Five Signs Your File Data is at Risk

Executive Summary

Persistent insider threats and regulatory compliance mandates make protecting sensitive file data a business requirement for virtually every organization. However, the sheer volume of file data and its rapid and continuous growth make it a challenge to secure properly. Below are five questions to help you assess your file security posture. If you aren't able to answer these five questions confidently, your file data is probably at risk.

Background

File data accounts for approximately 80% of business data, according to market analyst firm IDC[1], and is growing at 60% per year. For example, if you have a single terabyte of file data today, you'll have over 10 terabytes of file data five years from now. That staggering pace explains why most organizations are challenged to protect their sensitive file data.

File security dynamics:
Malicious insiders target valuable files
Regulations mandate protections
Volume and growth of files is overwhelming

Insider threats are the most significant factor driving the need to protect this data. The prospect of personal, financial and professional gain can drive insiders to abuse their access rights, as demonstrated by numerous media reports. Consider the following examples from the recent past:

- A former Ford Motor Company worker was indicted for taking over 4,000 documents with future automotive designs which he used to help land a new job with a different automaker in China[2].
- A former Goldman-Sachs worker was arrested for taking source code files for a Goldman-Sachs trading system so he could take them with him to his next employer[3].
- The Canada Revenue Agency, the national tax authority, reported that Agency workers had been downloading revenue and tax related files for personal gain and to provide preferential treatment for relatives and friends[4].

Regulatory compliance is another major file security driver. Numerous regulations address data security and, in general, do not limit their scope based on data format: they apply equally to files, databases and applications. Sarbanes Oxley (SOX), Payment Card Industry Data Security Standard (PCI DSS), the U.S. Health Insurance Portability and Accountability Act (HIPAA) and many others specify that sensitive data must be protected. So, even if an organization uses financial applications and databases to manage their finances, once financial data that is governed by SOX is exported to a spreadsheet for analysis or reporting, that spreadsheet is "in scope" for compliance purposes.

---

1 IDC: "2009 File-Based Storage Taxonomy", November 2009
2 Wall Street Journal, "Ex-Ford Worker Indicted in Secrets Theft", October 16, 2009
3 Wall Street Journal, "Ex-Goldman Worker Is Arrested", July 7, 2009
4 Toronto Star, "Rogue tax workers snooped on ex-spouses, family members", June 20, 2010

Is your file data at risk?

The following five sections pose questions to help you assess your file security posture, and provide insight into how organizations are approaching file security today. They also highlight ways to overcome the limitations of conventional file security approaches. If you find that your organization isn't able to address these five questions reliably, your file data is probably at risk, and you should consider taking additional measures to ensure business need-to-know access.

1. Who owns your file data?

Data owners:
Understand the value of the data
Know who should / shouldn't have access
Are difficult to identify for most data

Data owners are critical to accurately protecting files because they best understand the data and its relevance to your business. If you don't know who your data owners are, it's hard to properly secure your file data and to make well informed data management decisions. For example, on the security front, you won't be able to establish business need-to-know access without owner input. You can only guess who should and shouldn't have access to the files. It will also leave you at a loss operationally when it comes to data migrations, establishing archiving policies and knowing when file data is no longer relevant and can be purged. When data owners are known, they are able to work with the groups responsible for protecting and managing data, usually Compliance teams, Security staff, System Administrators, and Storage Administrators.

In most organizations, data owners are known for only a small fraction of file data. This is because as job roles and data change, it becomes difficult to tell who owns data at any instant. While there are some pieces of meta-data that provide clues to data ownership (e.g., file system owner information, file locations, file names, etc.), they are merely clues and, in the end, the hunt for owners usually leads to mass emails or phone calls looking for someone to claim ownership. That's simply not an efficient, scalable or easily repeatable business process.

Ultimately, the most efficient and accurate way to identify ownership is to determine who is actually using the files the most. The top data users are either going to be the data owners or, as the primary business users of the data, are going to able to identify an owner almost instantly.

2. Who is actually using your files?

Auditing access activity:
Slows down server performance
Generates mountains of details
Typically not done, or done too late

Establishing an audit log of who is accessing file data is vital for security, compliance and IT operations. Security organizations need an audit trail for forensic investigations into data breaches and other security incidents, and to spot activity that violates business policy. Compliance teams need file activity auditing to validate and document that access activity complies with regulations. And, IT operations staff use audit details to track down problems such as files that have been modified, deleted or gone missing (also known as File Integrity Monitoring). Auditing also forms the basis of owner identification, as discussed above. However, without a continuous audit trail, organizations cannot perform these tasks reliably or efficiently.

Continuous activity auditing for files is difficult because native operating system auditing imposes a tremendous performance impact on file servers and generates a huge volume of difficult-to-interpret audit records. Stand-alone auditing solutions often face these same limitations. As a result, organizations typically turn auditing on only after an incident has already occurred in the hope of catching a repeat offense. This

results in an incomplete audit trail and won't help with events that occur only once. For example, if an employee downloads valuable files before they resign from an organization, turning auditing on after the resignation has been tendered does not help document the breach.

The only viable solution is to deploy an auditing system that can capture all access activity without impacting server performance, and which can distil mountains of audit detail into actionable information.

3. Who has the potential to access your files?

File access rights visibility is required by numerous regulations that address data security and is, in general, a data security best practice. By understanding the file access rights as set today on the file systems in your organization, you can see what your de-facto access policy is. The rights may not reflect the desired state of security, but they will tell you what users actually have access to right now. That visibility is necessary to begin the process of remediating excessive access and for demonstrating compliance with data security regulations.

Many organizations use built-in operating system tools to establish rights visibility. These tools – such as the Microsoft Windows Permissions dialog box – work a file or folder at a time, and are tedious to use with a large volume of data or users. Therefore some businesses use scripts written by administrators to capture rights information. While scripts require less manual activity, they don't scale well and reporting and analysis have to be performed by another system. Commercial rights collection point products have the benefit of being professionally maintained but, like scripts, tend to have limited scalability and lack robust analysis and reporting.

Rights visibility, analysis and problem remediation require a system that can automatically gather file rights across an entire enterprise – on each file server and NAS device. These rights have to be collected on an ongoing basis, consolidated and stored for analysis and reporting.

4. Whose access rights should be revoked?

In most organizations, user access rights to file data are far in excess of what is required for business. This is because user rights are frequently granted, but seldom revoked. For example, access rights are granted to users when they join the company, start a new project, change job roles, etc. But, IT staff don't know when users are done with their file access needs, and users themselves don't call up the help desk and ask for their access to be revoked once their needs change. As a consequence, access rights become excessive and, over time, no longer based on a business need-to-know. This excess access increases the risk of malicious insiders accessing sensitive file data.

Identifying excessive access is hard for most organizations to do, in part because of the questions outlined in the earlier sections: First, it's hard to find data owners, the people who can intelligently review rights. And, second, even if the owners are known, it's cumbersome to develop a baseline snapshot of rights as they exist today.

Establishing rights review cycles helps maintain a business need-to-know level of access for file data. A rights review process requires establishing a baseline rights snapshot, providing data owners with file rights to review, capturing owner feedback on required rights changes, tasking administrators with making the changes, and producing subsequent snapshots for ongoing reviews.

5. How do you know when access rights or activity violate corporate policy?

File security policy enforcement:
Hard to implement
Usually manual, not automatic
Done periodically, not continuously

Rights review cycles are important, but most organizations don't conduct reviews more than quarterly. If rights are granted to the wrong people, sensitive data is left exposed to unauthorized access until the next review cycle. Even with quarterly reviews, it can be several months before anyone is aware of the situation. And, organizations that do not perform regular rights reviews may go even longer with sensitive data exposed. The same applies to actual file access activity: continuous auditing is critical, but unless the data is analyzed for policy violations, organizations may experience a data breach and not know about it until well after the fact.

In the vast majority of organizations, policy violations are caught only during formal rights review cycles or on a completely ad-hoc basis, such as in the course of performing related security or system administration activities.

The best way to detect policy violations is to thoroughly analyze user access rights and file access activity, and apply a set of checks to determine if a violation has occurred. If this can be done in an automated way and in real-time, actionable alerts can be generated so that administrators can take action as soon as problems are detected. This means, for example, that if a group is granted access to data they should not be able to access, an alert can be triggered at the time access is granted, without the need to wait for the next rights review cycle. Or, if someone actually accesses data they shouldn't, security staff can be informed in time to mitigate the impact.

Conclusions

Insider threats and regulatory compliance mandates place increasing pressures on organizations to be accountable for their file data. Yet the volume and growth of file data make it a challenge for organizations of all sizes to answer the questions above in a timely manner. The first step for any organization is to understand where they have gaps in addressing these five key questions. Because of the limitations of current approaches and point products, most organizations will find that the best way to answer these questions is with a solution that integrates file activity monitoring, user rights management for files and real-time policy enforcement. These capabilities not only provide organizations with the ability to address the questions above, but supply the foundation required to achieve optimal file security in an efficient, repeatable way.

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