



Low Hanging Fruit: Quick Wins with Data Loss Prevention

Version 2.0
Released: March 31, 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 McAfee



McAfee, Inc., headquartered in Santa Clara, California, is the world's largest dedicated security technology company. McAfee is relentlessly committed to tackling the world's toughest security challenges. The company delivers proactive and proven solutions and services that help secure systems and networks around the world, allowing users to safely connect to the Internet, browse and shop the web more securely. Backed by an award-winning research team, McAfee creates innovative products that empower home users, businesses, the public sector and service providers by enabling them to prove compliance with regulations, protect data, prevent disruptions, identify vulnerabilities, and continuously monitor and improve their security.

<http://www.mcafee.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/>


Introduction

There is no shortage of reasons to deploy Data Loss Prevention (DLP). Since its emergence nearly ten years ago, we've seen it grow from a technology mostly used to find accidental emails with credit card numbers, to becoming a fundamental adjunct for data security and compliance. When we released the first version of this report last year DLP was still mostly limited to early adopters, but a combination of changes in the threat environment, security priorities, and compliance needs are driving DLP deep into the mainstream of the security market.

Results from a security survey of over 1000 IT professionals performed in July and August of 2010 showed that Data Loss Prevention was widely deployed in organizations of all sizes, appearing in nearly half of all organizations in some form (network, endpoint, or storage). Even if we discount some results due to survey response bias, the results clearly show extensive adoption.

Organizations adopt DLP for a wide range of reasons. For some it's a direct response to a security incident; either one of their own, one at a peer, or a high-profile breach making the rounds in the mainstream media.

Others use it to support compliance, even though it's rarely a line item in any regulation. DLP helps reduce the manual costs of finding and securing sensitive information, reduces compliance-related incidents, and may even constrain your compliance scope by narrowing where regulated information is stored and used. We expect to see even greater use of DLP in the coming years to locate credit card numbers due to changes in the PCI standard.



Finally, we often see DLP used as a combination of an education and a visibility tool. Merely showing users when they accidentally or unknowingly place information at risk can dramatically reduce an organization's exposure. DLP is also the first *information security* tool that actually *understands the information*; providing never before possible insights as to how we use our content.

But deploying any new technology can be intimidating; especially if you're unfamiliar with it and haven't used it hands-on before. That's why we developed this *Quick Wins Process* to help you derive immediate value from your tool while still setting yourself up for long-term success.

The Quick Wins Process

Two of the most common criticisms of Data Loss Prevention (DLP) that comes up in user discussions are a) its complexity and b) the fear of false positives. Security professionals worry that DLP is an expensive widget that will fail to deliver the expected value -- turning into yet another black hole of productivity. But when used properly DLP provides rapid assessment and identification of data security issues not available with any other technology.

We don't mean to play down the real complexities you might encounter as you roll out a complete data protection program. Business use of information is itself complicated, and no tool designed to protect that data can simplify or mask the underlying business processes. However, there are steps you can take to obtain significant immediate value and security gains without blowing your productivity or wasting important resources.

In this paper we highlight the lowest hanging fruit for DLP, refined in conversations with hundreds of DLP users. These aren't meant to incorporate the entire DLP process, but to show you how to get real and immediate wins before you move on to more complex policies and use cases.

Establish Your Process


Nearly every DLP customer we have talked with has discovered actionable offenses committed by employees as soon as they turn the tool on. Some of these require little more than contacting a business unit to change a bad process, but quite a few result in security guards escorting people out of the building, or even legal action. In one case, a DLP vendor plugged in the tool for a lunchtime demonstration on the same day a senior executive decided to send proprietary information to a competitor. Needless to say, the vendor lost their hard drives that day to support the investigation and subsequent legal action, but they didn't seem too unhappy when they later acquired a new customer.

Even if you aren't planning on moving straight to enforcement mode, you need to put a process in place to manage the issues that will crop up once you activate your tool. The kinds of issues you need to figure out how to address in advance fall into two categories:

- **Business Process Failures:** Although you'll likely manage most business process issues as you roll out your sustained deployment, the odds are high some will be of such high concern they will require immediate remediation. These are often compliance related.
- **Egregious Employee Violations:** Most employee-related issues can be dealt with as you gradually shift into enforcement mode, but as in the example above, you will encounter situations requiring immediate action.

In terms of process, we recommend two tracks based on the nature of the incident. Business process failures usually involve escalation within security or IT, possible involvement of compliance or risk management, and engagement with the business unit itself. You are less concerned with getting someone in trouble than stopping the problem.

Employee violations, due to their legal sensitivity, require a more formal process. Typically you'll need to open an investigation and immediately escalate to management while engaging legal and human resources (since this might be a firing offense). Contingencies need to be established in case law enforcement is engaged, including plans to provide forensic evidence to law enforcement without having them walk out the door with your nice new DLP box and hard drives. Essentially you want to implement whatever process you already have in place for internal employee investigations and potential termination.



Prepare Your Directory Servers

One of the single most consistent problems with DLP deployments has nothing to do with DLP, and everything to do with the supporting directory (AD, LDAP, or whatever) infrastructure. Since with DLP we are concerned with user actions across networks, files, and systems (and on the network with multiple protocols), it's important to know exactly who is committing all these violations. With a file or email it's usually a straightforward process to identify the user based on their mail or network logon ID, but once you start monitoring anything else, such as web traffic, you need to correlate the user's network (IP) address back to their name.

This is built into nearly every DLP tool, so they can track what network addresses are assigned to users when they log onto the network or a service.

The more difficult problem tends to be the business process; correlating these technical IDs back to real human beings. Many organizations fail to keep their directory servers current, and as a result it can be hard to find the physical body behind a login. It gets even harder if you need to figure out their business unit, manager, and so on.

For a quick win, we suggest you focus predominantly on making sure you can track most users back to their real-world identities. Ideally your directory will also include role information so you can filter DLP policies violations based on business unit. Someone in HR or Legal usually has authorization for different sensitive information than people in IT and Customer Service, and if you have to manually figure all this out when a violation occurs, it will really hurt your efficiency later.

Integrate with Your Infrastructure

The last bit of preparation is to integrate with the important parts of your infrastructure. How you do this will vary depending on your initial focus (endpoint, network, or discovery). Remember, this all comes after you integrate with your directory servers.

The easiest deployments are typically on the network side, since you can run in monitoring mode without having to do too much integration. This might not be your top priority, but adding what's essentially an out of band network sniffer is very straightforward. Most organizations connect their DLP monitor to their network gateway using a SPAN or mirror port. If you have multiple locations, you'll probably need multiple DLP boxes and have to integrate them using the built-in multi-system management features common to most DLP tools.

Most organizations also integrate a bit more directly with email, since it is particularly effective without being especially difficult. The store-and-forward nature of email, compared to other real-time protocols, makes many types of analysis and blocking easier. Many DLP tools include an embedded mail server (MTA, or Mail Transport Agent) which you can simply add as another hop in the email chain, just like you probably deployed your spam filter.

Endpoint rollouts are a little tougher because you must deploy an agent onto every monitored system. The best way to do this (after testing) is to use whatever software deployment tool you currently use to push out updates and new software.

Content discovery -- scanning data at rest in storage -- can be a bit tougher, depending on how many servers you need to scan and who manages them. For quick wins, look for centralized storage where you can start scanning remotely through a file share, as opposed to widely distributed systems where you have to manually obtain access or install an agent. This reduces the political overhead and you only need an authorized user account for the file share to start the process.

You'll notice we haven't talked about all the possible DLP integration points, but instead focused on the main ones to get you up and running as quickly as possible. To recap:

- For all deployments: Directory services (usually your Active Directory and DHCP servers).
- For network deployments: Network gateways and mail servers.
- For endpoint deployments: Software distribution tools.
- For discovery/storage deployments: File shares on the key storage repositories (you generally only need a username/password pair to connect).

Deploying and Using Your DLP Solution

The differences between a long-term DLP deployment and our "Quick Wins" approach are goals and scope. With a traditional deployment we focus on comprehensive monitoring and protection of very specific data types. We know what we want to protect (at a granular level) and how we want to protect it, and we can focus on comprehensive policies with low false positives and a robust workflow. Every policy violation is reviewed to determine if it's an incident that requires a response.

In the Quick Wins approach we are concerned less about incident management, and more about gaining a rapid understanding of how information is used within our organization. There are two flavors to this approach -- one where we focus on a narrow data type, typically as an early step in a full enforcement process or to support a compliance need, and the other where we cast a wide net to help us understand general data usage to prioritize our efforts. Long-term deployments and Quick Wins are not mutually exclusive -- each targets a different goal and both can run concurrently or sequentially, depending on your resources.

Remember: even though we aren't talking about a full enforcement process, it is absolutely essential that your incident management workflow be ready to go when you encounter violations that demand immediate action!

Choose Your Flavor

The first step is to decide which of two general approaches to take:

- Single Type: In some organizations the primary driver behind the DLP deployment is protection of a single data type, often due to compliance requirements. This approach focuses only on that data type.
- Information Usage: This approach casts a wide net to help characterize how the organization uses information, and identify patterns of both legitimate use and abuse. This information is often very useful for prioritizing and informing additional data security efforts.

Choose Your Deployment Type

Depending on your DLP tool, it will be capable of monitoring and protecting information on the network, on endpoints, or in storage repositories -- or some combination of these. This gives us three pure deployment options and four possible combinations.

- **Network Focused:** Deploying DLP on the network in monitoring mode provides the broadest coverage with the least effort. Network monitoring is typically the fastest to get up and running due to lighter integration requirements. You can often plug in a server or appliance over a few hours or less, and instantly start evaluating results.
- **Endpoint Focused:** Starting with endpoints should give you a good idea of which employees are storing data locally or transferring it to portable storage. Some endpoint tools can also monitor network activity on the endpoint, but these capabilities vary widely. In terms of Quick Wins, endpoint deployments are generally focused on analyzing stored content on the endpoints.
- **Storage Focused:** Content discovery is the analysis of data at rest in storage repositories. Since it often requires considerable integration (at minimum, knowing the username and password to access a file share), these deployments, like endpoints, involve more effort. That said, its scan of major repositories is very useful, and in some organizations it's as important (or even more so) to understand stored data than to monitor information moving across the network.

Network-focused deployments tend to be the fastest to get up and running since they require the least amount of integration. But just because it's easier doesn't mean it is the best option in all circumstances- there are many valid cases to start with endpoint or storage. You might only have an endpoint tool, or need to focus on storage for a compliance audit.

Network deployments typically provide the most immediate information with the lowest effort, but depending on what tools you have available and your organization's priorities, it may make sense to start with endpoints or storage. Combinations are obviously possible, but we suggest you roll out multiple deployment types sequentially rather than in parallel to manage project scope.

Define Your Policies

The last step before hitting the "on" switch is to configure your policies to match your deployment flavor.

In a single type deployment, either choose an existing category that matches the data type in your tool, or quickly build your own policy. In our experience, pre-built categories common in most DLP tools are almost always available for the data types that commonly drive a DLP project. Don't worry about tuning the policy -- right now we just want to toss it out there and get as many results as possible. Yes, this is the exact opposite of our recommendations for a traditional, focused DLP deployment.

In an information usage deployment, turn on all the policies or enable promiscuous monitoring mode. Most DLP tools only record activity when there are policy violations, which is why you must enable the policies. A few tools can monitor general activity without relying on a policy trigger (either full content or metadata only). In both cases our goal is to collect as much information as possible to identify usage patterns and potential issues.

Monitor

Now it's time to turn on your tool and start collecting results.

Don't be shocked -- in both deployment types you will see a lot more information than in a focused deployment, including more potential false positives. Remember, you aren't concerned with managing every single incident, but want a broad understanding of what's going on on your network, in endpoints, or in storage.

Analyze

Now we get to the most important part of the process -- turning all that data into useful information.

Once we collect enough data, it's time to start the analysis process. Our goal is to identify broad patterns and identify any major issues. Here are some examples of what to look for:

- A business unit sending out sensitive data unprotected as part of a regularly scheduled job.
- Which data types broadly trigger the most violations.
- The volume of usage of certain content or files, which may help identify valuable assets that don't cleanly match a pre-defined policy.
- Particular users or business units with higher numbers of violations or unusual usage patterns.
- False positive patterns, for tuning long-term policies later.
- All DLP tools provide some level of reporting and analysis, but ideally your tool will allow you to set flexible criteria to support the analysis.

What Did We Achieve?

If you followed this process, by now you've created a base for your ongoing DLP usage while achieving valuable short-term goals. In a short amount of time you have:

- Established a flexible incident management process.
- Integrated with major infrastructure components.
- Assessed broad information usage.
- Set a foundation for later focused efforts and policy tuning to support long-term management.

Thus by following the Quick Wins process you can show immediate results while establishing the foundations of your program, all without overwhelming yourself by forcing unprepared action on all possible alerts before you understand information usage patterns.

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.