media requirement, with many of the SEC issuances being indifferent to the technology approach used by broker/dealers to meet this specific mandate. Solutions that stored data in a WORM format became an accepted form and the most recent

SEC interpretation release issued on May 7, 2003 even went so far as to say that the ‘non-erasable, non-rewriteable’ mandate could be met by an integrated (hardware and software) solution: ‘A broker-dealer would not violate the requirement in paragraph (f)(2)(ii)(A) of the rule if it used an electronic storage system that prevents the overwriting, erasing or otherwise altering of a record during its required retention period through the use of integrated hardware and software control codes. Rule 17a-4 requires broker-dealers to retain records for specified lengths of time. Therefore, it follows that the non-erasable and non-rewriteable aspect of their storage need not continue beyond that period.’⁷

The latest interpretation allows broker/dealers to use purpose-built archive software and disk storage that stores data in WORM format to meet the ‘non-rewriteable, non-erasable’ requirement. By using an integrated solution versus tape or optical—the more traditional WORM storage media—broker/dealers can keep their archives more accessible, better preparing them in the event of a regulatory inquiry.

VFS for Broker/Dealers

Broker/dealers that want to leverage Iron Mountain’s VFS Service to retain business records should evaluate the solution for compliance with the entire 17a-4 rule; however, ESG believes that there are two sections within the rule that broker/dealers should focus on when selecting an archive solution for records retention: 17CFR240.17a-4(f)(2) and 17CFR240.17a-4(f)(3). The following is a description of how VFS could be implemented to meet the requirements outlined in these sections. Please note that this reflects ESG’s opinions and not the opinions of the SEC. As such, ESG’s opinions should not be construed as legal advice, but rather as education material. (Note: ESG obtained the rules from the updated Code of Federal Regulations located at <http://www.gpoaccess.gov/cfr/index.html>.)

Rule

17CFR240.17a-4(f)(2) - *If electronic storage media is used by a member, broker, or dealer, it shall comply with the following requirements:*

- (i) *The member, broker, or dealer must notify its examining authority designated pursuant to section 17(d) of the Act (15 U.S.C. 78q(d)) prior to employing electronic storage media. If employing any electronic storage media other than optical disk technology (including CD-ROM), the member, broker, or dealer must notify its designated examining authority at least 90 days prior to employing such storage media. In either case, the member, broker, or dealer must provide its own representation or one from the storage medium vendor or other third party with appropriate expertise that the selected storage media meets the conditions set forth in this paragraph (f)(2).*
- (ii) *The electronic storage media must:*
 - (A) *Preserve the records exclusively in a non-rewriteable, non-erasable format;*
 - (B) *Verify automatically the quality and accuracy of the storage media recording process;*
 - (C) *Serialize the original and, if applicable, duplicate units of storage media, and time-date for the required period of retention the information placed on such electronic storage media; and*
 - (D) *Have the capacity to readily download indexes and records preserved on the electronic storage media to any medium acceptable under this paragraph (f) as required by the Commission or the self-regulatory organizations of which the member, broker, or dealer is a member.*

ESG Opinion

The VFS Service would fall under the most recent SEC interpretation for the ‘non-erasable, non-rewriteable’ requirement. Broker/dealers can store records within a file system that has a retention policy capable of being enforced on a per file basis. The records will be kept in ‘Read-only’ format, preventing deletion or modification. Only authorized users can create retention policies and once a retention period on a file share is established, it cannot be reduced nor can the file share be deleted. The VFS Appliance is extremely difficult to access as the appliance has been hardened, making it difficult for anyone to make changes to retention policies or records managed by the solution.

Constant data integrity checks via the digital signature allow the VFS Service to verify the accuracy and quality of the record archiving process. The digital signature is based on the record and the date the record was archived

⁷ Source: <http://www.sec.gov/rules/interp/34-47806.htm>.

to the VFS. Iron Mountain uses a universal clock service to timestamp all of the archived records and it is this timestamp that the beginning of the retention policy is based upon.

The VFS Appliance tracks where all the business records are stored, supported by full auditing capabilities including the file's location, who has access to it, and whether or not any actions—other than read requests—(deletions, modification, etc.) occur against the data. The audit capabilities also include notifications on the access to VFS Appliance components so broker/dealers can see if any configuration changes were made to the system. All of the audit trail reports are viewable by submitting a special request to Iron Mountain.

To access the records, a customer can utilize file paths to locate a specific file. In addition, the file system paths (record paths that can serve as the index) can be viewed by authorized employees. The file system paths and the business records themselves can be exported onto any portable storage device with the file system properties—namely, the retention policies—intact. Customers may choose to use Iron Mountain's Data Shuttle option to download content for regulators, if necessary.

Rule

17CFR240.17a-417a-4(f)(3) - *If a member, broker, or dealer uses micrographic media or electronic storage media, it shall:*

- (i) *At all times have available, for examination by the staffs of the Commission and self-regulatory organizations of which it is a member, facilities for immediate, easily readable projection or production of micrographic media or electronic storage media images and for producing easily readable images.*
- (ii) *Be ready at all times to provide, and immediately provide, any facsimile enlargement which the staffs of the Commission, any self-regulatory organization of which it is a member, or any State securities regulator having jurisdiction over the member, broker or dealer may request.*
- (iii) *Store separately from the original, a duplicate copy of the record stored on any medium acceptable under Sec. 240.17a-4 for the time required.*
- (iv) *Organize and index accurately all information maintained on both original and any duplicate storage media.*
 - (A) *At all times, a member, broker, or dealer must be able to have such indexes available for examination by the staffs of the Commission and the self-regulatory organizations of which the broker or dealer is a member.*
 - (B) *Each index must be duplicated and the duplicate copies must be stored separately from the original copy of each index.*
 - (C) *Original and duplicate indexes must be preserved for the time required for the indexed records.*
- (v) *The member, broker, or dealer must have in place an audit system providing for accountability regarding inputting of records required to be maintained and preserved pursuant to Sec. Sec. 240.17a-3 and 240.17a-4 to electronic storage media and inputting of any changes made to every original and duplicate record maintained and preserved thereby.*
 - (A) *At all times, a member, broker, or dealer must be able to have the results of such audit system available for examination by the staffs of the Commission and the self-regulatory organizations of which the broker or dealer is a member.*
 - (B) *The audit results must be preserved for the time required for the audited records.*
- (vi) *The member, broker, or dealer must maintain, keep current, and provide promptly upon request by the staffs of the Commission or the self-regulatory organizations of which the member, broker, or broker-dealer is a member all information necessary to access records and indexes stored on the electronic storage media; or place in escrow and keep current a copy of the physical and logical file format of the electronic storage media, the field format of all different information types written on the electronic storage media and the source code, together with the appropriate documentation and information necessary to access records and indexes.*

ESG Opinion

A retention policy is assigned to a file system so that all parts of the file system, including the file system paths and the records themselves, are saved in read-only format for the duration of the retention period. With all of the records stored in standard file systems, they can be read or printed with the appropriate VFS permissions.

By default, the VFS Service creates two copies of all the file systems and business records saved within those file systems. The copies of the file system (and file system paths) and the business record are stored in separate geographical data centers. The VFS Appliance tracks where all of the business records are kept for easy retrieval purposes.

The constant process of checking record integrity, along with VFS Appliance access and component logs, provides a comprehensive audit trail for broker/dealers to report on the authenticity of the information being retained. Customers can generate reports to show what data was archived, if any data was missed, and prove that the data that was stored is the same data being retrieved or produced (if necessary).

Iron Mountain maintains and escrows the source code of the VFS Appliance and can provide the logical data structure to the broker/dealer if regulators request it.

Summary

The aforementioned rules are only a subset of the requirements in 17a-4. Broker/dealers must conduct their own due diligence. If a broker/dealer chooses to deploy Iron Mountain's VFS service to retain certain business records, the broker/dealer must notify its examining authority 90 days prior to implementation.

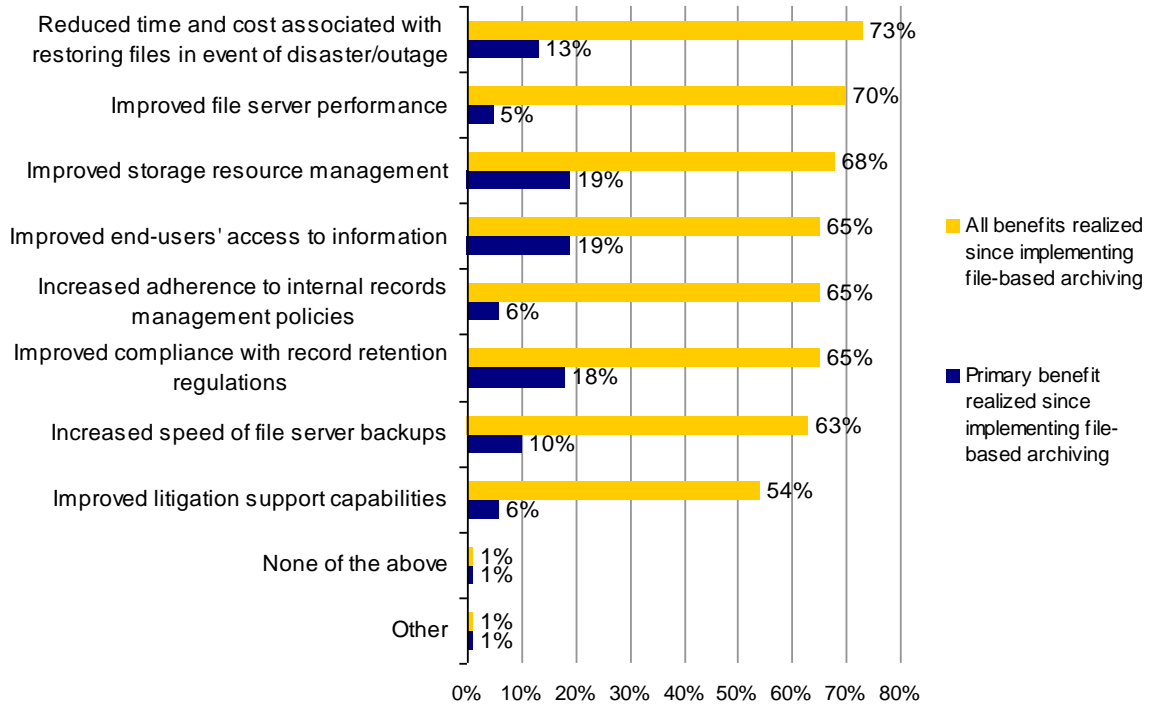
Conclusion

Unstructured data growth is unlikely to subside any time soon, forcing IT to find budget for new file servers to save the information. Record retention regulations and electronic discovery mandates, also unlikely to decrease, compound storage capital and operating costs because they mandate keeping data for longer periods of time. IT needs ways to identify subsets of files that are business records, subject to electronic discovery requests, or inactive and move this data to a separate, lower cost environment.

File archiving solutions provide a means to accomplish these tasks, with the added benefit of managing retention policies. Existing file archive customers attest to the benefits of these solutions (see Figure 2), however, companies are realizing that the archive storage environment can be just as complex to manage as the primary one due to the amount of data being saved. Iron Mountain addresses this challenge with its VFS service. Combining Iron Mountain's service delivery capabilities—inclusive of its extremely secure data centers—with its unique VFS Appliance, the VFS service delivers all of the benefits of file archiving without the internal IT operational hassle.

FIGURE 2. FILE-BASED ARCHIVING BENEFITS REPORTED BY CURRENT USERS

What benefits has your organization realized since implementing processes and technologies to archive file-based content? (Percent of respondents, N = 93)



Source: ESG Research Report, 2007 File Archiving Survey, December 2007

The VFS service is ideal to underpin electronic records management programs because it allows customers to set and enforce retention periods on a per file basis. VFS can store data in WORM format, which prevents information from being altered or deleted during the assigned retention period, facilitating compliance with strict records retention regulations such as SEC Rule 17a-4, as well as legal preservation requirements.

With so many possible benefits, organizations should not be considering whether or not they should archive files; rather, they should be determining what solution best meets their requirements. There are plenty of file archive options in the marketplace, inclusive of product (on premise) and service offerings. Because it is simple to implement, requires minimal IT resources to operate, and is supported by Iron Mountain's service delivery capabilities, the VFS service should be one of the options considered.



20 Asylum Street
 Milford, MA 01757
 Tel: 508-482-0188
 Fax: 508-482-0218

www.enterprisestrategygroup.com