

10 Building Blocks for Securing File Data

Introduction

Securing file data has never been more important or more challenging for organizations. Files dominate the data center, with analyst firm IDC¹ estimating that unstructured data accounts for approximately 80% of all enterprise data and is growing at 60% per year. Managing this data is challenging not only because of its volume, but also because there is a lack of basic visibility and control over file data for most organizations, despite many decades of file use and a vast array of file protection technologies.

From file system access control lists (ACLs), to encryption, to data loss prevention solutions, many technologies are available to help secure file data. But, even organizations with these technologies in place find it challenging to protect their file data. This is because three fundamental capabilities are lacking in most organizations:

1. *Operationally efficient file activity monitoring and auditing*
2. *Scalable user rights management for files*
3. *Automated business policy enforcement for file data*

These three capabilities are core components of the emerging File Activity Monitoring market, and form the basis of a phased approach to file security.

This guide describes ten phases for securing file data, including how and when to use these basic capabilities, as well as when to deploy other complementary technologies.



¹ IDC "2009 File-Based Storage Taxonomy", November 2009

Phased Approach

Given the volume of file data, the large number of file protection technologies available to companies, and IT budget constraints, it is imperative to identify which file security solutions to deploy, and when, in order to maximize security while minimizing capital and operating expenses.

One of the areas where there has been a great deal of interest and media coverage is Data Loss Prevention (DLP) products. DLP products play an important role in an overall file data protection strategy and, according to leading analyst firms, are particularly good at preventing inadvertent or accidental leaks. Still, an effective DLP deployment can involve lengthy roll-out and tuning, which is why analysts and businesses deploying DLP recommend a phased deployment approach.

The phased approach to file security outlined below optimizes the balance between security and costs by applying File Activity Monitoring and its three critical levels of file data control: file activity monitoring, user rights management and policy-based controls. File Activity Monitoring complements existing file security solutions, providing a comprehensive record of every file access by every user, an efficient way to align access rights with business policy, and an automated way to alert on, and respond to, policy violations.

Phase 1: Start in the Datacenter

The most concentrated repository for sensitive file data in organizations is the shared file servers and network attached storage devices in the datacenter. Therefore, placing effective File Activity Monitoring controls on this centralized data provides a great deal of leverage. For example, reducing access rights to business need-to-know levels in the datacenter helps prevent users without a business need from ever having access to the data in the first place. Limiting access to those users with legitimate business needs helps limit the spread of sensitive data. And, by auditing access activity to this data, organizations establish a forensic audit trail that can be used in the event that a legitimate user chooses to abuse their access rights.

Preventing Data Breaches with File Activity Monitoring

File Activity Monitoring delivers important capabilities not found in DLP solutions, in part by taking a different approach to data breach prevention. The table below highlights some of the areas where these solutions differ from each other.

Capability	DLP	File Activity Monitoring
Transparent to file data users	Intentionally non-transparent; meant to change user behavior. Requires business involvement in the planning and implementation phases to be successful.	Transparent to file data end users. Does not represent a cultural shift or require changes to business policies.
Data owner identification	While DLP implementations require working with business data owners, DLP solutions do not identify the owners of data.	Helps identify file data owners, who can in turn facilitate DLP planning and implementation and can participate in data owner rights reviews, or "attestations".
User rights management	DLP solutions do not address user rights management.	Helps establish business need-to-know access for file data. Audits user rights, facilitates rights reviews and identifies excessive rights and dormant users.
Comprehensive access activity audit trail	Typically monitors only that data designated as sensitive.	Monitors all file access activity by all users to produce a comprehensive forensic audit record.
Address unintentional and malicious activity	Primarily targets data loss due to accidents and ignorance.	Addresses malicious insider threats by establishing business need-to-know access, monitoring all user activity and enforcing security policies.

Phase 2: Begin with the Data you Understand

A logical place to start within the datacenter is with data that you already know is sensitive, and for whom an owner is known. Virtually every organization has data that falls into this category. Examples include regulated file data such as credit card numbers, customer information and financial data, and sensitive information such as legal documents, business plans, and intellectual property. Focusing on this data first enables you to have an impact immediately. The return on your efforts will be quick for two key reasons. First, you avoid the typically time consuming task of content discovery because you are working with data stored in a well-known location (e.g., in a certain set of shares or folders), that you already know is sensitive. The second factor contributing to immediate results and time savings is that you are working on data with known owners. Because data owners are the individuals who best understand the business relevance of their data, you know exactly who to work with to develop proper protection and control strategies.

Phase 3: Monitor Data Access

The third phase of securing file data is to monitor all file access activity. This provides instant value by creating an audit trail that can be used at any point for ad-hoc forensic analysis and reporting. Leading File Activity Monitoring solutions deliver transparent auditing that does not require modification of file servers, clients or applications, nor does it impact performance the way native auditing does. In addition to the forensic value of the audit trail, access activity provides intelligence about who is using, and not using, file data. This information can be used to recognize dormant users and unused access rights, as described below in Phase 6 on establishing business need-to-know access. Additionally, as you move beyond data with known owners, the audit trail helps identify data owners: users accessing specific files or folders most often are frequently the owner or can easily identify the owner.

Phase 4: Classify Data

Data classification is helpful for managing sensitive data, and is especially powerful when you have the ability to take action on the classification information via File Activity Monitoring policies. By selecting data in phase 2 that is known to be sensitive, you are able to postpone the need to perform content-based data discovery until a later phase (see phase 9 below on discovering sensitive data based on content). Depending upon available resources, this data can even be analyzed and protected at the same time content-based discovery is being rolled out and tuned. In any case, you can simply designate the data from phase 2 as sensitive based on metadata such as its location and type (e.g., all Excel files in the folder named "Transactions" are classified as "PCI Data"), and then use this classification information in an actionable way, as discussed next in phase 5.

Phase 5: Implement Security Policies

File Activity Monitoring solutions typically have policies that can be triggered by a wide range of information, such as the name of the user accessing a file, their department, the location of the file, the time of day, the file's classification, etc. Organizations should use these policies to automate business process enforcement and alert on violations. For example, policies can be set to respond when users outside of certain, authorized departments access PCI data, or when access rights to sensitive data are modified. File Activity Monitoring policies that are enforced in real-time allow organizations to respond instantaneously to policy violations, even blocking access if desired.

Phase 6: Establish Business Need-to-Know Access

Every organization with a shared file system already depends on access control lists (ACLs) to control access to shared file data. The most significant challenge with this basic, built-in security mechanism is that ACLs rapidly fall out-of-synch with business changes and no longer reflect business need-to-know access. Therefore, phase 6 involves auditing and managing user rights at the ACL level, which is a fundamental File Activity Monitoring capability. File Activity Monitoring solutions analyze file system ACLs, and the corresponding users and groups from directory services (e.g., Active Directory), to build a comprehensive view of who has access to files and what type of access they have. This simplifies the processes of identifying excessive access. For example, data that is accessible by groups such as the “Everyone” and “Domain Users” groups in Active Directory is probably overly exposed.

In addition to identifying broad categories of excessive access such as Everyone group access, File Activity Monitoring solutions combine auditing details with rights information to help organizations identify rights that have been granted, but are never used. For example, a user with “delete” rights to a certain folder who never deletes data there may not need that level of access rights. Similarly, if a user with access rights is never seen accessing data, it may be an indication that the user is actually dormant and should be decommissioned from the directory services.

The final step in ensuring business need-to-know access is to engage data owners in performing a user rights review (these are sometimes referred to as “attestations” in industry regulations, and are also called “entitlement reviews”). During this process, the collected rights information is reviewed by data owners who can then validate whether or not access is aligned with business requirements, and make corrective recommendations.

Organizations can leverage the classification done in phase 4 to prioritize and focus the rights review process on the most sensitive information. For example, data classified as being in scope for PCI can be analyzed for excessive access, and permissions can be reduced there first to ensure the organization remains compliant.

Phase 7: Monitor Permission Changes

Once an organization has established business need-to-know access for their data, the biggest challenge will be to keep access aligned with business requirements. Access rights to file data are constantly in flux due to changing job roles and responsibilities, the addition of new data and projects requiring user collaboration. Furthermore, rights changes are very often made by the IT Help Desk as part of a process that doesn’t involve the Security staff and, in many cases, may not involve the data owners. The only way to keep pace with these changes is via automated monitoring. By establishing security policies within a File Activity Monitoring solution, stakeholders such as Security staff and data owners can receive real-time information about permission changes to sensitive business data. With a timely notification, stakeholders can investigate and correct permission changes that don’t align with business need-to-know access.

Phase 8: Identify Sensitive Data Based on Context

Having used the phases above to secure data that is already known to be sensitive, you can then focus attention on other data. A logical place to begin is with data that is heavily used by the same owners you have already been working with in the earlier phases. With a full audit trail from a File Activity Monitoring solution, reports can be run and shared with data owners to highlight the data they use most frequently. Any of this data that is identified as sensitive can then have phases 4 through 7 applied to it as well to ensure it is properly protected.

Phase 9: Discover Sensitive Data Based on Content

Locating pockets of sensitive data within the datacenter that were not identified through the phases above requires content-based discovery. Most companies that undertake this process with unstructured data use a Data Loss Prevention (DLP) solution. DLP solutions frequently provide a data discovery component as either a core or optional module. Additional approaches include stand-alone content discovery solutions, enterprise search products, e-Discovery solutions, etc. Once data has been discovered with one of these approaches, Phases 4 through 7 can be applied to protect the data. For data without known owners, audit data can be used to identify the owners which, as noted previously, are often those users accessing specific files or folders most often.

Phase 10: Beyond the Datacenter

Limiting access and monitoring use of file data in the datacenter is a critical starting point for effective file security. With that data protected using the phases outlined above, organizations can turn their attention to file data located on end-points. DLP solutions are the most popular approach to securing this data, and provide protections such as monitoring and control over what data can be copied to removable storage devices, printed, attached to email, etc. Enterprise Rights Management solutions are another end-point protection measure, and provide the ability to manage and enforce fine grained access control of sensitive files within client applications. For example, these solutions can control who has the right to view a document, edit content, or copy content to a clipboard.

While end-point solutions are constantly evolving, end-points remain a challenging location from which to control file data. Common approaches typically involve the installation and management of client-side security software, which can result in administrative challenges related to interactions between those security technologies and other agents and applications required on the end-points. A final challenge is that once sensitive data makes its way out of the datacenter and onto a mobile device, laptop or desktop, information leakage is hard to control because malicious insiders can take screen captures, photos, or simply make written or mental notes.

Conclusion

Securing file data has become a top priority for organizations facing a mounting volume of sensitive unstructured data. Given a basic lack of visibility and control over files, businesses often spend a great deal of capital and operational expenses working with file protection technologies that ultimately leave their file data vulnerable to malicious insiders and falling short of compliance mandates.

The phased approach described in this guide, when implemented using a File Activity Monitoring solution, enables organizations to maximize security while minimizing capital and operating expenses. File data is secured starting from the core – the datacenter – to provide the greatest security, leverage and efficiency.

Imperva SecureSphere File Security makes it easy to implement a phased approach to file security with capabilities that include file activity monitoring, user rights management, automated, real-time policy enforcement, and enterprise-class reporting and scalability.

About Imperva

Imperva is the global leader in data security. Our customers include leading enterprises, government organizations, and managed service providers who rely on Imperva to prevent sensitive data theft by hackers and insiders. The award-winning Imperva SecureSphere is the only solution that delivers full activity monitoring for databases, Web applications and file systems.

To learn more about Imperva's solution visit <http://www.imperva.com>.

Imperva

Headquarters
3400 Bridge Parkway, Suite 200
Redwood Shores, CA 94065
Tel: +1-650-345-9000
Fax: +1-650-345-9004

Toll Free (U.S. only): +1-866-926-4678
www.imperva.com

© Copyright 2011, Imperva

All rights reserved. Imperva, SecureSphere, and "Protecting the Data That Drives Business" are registered trademarks of Imperva.

All other brand or product names are trademarks or registered trademarks of their respective holders. #WP-10-BUILDING-BLOCKS-SECURING-FILE-DATA-0211rev1

