



# Understanding and Selecting a File Activity Monitoring Solution

Version 1.0  
Released: May 16, 2011

## Author's Note

The content in this report was developed independently of any sponsors. It is based on material originally posted on the [Securosis blog](#) but has been enhanced, reviewed, and professionally edited.

Special thanks to Chris Pepper for editing and content support.

## Licensed by Imperva



Imperva is the global leader in data security. With more than 1,300 direct customers and 25,000 cloud customers, our customers include leading enterprises, government organizations, and managed service providers who rely on Imperva to prevent sensitive data theft from hackers and insiders. Our

SecureSphere File Activity Monitoring solution delivers real-time file activity monitoring, user rights management and policy based controls. With these capabilities, SecureSphere ensures unstructured data is protected and transforms time-consuming, error-prone file management tasks into workable processes.

Visit Imperva [www.imperva.com](http://www.imperva.com)

## Copyright

This report is licensed under Creative Commons Attribution-Noncommercial-No Derivative Works 3.0.

<http://creativecommons.org/licenses/by-nc-nd/3.0/us/>

# Table of Contents

<b>Introduction</b>	<b>1</b>
A new approach to an old problem	1
Defining FAM	2
Market Drivers, Business Justifications, and Use Cases	2
<b>Technical Architecture</b>	<b>5</b>
Central Management Server	5
Sensors	6
Directory Integration	6
Capturing Access Controls (File Permissions)	6
<b>Core Features and Administration</b>	<b>8</b>
Entitlement (Permission/Rights) Analysis and Management	8
Secure Aggregation and Correlation	9
Activity Analysis	9
Data Owner Identification	10
<b>FAM: Policy Creation, Workflow, and Reporting</b>	<b>11</b>
Policy Creation	11
Workflow	11
Reporting	12
<b>Selection Process</b>	<b>13</b>
Define Needs	13
Formalize Requirements	14
Evaluate Products	14
Internal Testing	14
<b>Conclusion</b>	<b>16</b>
<b>Who We Are</b>	<b>17</b>
About the Authors	17
About Securosis	17

# Introduction

## A new approach to an old problem

One of the more pernicious problems in information security is allowing someone to perform something they are authorized to do, but catching when they do it in a potentially harmful way. For example, in most business environments it's important to allow users broad access to sensitive information, but this exposes us to all sorts of data loss/leakage scenarios. We want to know when a sales executive crosses the line from accessing customer information as part of their job to siphoning it for a competitor.

In recent years we have adopted tools such as Data Loss Prevention to help detect leaks of defined information, and Database Activity Monitoring to expose deep database activity and help detect unusual activity. But despite these developments, one major blind spot remains: monitoring and protecting enterprise file repositories.

Existing system and file logs rarely offer sufficient detail to truly track activity, generally don't correlate across multiple repository types, don't tie users to roles/groups, and don't support policy-based alerts, and don't provide comprehensive access rights audits. Even existing log management and Security Information and Event Management tools can't provide this level of information.

Four years ago when we initially developed the [Data Security Lifecycle](#), we referred to a technology we called File Activity Monitoring. At the time we saw it as similar to Database Activity Monitoring, in that it would give us the same insight into file usage DAM provides for database access. Although the technology didn't yet exist it seemed like a very logical next step from DLP and DAM.

Over the past two years the first FAM products have entered the market, and although market demand is nascent, numerous discussions with a variety of organizations show that interest and awareness are growing. FAM addresses a problem many organizations are now starting to tackle, and the time is right to dig into the technology and learn what it provides, how it works, and what to look for.

Imagine having a tool to detect an administrator suddenly copying the entire directory containing the latest engineering plans, or a user with rights to a file outside their business unit accessing it for the first time in 3 years. Or imagine being able to hand an auditor a list of all access, by user, to patient record files. Those are merely a few of the potential uses for FAM.

## Defining FAM

We define FAM as:

*Products that monitor and record all activity within designated file repositories at the user level, and generate alerts on policy violations.*

This leads to the key defining characteristics:

- Products are able to monitor a variety of file repositories, which include at minimum standard network file shares (SMB/CIFS). They may additionally support document management systems and other network file systems.
- Products are able to collect all activity, including file opens, transfers, saves, deletions, and additions.
- Activity can be recorded and centralized across multiple repositories with a single FAM installation (possibly comprised of multiple products, depending on network topology).
- Recorded activity is correlated to users through directory integration, and the product *should* understand file entitlements and user/group/role relationships.
- Alerts can be generated based on policy violations, such as an unusual volume of activity by user or file/directory.
- Activity reports can be generated for compliance and other needs.

You might think much of this should be possible with DLP, but unlike DLP, File Activity Monitoring doesn't require content analysis (although FAM may be part of, or integrated with, a DLP solution). FAM expands the data security arsenal by allowing us to understand how users interact with files, and identify issues even when we don't know their contents. DLP, DAM, and FAM are all highly complementary.

## Market Drivers, Business Justifications, and Use Cases

Now that we have defined File Activity Monitoring it's time to talk about why people buy it, how it's used, and why you might want it.

### Market Drivers

As we mentioned earlier, the term FAM first appeared in our Data Security Lifecycle. Although some people were tossing the general idea around, there wasn't a single product on the market. A few vendors were considering introducing something, but conversations with users showed a clear lack of market demand.

This changed dramatically over the past two years due to a combination of indirect compliance needs, headline-driven security concerns, and gaps in existing security tools. Although the FAM market is completely nascent, interest is slowly growing as organizations look for better handles on their unstructured file repositories.

We see three main market drivers:

- **As an offshoot of compliance:** Few regulations require continuous monitoring of user access to files, but quite a few require some level of access control auditing, particularly for sensitive files. As you'll see later, most FAM tools also include entitlement assessment, and they monitor and report clearly on activity. We see some organizations initially consider FAM to help generate compliance reports, and later activate additional capabilities to improve security.
- **Security concerns:** The combination of APT-style attacks against sensitive data repositories and headline-grabbing cases such as WikiLeaks are driving strong interest in gaining control over file repositories.

- **To increase visibility:** Although few FAM deployments start with a goal of providing visibility into file usage, once deployment starts, organizations often use it to gain a better understanding of how files are used within the organization, even if this isn't to meet a compliance or security need.

FAM, like its cousin Database Activity Monitoring, typically starts as a smaller project to protect a highly sensitive repository and then grows to expand coverage as it proves its value. Since it isn't generally required directly for compliance, we don't expect the market to explode, but rather to grow steadily.

## Business Justifications

If we turn around the market drivers, four key business justifications emerge for deployment of FAM:

- To meet a compliance obligation or reduce compliance costs. For example, to generate reports on who has access to sensitive information, or who accessed regulated files over a particular time period.
- To reduce the risk of major data breaches. While FAM can't protect every file in the enterprise, it provides significant protection for the major file repositories that turn a self-constrained data breach into an unmitigated disaster. You'll still lose files, but not necessarily the entire vault.
- To reduce file management costs. Even if you use document management systems, few tools provide as much insight into file usage as FAM. By tying usage, entitlements, and user/group activity to repositories and individual files FAM enables robust analysis to support other document management initiatives such as consolidation.
- To support content discovery. Surprisingly, many content discovery tools (particularly Data Loss Prevention) and manual processes struggle to identify file owners. FAM uses a combination of entitlement analysis and activity monitoring to help determine who owns each file.

## Example Use Cases

By now you should have a good idea how FAM can be used, but here are a few sample use cases:

- Company A deployed FAM to protect sensitive engineering documents from external attacks and insider abuse. They monitor the shared engineering file share and generate a security alert if more than 5 documents are accessed in less than 5 minutes; then block copying of the entire directory.
- A pharmaceutical company uses FAM to meet compliance requirements for drug studies. The tool generates a quarterly report of all access to study files and generates security alerts when IT administrators access files.
- Company C recently performed a large content discovery project to locate all regulated Personally Identifiable Information, but struggled to determine file owners. Their goal is to reduce sensitive data proliferation, but simple file permissions rarely indicate the file owner, which is needed before removing or consolidating data. With FAM they monitor discovered files to determine the most common accessors – who are often the owners.
- Company D has had problems with sales executives sucking down proprietary customer information before taking jobs with competitors. They use FAM to generate alerts based on both high-volume access and authorized users accessing older files they've never touched before.

As you can see, tying users to activity, combined with the capability to alert or block based on flexible usage policies, makes FAM interesting. Imagine being able to kick off a security investigation based on a large number of file access, or low-and-slow access by a service or administrative account.

## File Activity Monitoring vs. Data Loss Prevention

The relationship between FAM and DLP is interesting. These two technologies are extremely complementary – so much so that in one case (as of this writing) FAM is a feature of a DLP product – but they also achieve slightly different goals.

Securosis, L.L.C.

The core value of DLP is its *content analysis* capabilities: the ability to dig into a file and understand its content. FAM, on the other hand, doesn't necessarily need to know the contents of a file or repository to provide value. Certain access patterns themselves often indicate a security problem, and knowing the exact file contents may not be necessary for compliance initiatives such as access auditing.

FAM and DLP work extremely well together, but each provides plenty of value on its own.

# Technical Architecture

FAM is a relatively new technology, but we already see a few consistent architectural models emerging. The key components are a central management server, sensors, and connectors to the directory infrastructure.

## Central Management Server

The core function of FAM is to monitor user activity on file repositories. While conceptually simple, this information is only sometimes available natively from the repository, and enterprises store their sensitive documents and files using a variety of different technologies.

This leads to three main deployment options – each of which starts with a central management server or appliance:

- **Single Server/Appliance:** A single server or appliance serves as both the sensor/collection point and management console. This configuration is typically used for smaller deployments and when installing collection agents isn't possible.
- **Two-tier Architecture:** This consists of a central management server and remote collection points/sensors. The central server may or may not monitor directly — but either way it aggregates information from remote systems, manages policies, and generates alerts. The remote collectors may use any of the collection techniques we will discuss later, and always feed data back to the central server.
- **Hierarchical Architecture:** Collection points/sensors aggregate to business-level or geographically distributed management servers, which in turn report to an enterprise management server. Hierarchical deployments are best suited for large enterprises, which may have different business unit or geographic needs. They can also be configured to only pass certain kinds of data between tiers to handle large volumes of information, to support privacy by unit or geography, and to support different policy requirements.

Whichever deployment architecture you choose, the central server aggregates all collected data (except deliberately excluded data), performs policy-based alerting, and manages reporting and workflow.

The server itself may be available in one of three flavors (or for hierarchical deployments, a combination of the three):

- Dedicated appliance
- Software/server
- Virtual appliance

Which flavors are available depends on the vendor, but most offer at least one native option (appliance/software) and a virtual appliance.

If the product supports blocking, this is usually handled by configuring it as a transparent bridge or in the server agent, which we will discuss in a moment.

## Sensors

The next component is the sensors used to collect activity. Remember that this is a data-center oriented technology, so we focus on the *file repositories*, rather than the *file access points (endpoints)*. There are three primary homes for files:

- Server-based file shares (Windows and UNIX/Linux)
- Network Attached Storage (NAS)
- Document Management Systems (including SharePoint)

SANs are generally accessed through servers attached to a controller/logical unit or document management systems, so FAM systems focus on the file server/DMS and ignore the storage backend.

FAM tools use one of three options to handle all these technologies:

- **Network monitoring:** Passive monitoring of the network outside the repository, which may be performed in bridge mode or in parallel, by sniffing at a SPAN or mirror port on the local network segment. The FAM sensor or server/appliance only sniffs for relevant traffic (typically the CIFS protocol, and possibly others such as WebDAV).
- **Server agent:** This is an operating system-specific agent that monitors file access on the server (usually Windows or UNIX/Linux). The agent does the monitoring directly, and does not rely on native OS audit logs.
- **Application integration:** Certain NAS products and document management systems support native auditing well beyond what's normally provided by operating systems. In these cases, the FAM product may integrate via an agent, extension, or administrative API.

The role of the sensor is to collect activity information: who accessed the file, what they did with them (open, delete, etc.), and when. The sensor should also track important information such as permission changes.

## Directory Integration

This is technically a function of the central management server, but may involve plugins or agents to communicate with directory servers.

Directory integration is one of the most important functions of a File Activity Monitor. Without it the collected activity isn't nearly as valuable. As you'll see when we talk about the different functions of the technology, one of the most useful is its ability to manage user entitlements and scan for things such as excessive permissions.

You can assume Active Directory is supported, and LDAP is likely, but if you have an unusual directory server be sure to check with the vendor before buying any FAM products.

Roles and permissions change on a constant basis, so it's important for this data flow to be as close to real-time as possible so the FAM tool knows, at all times, the actual group/role status of users. Since directory information is often so out of date, some FAM tools can integrate with other systems, such as Human Resources platforms, for more current information.

## Capturing Access Controls (File Permissions)

Although this isn't a separate architecture component, all File Activity Monitors are able to capture and analyze existing file permissions — something else we will discuss later.

**Tip-** *the cleaner your directory, the better your ability to correlate activity and rights with actual employees. You may want to integrate with human resources systems if you need a better authoritative source.*

authoritative source.  
you need a better  
resources systems if

Securosis, L.L.C.

This is done by granting administrator or file owner permissions to the FAM server or sensor, which then captures file permissions and sends them back to the management server. Changes are then synchronized in real time through monitoring, and in some cases the FAM is used to manage future privilege changes.

That's it for the base architecture — in the next section we'll start talking about all the nifty features that run on these components.

# Core Features and Administration

Now that we understand the technical architecture, let's look at the principal features seen across most File Activity Monitoring tools.

## Entitlement (Permission/Rights) Analysis and Management

One of the most important features in most FAM products is *entitlement (permission) analysis*. The tool collects all the file and directory permissions for the repository, ties them back to users and groups via directory integration, and generates a variety of reports. Knowing that an IP address tried to access a file is interesting; but practicality really requires policies to account for users, roles, and their mappings to real-world contexts such as business units.

As we mentioned in the technical architecture section; all FAM products integrate with directory servers to gather user, group, and role information. This is the only way tools can gather sufficient context to support security requirements, such as tracing activity back to a real employee rather than just a username that might not identify the person behind it. (Not that FAM is magic – if your directories don't contain sufficient information for these mappings you might still have a lot of work to trace back identities).

At the most basic level a FAM tool uses this integration to perform at least some minimal analysis on users and groups. The most common is permission analysis – providing complete reports on which users and groups have rights to which directories/repositories/files. This is often a primary driver for buying a FAM tool in the first place, as such reports are often required for compliance.

Some tools include more advanced analysis to identify entitlement issues – particularly rights conflicts. For example, you may be able to identify which users in accounting also have engineering rights. Or list users with multiple roles that violate conflict of interest policies. While *useful* for security, these capabilities can be *crucial* for finding and fixing compliance issues.

A typical rights analysis collects existing rights, maps them to users and groups, helps identify excessive permissions, and identifies unneeded rights. Examples include:

- Determine which users outside engineering have rights to engineering documents.
- Find which users with access to healthcare records also have access to change privileges, but aren't in an administrative group.
- Identify all files and repositories the accounting group has access to, and then which other groups also have access to those files.
- Identify dormant users in the directory who still have access to files.

Finally, the tool may allow you to manage permissions internally so you don't have to manually connect to servers in order to make entitlement changes.

## Secure Aggregation and Correlation

As useful as FAM is for a single repository, its real power appears as you monitor larger swaths of your organization and can centrally manage permissions, activities, and policies.

FAM tools use a similar architecture to Database Activity Monitoring – with multiple sensors of different types sending data back to the central management server. This information is normalized, stored in a secure repository, and available for a variety of analyses and reports. The information is also analyzed in real time for policy violations and (possible) enforcement actions, as we will discuss later.

The tools don't care if one server is a NAS, another a Windows server, and the last a supported document management system – they can review them all consistently.

This aggregation also supports correlation – you can build policies based on activities occurring across different repositories and users. For example, you can alert on unusual activity by a single user across multiple file servers, or on multiple user accounts all accessing a single file in one location.

Essentially, the FAM tool gives you a big picture view of all file activity across monitored repositories, with various ways to build alerts and analyze the data, from a central management server. If your product supports multiple file protocols, it will present them in a consistent, activity-based format (e.g., open, delete, privilege change, etc.).

## Activity Analysis

While understanding permissions and collecting activity are great, and may be all you need for a compliance project, the real power of FAM is its capability to monitor all file activity (at the repository level) in real time, and generate alerts, or block activity based on security policies.

Going back to our technical architecture: activity is collected via network monitoring, software agents, or other application integration. The management server then analyzes this activity for policy warnings and violations such as:

- A user accessing a repository they have access to, but had not accessed within the past 180 days.
- A sales employee downloading more than 5 customer files in a single day.
- Any administrator account accessing files in a sensitive repository.
- A new user (or group) being given rights to a sensitive directory.
- Any user account copying an entire directory from an engineering server.
- A service account accessing files.

Some tools allow you to define policies based on a sensitivity tag for the repository and user groups (or business units), instead of having to manually build policies on a per-repository or per-directory level.

This analysis doesn't necessarily need to happen in real time – it can also be done on a scheduled or *ad hoc* basis to support a specific requirement, such as an auditor who wants to know who accessed a file, or as part of an incident investigation. We'll talk more about reporting later.

## Data Owner Identification

Although every file has an 'owner', translating that to an actual person is often a herculean task. Another primary driver of File Activity Monitoring is to help organizations identify file owners.

This is typically managed through a combination of privilege and activity analysis. Privileges might reveal a file owner, but activity may be more useful. You could build a report showing the users who most often access a file, then correlate it against ownership permissions for a short list of likely file owners.

This is, of course, much simpler if the tool was already monitoring the repository and can identify who created the file initially.

# FAM: Policy Creation, Workflow, and Reporting

Now that we have covered the base features it's time to consider how these tie in with policies, workflow, and reporting. We'll focus on the features needed to support these processes rather than defining the processes themselves.

## Policy Creation

File Activity Monitoring products offer two major categories of policies:

- Entitlement (Permissions/Access Control) policies. These define which users can access which repositories and types of data. They define rules for things such as orphaned user accounts, separation of duties, role/group conflicts, and other situations that don't require real-time file activity.
- Activity-based policies. These alert and block based on real-time user activity.

When evaluating products, look for a few key features to help with policy creation and management:

- Policy templates to serve as examples and baselines for building your own policies.
- A clean user interface that allows you to understand business context. For example it should allow you to group categories, pool users and groups to speed up policy application (e.g., combine all the different accounting related groups into "Accounting"), and group and label repositories. This is especially important given the volume of entries to manage when you integrate with large user directories and multi-terabyte repositories.
- New policy wizards to speed up policy creation.
- Hierarchical management for multiple FAMs in the same organization.
- Role-based administration, including roles for super administrators and assigning policies to sub-administrators.
- Policy backup and restore.

## Workflow

As with policy creation, workflow requirements focus on the two major functions of FAM: entitlement management and activity monitoring.

### Entitlement Management

This workflow should support a closed-loop process for collection of privileges, analysis, and application of policy-based changes. Your tool should do more than merely collect access rights – it should help you build a process to ensure that access controls match policies. This typically requires a combination of different workflows for different goals – including identification of orphan accounts with access to sensitive data, excessive privileges, conflict of interest/separation of duties based on user groups, and restricting access to sensitive repositories.

Products and policies vary, but they tend to fit a common pattern:

- Collect existing entitlements.
- Analyze based on policies.
- Apply corrective actions (either building an alerting/blocking policy or changing privileges).
- Generate a report on identified and remediated issues.

The workflow should also link into data owner identification because this must often be understood before changing rights.

## Activity Monitoring and Protection

The activity monitoring workflow is very different than entitlement management. Here the focus is on handling alerts and incidents in real time. The key interface is the *incident handling queue* common to most security tools. The queue lists incidents and supports various sorting and filtering options. Its workflow tends to follow the following structure:

1. Incident occurs and alert appears in the queue. It is displayed with the user, policy violated, and repository or file involved.
2. The incident handler can investigate further by filtering for other activity involving that user, that repository, or that policy over a particular time period (or various combinations).
3. The handler can assign or escalate the incident to someone else, close the incident, or take corrective actions such as adjusting file permissions.

The key to keeping this efficient is not requiring the incident handler to jump around the user interface in a manual process. For example, clicking on an incident should show its details and then links to see other related incidents by user, policy, and repository.

Incidents should also be grouped logically – an attempt to copy an entire directory should appear as one incident, not one incident for each of 1,000 files in the repository.

Any FAM product may also include additional workflows, such as for identifying file owners.

## Reporting

One of the most important functions of any File Activity Monitoring product is robust reporting – this is particularly important for meeting compliance requirements.

Aside from a repository of predefined reports for common requirements such as PCI and HIPAA, the tool should allow you to generate arbitrary reports. (We hate to list that as a requirement, but we still occasionally see security tools that don't support creation of arbitrary reports).

# Selection Process

## Define Needs

The first step in the process is to determine your needs, keeping in mind that there are two main drivers for File Activity Monitoring projects, and it's important to understand the differences and priorities between them:

- Entitlement management
- Activity monitoring

Most use cases for FAM fall into one of these two categories, such as data owner identification. It's easy to say "Our goal is to audit all user access to files," but we recommend you get more specific. Why are you monitoring? Is your primary need security or compliance? Are there specific business unit requirements? These answers help pick the best solution for individual requirements.

We recommend the following process for this step:

1. **Create a selection committee:** File Activity Monitoring initiatives tend to directly involve three major technical stakeholders, as well as compliance/legal. On the IT side we typically see security and server and/or storage management involved. This varies considerably, based on organization size and storage complexity. For example, it might be the document management system administrators, SharePoint administrators, NAS/storage management, and server administration. The key is to involve the major administrative leads for your storage repositories. You may also need to involve network operations if you plan to use network monitoring.
2. **Define the systems and platforms to protect:** FAM projects are typically driven by a clear audit or security goal tied to particular storage repositories. In this stage, detail the scope of what will be protected and the technical specifics of the platforms involved. You'll use this list to determine technical requirements and prioritize features and platform support later. Remember that needs grow over time, so break the list into a group of high-priority systems with immediate requirements, and a second group summarizing all major platforms you may need to protect later.
3. **Determine protection and compliance requirements:** For some repositories you might want strict preventative security controls, while others may just need comprehensive activity monitoring or entitlement management to satisfy a compliance requirement. In this step, map your protection and compliance needs to the platforms and repositories from the previous step. This will help you determine everything from technical requirements to process workflow.
4. **Outline process workflow and reporting requirements:** File Activity Monitoring workflow varies by use. You will want to define different workflows for entitlement management and activity monitoring, as they may involve different people — that way you can define what you need instead of having the tool determine your process. In most cases audit, legal, or compliance, have at least some sort of reporting role. Different FAM tools have different strengths and weaknesses in management interfaces, reporting, and internal workflow — so think through the process before defining technical requirements to prevent headaches down the road.

By the end of this phase you should have defined key stakeholders, convened a selection team, prioritized the systems to protect, determined protection requirements, and roughed out process workflow.

## Formalize Requirements

This phase can be performed by a smaller team working under the mandate of the selection committee. Here the generic needs from phase 1 are translated into specific technical features, and any additional requirements are considered. This is the time to come up with criteria for directory integration, repository platform support, data storage, hierarchical deployments, change management integration, and so on. You can always refine these requirements after you begin the selection process and get a better feel for how the products work.

At the conclusion of this stage you will have a formal RFI (Request For Information) for vendors, and a rough RFP (Request For Proposals) to clean up and formally issue in the evaluation phase.

## Evaluate Products

As with any product, it can be difficult to cut through the marketing materials and figure out whether a product really meets your needs. The following steps should minimize your risk and help you feel confident in your final decision:

1. **Issue the RFI:** Larger organizations should issue an RFI through established channels and contact the leading FAM vendors directly. If you're a smaller organization start by sending your RFI to a trusted VAR and email the FAM vendors which appear appropriate for your organization.
2. **Perform a paper evaluation:** Before bringing anyone in, match any materials from the vendor or other sources to your RFI and draft RFP. Currently few vendors are in the FAM market so your choices will be limited, but you should be fully prepared before you go into any sales situations. Also use outside research sources and product comparisons.
3. **Bring in vendors for on-site presentations and demonstrations:** Instead of a generic demonstration, ask each vendor to walk through your specific use cases. Don't expect a full response to your draft RFP – these meetings are to help you understand the different options and eventually your requirements.
4. **Finalize your RFP and issue it to your short list of vendors:** At this point you should completely understand your specific requirements and issue a formal, final RFP.
5. **Assess RFP responses and begin product testing:** Review the RFP results and drop anyone who doesn't meet any of your hard requirements (such as platform support), as opposed to 'nice-to-have' features. Then bring in any remaining products for in-house testing. You will want to replicate your highest volume system and its traffic if at all possible. Build a few basic policies that match your use cases, and then violate them, so you get a feel for policy creation and workflow.
6. **Select, negotiate, and buy:** Finish testing, take the results to the full selection committee, and begin negotiating with your top choice.

## Internal Testing

In-house testing is the last chance to find problems in your selection process. Make sure you test the products as thoroughly as possible. And keep in mind that smaller organizations may not have the resources or even the opportunity to test before purchasing. Key aspects to test are:

- **Platform support and installation:** Determine agent or integration compatibility (if needed) with your repositories. If you plan to use agents or integrate with a document management system, this is one of the most important steps.
- **Performance:** Is network or agent performance acceptable for your environment? Are there other operational considerations driving you toward one model or another? Don't set arbitrary standards – monitor performance on

production systems to ensure your tests represent operational requirements. This is most important if you have massive repositories with high-volume access.

- **Policy creation and management:** Create policies to understand the process and its complexity. Will built-in policies satisfy your requirements? Are there wizards and less-technical options for non-technical experts to create policies? Then violate policies and try to evade or overwhelm the tool to learn its limitations.
- **Incident workflow:** Review the working interface with those employees who will be responsible for enforcement.
- **Directory integration**
- **Entitlement workflow:** Is it a closed loop or manual?

# Conclusion

The main question I have asked myself about File Activity Monitoring is why did it take so long? From a technology perspective it's a relatively straightforward problem, albeit one with plenty of nuance. In terms of need it seems that understanding and monitoring how users interact with files should have been high on the list for both compliance and security. It certainly wasn't because no one thought about it — I know of at least two vendors which started testing out the idea years before the first products hit the market — both informed me that they found little interest among clients and prospects.

One big reason is likely priorities — while we always knew protecting internal files was important, everything from attacks, to regulations such as SOX and PCI, to protecting ongoing infrastructure upgrades, has been so demanding we couldn't really spare the time for anything else. Sure, documents were a 'priority', but auditors are much noisier.

Another reason may be the sheer intimidation factor of dealing with internal file stores. Many security professionals I've spoken with cringe at the thought of large scale data identification and classification projects — even supported by tools such as DLP.

But the advantage of File Activity Monitoring is that we can start protecting our content without having to tackle these massive problems. With FAM we can quickly assess information usage and focus on those points with the most business activity, which are likely where the most important information resides. We can build rules around usage patterns to help identify problems before they get too bad. Everything from sales folks snarfing down volumes of data before moving on to the next job, to brand new accounts created on some server that suddenly accesses a file store with administrative privileges.

FAM is especially interesting for two reasons. First, the entitlement management functionality is incredibly valuable for any organization that needs to gain a handle on which users have access to which repositories — especially for compliance. Second is the intersection of transparency, aggregation, correlation, and policy-based alerting to improve security. We can now collect, analyze, and alert on (or block) user activity without overly interfering with normal business processes.

FAM simply provides better visibility, more efficiently, into file-based content than we have ever had before.

# Who We Are

## About the Authors

### **Rich Mogull, Analyst/CEO**

Rich has twenty years experience in information security, physical security, and risk management. He specializes in data security, application security, emerging security technologies, and security management. Prior to founding Securosis, Rich was a Research Vice President at Gartner on the security team, where he also served as research co-chair for the Gartner Security Summit. Prior to his seven years at Gartner, Rich worked as an independent consultant, web application developer, software development manager at the University of Colorado, and systems and network administrator. Rich is the Security Editor of *TidBITS*, a monthly columnist for *Dark Reading*, and a frequent contributor to publications ranging from Information Security Magazine to *Macworld*. He is a frequent industry speaker at events including the RSA Security Conference and DefCon, and has spoken on every continent except Antarctica (where he's happy to speak for free — assuming travel is covered).

## About Securosis

Securosis, L.L.C. is an independent research and analysis firm dedicated to thought leadership, objectivity, and transparency. Our analysts have all held executive level positions and are dedicated to providing high-value, pragmatic advisory services.

We provide services in four main areas:

- Publishing and speaking: Including independent objective white papers, webcasts, and in-person presentations.
- Strategic consulting for end users: Including product selection assistance, technology and architecture strategy, education, security management evaluations, and risk assessments.
- Strategic consulting for vendors: Including market and product analysis and strategy, technology guidance, product evaluations, and merger and acquisition assessments.
- Investor consulting: Technical due diligence including product and market evaluations, available in conjunction with deep product assessments with our research partners.

Our clients range from stealth startups to some of the best known technology vendors and end users. Clients include large financial institutions, institutional investors, mid-sized enterprises, and major security vendors.

Securosis has partnered with security testing labs to provide unique product evaluations that combine in-depth technical analysis with high-level product, architecture, and market analysis.