TechTarget
Security Media

SearchSecurity.com SearchSecurity.co.UK SearchSecurity.in SearchFinancialSecurity.com SearchMidmarketSecurity.com
SearchSecurityChannel.com INFORMATION SECURITY DECISIONS FINANCIAL INFORMATION SECURITY DECISIONS INFORMATION SECURITY

SearchSecurityChannel.com Pocket E-Guide

Real World Data Loss Prevention (DLP) Benefits

In this expert E-Guide, Rich Mogull reviews data breach prevention techniques and policies to help ensure that your organization doesn't make headlines for the wrong reasons. Mogull also presents dozens of real world Data Loss Prevention (DLP) usage scenarios to display how DLP works. In the second article, he shares a collection of lessons learned to help you avoid common pitfalls while deriving maximum value from DLP solutions.

TechTarget
The Technology Media ROI Experts

SearchSecurityChannel.com Pocket E-Guide

Real World Data Loss Prevention (DLP) Benefits

Table of Contents:

Data breach prevention techniques: Helping customers avoid data breaches

by Allen Zuk, Contributor

No one wants to read about their organization—or that of their customers—in the headlines following a breach of customer data or other sensitive information. And now that the Privacy Rights Clearing House maintains a comprehensive list of all known data breaches since 2005, major breaches live on in infamy long after the incident. Even more embarrassing is that most breaches are preventable.

In this tip, we'll review data breach prevention techniques and policies that can help ensure your customers don't make headlines for the wrong reasons.

Data breach prevention techniques

There are numerous techniques and a variety of tools that can help stop leakage or loss of information. In the following sections, we will briefly discuss each method and what solution providers can do to help their clients implement stop gaps to improve their overall information security posture.

Information security policies

Instituting information security policies and procedures is the least expensive way to help combat data loss. Policies and procedures are developed to instill a common set of principles for all personnel. That being said, policies and guidelines are also infrequently enforced. If staff members are not educated on these policies and guidelines, then enforcement becomes almost impossible.

To start, help your customers by either assisting them with conducting an information security policy assessment or by offering them the service. Solution providers will need to be well-versed in the use of information security baseline standards, such as ISO 27002 (formerly ISO 17799) and COBIT. Having a thorough understanding of these guidelines will help you (the solution provider) position yourself as a trusted advisor to the client.

Solution providers should have solid policy-writing skills and knowledge of the various data breach laws as well as those that are being drafted. Solution providers also need to be aware of the various State Security Breach Notification Laws that are in existence and to be able to articulate and integrate these with their clients' information security policies. Other best practices include making sure customers update their antivirus software, and maintaining an exit policy for employees that ensures privileges are revoked.

Emerging technologies

There are a variety of products available in the DLP (data loss prevention) category that combine software management and policy implementation and control. These products provide an "automated" mechanism that responds to defined attributes for policy management. In the simplest terms, these products allow the administrator to define criteria that determines how information will flow in and out of an organization as well as provide an audit trail and an alert notification process for exclusionary requirements.

DLP vendors that offer such products include NextLabs Inc., Orchestria Corp., Proofpoint Inc., Vericept Corp., Verdasys Inc. and Symantec Corp. (via its acquisition of Vontu). Their technologies utilize customizable "policies" defined by the organization to monitor, report against, redirect and stop data flow within the organization's network and computing systems. When enabled, these products could, for instance, disable USB ports or prevent laptops from accessing the network.

Solution providers should be well versed in the use and application of these tools to assist their clients with policy development and implementation. Organizations often face challenges implementing and managing their data loss prevention programs, and solution providers should be prepared to fill those gaps.

Information technology risk management assessments

An information technology risk management assessment can be used to assess a company's information security posture. An information technology risk management assessment gauges the effectiveness of IT security controls and ensures the implemented security technologies do not introduce unnecessary risk and exposure to the business.

The risk management assessment includes two core program components: the first is an organization's current maturity posture snapshot for security and management of the technologies implemented, second is a detailed gap analysis report that includes a mitigation roadmap containing recommendations for continuous improvement.

The Information risk management maturity matrix diagram illustrates a sample maturity matrix that is used to evaluate the organizations current and desired posture for IT/IS security and management.

Partner with third-party vendors that specialize in conducting risk management assessments. Leverage your relationship with the customer and introduce your extended advisory support by offering strategic assessments of the customer's IT risk posture. Demonstrate the value these assessments have with simple, yet,effective "heat maps." Heat maps are high-impact illustrations that pinpoint specific gaps or deficiencies visually so the client sees where they need to focus resources immediately. This heat map illustrates a sample "heat map" highlighting areas of severe deficiency (red), minimal deficiency (yellow) and no deficiency (green).

Conclusion

While it is nearly impossible to completely stop all data loss and data leakage, there are a variety of options to mitigate the risk and exposure. However, this is not to say that solution providers should just simply throw an assortment of tools, policies and approaches at the problems.

The best value a solution provider can bring to the customer is to understand the organization, its challenges and obstacles, and develop a strategy that integrates fundamental policies for awareness and education with technologies aimed at preventing the unauthorized removal of corporate information assets, and a comprehensive IT risk management assessment to reduce the risk of breaches and exposure.

SOPHOS

INVINCIBLE

WE PROTECT YOUR DEALS, EVEN IN THE MOST HOSTILE WATERS

Deal Protection | Escalating Discounts | Margin Protection | Incumbent Renewals | sophos.com/partners

Data loss prevention benefits in the real world

by Rich Mogull

The reality of any new technology, security or otherwise, rarely lives up to its promise. Once you move past the bright sheen of the product brochures and top-level user interfaces, only the practicalities of implementing the product in the real world remain. This is especially true of newer technologies we have little prior experience with, where our product expectations are defined by marketing, the press, and the rare peer reference. It's only after these tools are tested in the real world, under full production conditions, that we really start learning how to either best implement them, or kick them back to the vendor for a little more polish (and a compelling business use).

Data loss prevention (DLP) is one of the most promising, and least understood, security technologies to emerge during the last few years. It dangles promises of ubiquitous content protection before our eyes, with shadows of complexity and costs glooming over its shoulder. As with everything, the reality is somewhere in-between. We've interviewed dozens of DLP users (including our own contacts, random volunteers and vendor references) to find out how DLP works in the trenches of the real world. The result is a collection of lessons learned and use cases to help you avoid common pitfalls while deriving maximum value.

Lesson 1: Users are confused by a confusing market

One of the more significant findings when researching this article was discovering extensive confusion as to just what comprised a DLP solution. In large part this is due to competing, and contradictory, messages from the vendor community. Data loss prevention is a generic term, and it's been used to brand everything from full DLP suites, to encryption, to USB port blocking. By our informal estimate, only 40 percent of the DLP users we talked to use a full DLP product. Of the rest, USB, file and drive encryption, and email filtering were cited as the most common data protection techniques. Many of those users knew they weren't really doing data loss prevention, but they cited cost and complexity as their concerns with using a full DLP product (which protects information on the network, in stored data, and on endpoints using deep content analysis).

One large airline we spoke with is using a generic network sniffer/forensics tool with some basic keyword policies instead of a DLP solution. But this approach has severe flaws, with the security manager saying, "I'm not sure we can actually see everything going on." They are also looking to add USB port blocking, but more to protect against malicious software than to limit data loss. They do expect to look at DLP in 2010 or 2011.

Even though there are only around a dozen full-suite DLP solutions on the market, nearly every major (and many minor) security vendor claims some sort of DLP capability. We call those tools that offer some sort of content awareness—such as regular expressions—on single channel, such as email, "DLP as a feature." But many tools claiming DLP don't even offer that basic functionality. When standard encryption tools market themselves as DLP, it's no wonder customers are confused.

Lesson 2: Full DLP solutions take more effort to deploy, but are more effective and easier to manage

Although you can whack a nail with a big enough wrench, it won't ever work as well as a hammer, and can't touch the efficiency of a nail gun. DLP as a feature does have its place--particularly for clients on a budget, or with only basic data protection needs, but our interviews consistently showed higher satisfaction among those using dedicated DLP suites. Many of the clients using DLP features described it as a temporary measure until they were ready to consider full DLP. One user stated, "We are watching the marketplace closely, but don't want to be an early adopter."

The tradeoff is that dedicated DLP does take more effort to deploy than merely flipping on a feature switch in another product, but deployment requirements are fairly low. On average, a 5,000-person organization can deploy network monitoring with email filtering in a few hours or days, using one to three Full Time Equivalents. Additional network blocking (Web and FTP) usually requires integration with an existing Web gateway, and deployment complexity scales almost linearly based on the number of egress points. Content discovery (data-at-rest scanning) is more resource intensive since you need to manually add storage repositories to scan. Each repository may only take a few seconds to minutes to add, but you have to first identify them and obtain administrative credentials. Endpoint monitoring takes time to test the software on your standard images, then deploys exactly like any other endpoint tool.

Egress Filtering Made Easy

Data loss prevention technology is designed to mitigate the threat posed by data exfiltration on a network. Follow these three steps to lessen your risk:

ENSURE your DLP has access to any outbound connections that might originate from your transaction processing network, especially dedicated pipes that are not monitored by anything on the standard enterprise gateway

DON'T RESTRICT your DLP tool to only certain types of network traffic or only protocols running on standard ports. Attackers will use different combinations to move stolen data off your network.

COMBINE your DLP with a network proxy. This is crucial to properly manage egress filtering, enabling DLP to block as much as possible.

SET your DLP to alert when it detects encrypted files; this forces attackers to use non-standard encryption.

Read the full text of this tip on SearchFinancialSecurity.com

—Rich Mogull

However, full DLP solutions include much more efficient workflow for managing policy violations, especially if compliance, human resources, or legal will be involved. They also allow users to create a single data protection policy, and then apply it across multiple channels, rather than defining the information in multiple tools.

Lesson 3: Set the right expectations and workflow early

While deploying the technology is fairly straightforward, many organizations find they struggle more with setting the right expectations, defining workflow and building policies. We once had a client install one vendor's product and start monitoring using default policies without defining any incident management procedures or workflow. They stated: "We don't want to snoop on employees, so we don't worry about involving management or human resources," without realizing that most data leaks come from employees, with likely legal and HR implications.

A typical mistake is failing to define what types of data you want to protect, and how you want to protect it, before buying a tool (then being disappointed in the result). The other major pre-selection mistake is failing to engage business unit managers outside of security. One reference purchased an endpoint-only tool to prevent information leaks onto USB devices, only to find the tool shelved once sales management started receiving complaints.

When expectations are set properly, and the tool and policies deployed in a phased manner, DLP projects tend to go smoothly with minimal overhead. On average, a 10,000 employee organization with a handful of policies only requires 1-3 FTEs, usually split part-time under multiple employees, to manage policy violations. When basic policies are deployed, such as credit card protection, that same team may handle organizations up to 50,000 or more employees. On the other end, using poorly tuned or low threshold policies will require more incident managers, and one risk-averse organization purposely chose higher false positives for greater data visibility.

Lesson 4: Poor identity management hinders good DLP

One of the largest obstacles to a successful DLP deployment is poor identity management, especially in content discovery deployments. If you locate a file with sensitive data in an unapproved location, a poor directory infrastructure may make it nearly impossible to identify the file's owner. Without being able to identify the owner, protecting the file could break a legitimate business process. One health care organization reported that it might take it days to track down a single file's owner. Even though the DLP solution scans their infrastructure relatively quickly, the manual delays of tracking down users and managers (due to a haphazard Active Directory deployment) has dragged out their project by many months.

In another case, we received a report of an organization that almost fired the wrong employee when the IP address was tied back to the wrong user, due to a contractor improperly connecting their system.

Lesson 5: False positives are a manageable concern

The single most common worry over deploying DLP is the time required to manage false positives. The assumption is that policies based on keywords or generic 16-digit (credit card) numbers will constantly trigger false positives. But in real world deployments, users find false positives to be minimal as long as the right content analysis technique is used and policies are properly tuned.

For structured data, such as credit card and account numbers, most DLP tools have a range of advanced techniques to limit false positives. You can choose to only protect numbers from an internal database (instead of a generic expression), or set thresholds that alert only for multiple violations in a single document. For unstructured data, such as documents, DLP solutions use techniques such as partial document matching to alert only if a portion of a protected document (usually a few sentences) is found, as opposed to keywords.

One large financial institution reported much fewer false positives once they built DLP policies on their live databases. This same institution also highlighted the importance of real false positives vs. "false" false positives. "False" false positive happen when you alert on a real credit card number, but it isn't one you care about (such as an employee on Amazon). A mid-sized credit union reported that while they see some false positives, the vast majority are ones they want to see and evaluate.

Lesson 6: Progressive deployments are most effective

Nearly every organization we talked with reported deploying DLP in stages; starting with one component and policy, then slowly expanding. This allowed them to better understand the new technology, tune internal workflows and processes and optimize policies.

Initial deployments tend to start as either network-centric or discovery-centric. With a network-centric deployment the organization starts with basic network monitoring, and then typically expands into email. Some organizations continue to expand into other network blocking, via gateway integration. In a discovery-centric deployment the organization starts with data-at-rest scanning, usually on servers and storage repositories, and then grows into endpoint scanning. This initial phase usually lasts one to two years (defined by budget cycles), then expands into the opposing channel or endpoint enforcement. We didn't find many DLP endpoint-centric initial deployments, perhaps because many organizations start with USB port blocking and encryption on the endpoint before moving into DLP.

In all cases, organizations report finding it better to start with a narrow set of policies and then expanding once incident types and volumes are better understood. On average, DLP managers said it takes about three to six months to tune a new policy, depending on its complexity. Simple policies, such as protecting a single collection of documents, require very little tuning, but more complex policies take time to refine. The general rule is to deploy any policy in monitoring mode and tune it to meet business objectives before moving into active enforcement/blocking. User notification, education, and disciplinary action during the monitoring phase materially lower violation counts and prepare the organization for potential business process impact.

Lesson 7: Endpoint DLP is still more limited than network or discovery

In network and discovery deployments, the DLP solution runs on high-powered, dedicated hardware. On the endpoint, the DLP agent must share resources with all the other cruft we load onto enterprise desktops and laptops. Thus, endpoint tools are more limited as to the type and number of policies they can run. Woe be on the DLP manager that attempts to load a policy containing the hashes for the entire customer database onto the sales team's laptop.

Not that endpoint DLP is unmanageable or too limited to be useful. Some tools communicate back to the central DLP server for content analysis when the system is on the same network. Since, in that situation, all email and network traffic are already monitored by the central server, only limited kinds of activities (like writes to USB drives) need to be offloaded. This also works well for endpoint discovery, where the local agent coordinates with the server for minimal impact. A few tools even support adaptive policies--where a smaller policy, such as a less-accurate regular expression, is only used when the endpoint can't see the DLP server. Yes, there will be more false positives, but remote activity can still be monitored and enforced.

Most DLP suite vendors started focusing more on the endpoint in 2008, but overall we see far less consistency across the different products than we do for network and discovery.

Lesson 8: Content discovery is hot

A security manager for a group of casinos reported they decided to start with content discovery over network monitoring. "We want a full solution, but the largest benefit will be in discovery. We just want to know where everything is. It's breach avoidance."

When interviewing independent references, fully half of them stated they started with, or are considering, data at rest scanning before network monitoring. Of this group, reducing PCI compliance costs and risk is the single biggest driver. Using DLP content discovery, they can inventory their environment for sensitive data to protect, reduce audit costs, and cut down on unneeded data exposure. Reduced audit costs alone, over time, can sometimes offset the total cost of the DLP tool.

Across all of our interviews two key trends emerged. DLP is clearly a viable option for real-world data protection, and many see it forming the core of their data protection initiatives. It can identify where your data is located, where it's moving, and how it's being used. On the other hand, few organizations are deploying DLP to its full capabilities, and products aren't the magical panacea often presented in sales meetings.

Rich Mogull is the founder of consultancy Securosis. Send comments on this article to feedback@infosecuritymag.com

Resources from Sophos

SOPHOS | Partner Program

How to protect against data theft and ensure that it remains confidential

Prescription for HIPPA Compliance

Stopping data leakage: Making the most of your security budget

About Sophos

Sophos enables enterprises to secure and control their IT infrastructure. Our network access control, endpoint, web, email and encryption solutions simplify security to provide integrated defenses against malware, spyware, intrusions, unwanted applications, spam, policy abuse, data leakage and compliance drift. We protect over 100 million users in nearly 150 countries.