



Quick Wins with DLP Light

Version 1.0
August 15, 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 Trend Micro



Trend Micro enables enterprises to encrypt, control, and protect confidential data—endpoint to cloud—with an ideal blend of data loss prevention, encryption, and device control solutions. The Trend Micro Data Protection portfolio offers standalone and integrated data protection solutions that deliver a consistent policy framework across endpoints, servers and gateways, and the automation and workflow drives down administrative burden. Virtual appliances allow companies to leverage existing infrastructure without having to purchase expensive, proprietary hardware appliances. And Trend Micro makes data protection easier with out-of-the box templates, pre-defined policies, active directory integration, and policy-driven, key management. This comprehensive portfolio keeps data safe from intentional or unintentional harm wherever it resides, protecting it against damage, loss, and unauthorized access.

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

Our entire profession is called “information security”, but surprisingly few of our technologies focus on actually protecting the data itself, as opposed to the infrastructure surrounding it. Data Loss Prevention emerged nearly 10 years ago to address exactly this problem. By peering inside files, network traffic, and other sources – and understanding both content and context – DLP provides new capabilities comparable to when we first started looking inside network packets.

The Data Loss Prevention market is split into two broad categories of tools – full suites dedicated to DLP, and what we call “DLP Light”.

There is plenty of confusion about the differences between these approaches – and even their definitions. In this paper we focus on DLP Light: what it is, how it works, and how to take advantage of it quickly. (For more information on full-suite Data Loss Prevention, see our white paper [Understanding and Selecting a Data Loss Prevention Solution](#).)

Defining DLP Light

We are talking about a subset of Data Loss Prevention, so let’s start with our definition of DLP:

Products that, based on central policies, identify, monitor, and protect data at rest, in motion, and in use, through deep content analysis.

A full DLP suite includes network, storage, and endpoint capabilities; as well as a range of deep content analysis techniques such as document fingerprinting.

DLP Light tools include a *subset* of those capabilities; they are generally features of, or integrated with, other security products. DLP Light is useful for organizations wanting to leverage existing tools to save costs, dip their toes into DLP before a bigger deployment, assess their need for DLP, or assess their data security problems. DLP Light tools tend to have some or all the following characteristics:

- *Focused on a subset of ‘channels’.* A DLP Light tool might focus on portable storage, email, web traffic, other channels, or a combination.
- *Fewer/simpler content analysis techniques.* Rather than providing a wide range of deep content analysis techniques, many of which are resource-intensive, DLP Light products tend to offer a smaller set of techniques, or even a single method.
- *Less dedicated workflow.* DLP Light tools are often integrated with, or features of, other security tools.

The key defining characteristic of both full DLP and DLP Light is **content analysis**. If a tool can peer into network traffic or a file and sniff out something like a credit card number, it’s DLP. If all it does is rely on tagging/labeling, metadata, or contextual information... it isn’t DLP.

Quick Wins with DLP Light: Technologies and Architectures

DLP Light tools cover a wide range of technologies, architectures, and integration points. We can't highlight them all, so here are the core features and common architectures.

Content Analysis and Workflow

Content analysis support is the single defining element for Data Loss Prevention – 'Light' or otherwise. Without content analysis we don't consider a tool or feature DLP, even if it helps to "prevent data loss".

Most DLP Light tools start with some form of rule/pattern matching – usually regular expressions, often with some additional contextual analysis. This base feature covers everything from keywords to credit card numbers. Most customers don't want to build their own rules, so the tools come with pre-built policies, which are sometimes updated as part of a maintenance contract or license renewal.

A Regular Expression for Credit Cards

```
^(?:(<Visa>4\d{3})|(?:<Mastercard>5[1-5]\d{2})|(?:<Discover>6011)|(?:<DinersClub>(?:3[68]\d{2})|(?:30[0-5]\d))|(?:<AmericanExpress>3[47]\d{2}))([ -]?(?:<DinersClub>(?:\d{6}\1\d{4})|(?:<AmericanExpress>(?:\d{6}\1\d{5})|(?:\d{4}\1\d{4}\1\d{4}))))$
```

DLP Light tools and features may or may not have their own workflow engine and user interface for managing incidents. Most don't include workflow, instead integrating policy alerts into whatever existing console and workflow the tool uses for its primary function. This isn't necessarily better or worse – it depends on your requirements.

Network Features and Integration

DLP features are increasingly integrated into existing network security tools, especially email security gateways.

- **Email Security Gateways:** Email gateways are one of the main integration points with full DLP solutions: all the policies and workflow are managed on the DLP side, but analysis and enforcement are integrated with the gateway directly rather than requiring a separate mail hop. Depending on your specific tool, internal email may be covered.
- **Web Security Gateways:** Some web gateways now directly enforce DLP policies on the content they proxy. These are also the tools you will use to reverse proxy SSL connections to monitor encrypted communications,.
- **Unified Threat Management:** UTMs provide broad network security coverage, and may include network DLP.
- **Intrusion Detection and Prevention Systems:** IDS/IPS tools already perform content inspection, and so are an easy place to add DLP analysis (which is usually very primitive rules-sets extensions).
- **SIEM and Log Management:** All major SIEM tools can accept alerts from DLP solutions, and many can correlate them with other collected activity. Some SIEM tools also offer DLP features, depending on what kinds of activity they can collect for content analysis..

- **Email Servers/Internal Mail:** While less common, some email servers include basic filtering capabilities, content or context-based encryption integration, or content-based archiving support..

Endpoint Features and Integration

DLP features have appeared in various endpoint tools other than dedicated DLP products since practically before there was a DLP market.

- **Endpoint Protection Platforms:** EPP is the term for comprehensive endpoint suites that start with anti-virus, and may also include portable device control, intrusion prevention, anti-spam, remote access, etc. Many EPP vendors offer basic DLP features – most often local file or portable storage monitoring. Some support for network monitoring.
- **USB/Portable Device Control:** Some of these tools offer basic DLP capabilities, and we are seeing others evolve to offer somewhat extensive endpoint DLP coverage.

Overall, most people deploying DLP features on an endpoint (without a dedicated DLP solution) are focused on scanning the local hard drive and/or monitoring/filtering file transfers to portable storage.

Storage Features and Integration

We don't see nearly as much DLP Light in storage as in networking and endpoints – in large part because there aren't as many clear security integration points.

- **Database Activity Monitoring and Database Vulnerability Assessment:** DAM products, many of which now include or integrate with Database Vulnerability Assessment tools, now sometimes include content analysis.
- **Vulnerability Assessment:** Some vulnerability assessment tools can scan for basic DLP policy violations if they include the ability to passively monitor network traffic or scan storage.
- **Content Classification, Forensics, and Electronic Discovery:** These tools aren't dedicated to DLP, but we sometimes see them positioned as offering DLP features.
- **Document Management and Collaboration Servers:** Integration into SharePoint or other content management systems.

DLP Light Software as a Service (SaaS)

Although currently no completely SaaS-based DLP services are available, some early SaaS offerings are available for limited DLP deployments.

- **DLP for email:** Many organizations are opting for SaaS-based email security rather than installing internal gateways, or combining the two. This is clearly a valuable and straightforward integration point for monitoring outbound email and performing DLP.
- **Content Discovery:** While still fairly new to the market, it's possible to install an endpoint (or server, at least on Windows) agent that scans locally and reports to a cloud-based DLP service. These target smaller to mid-size organizations which don't want the overhead of a full DLP solution, and don't have particularly deep requirements.
- **DLP for Web Filtering:** As with email, we see organizations adopting cloud-based web content filtering to block web-based attacks before they hit the local network, and to better support remote users and locations.

There are even more options on the market, including the occasional free or open source tool, but most fall within these major areas.

Quick Wins with DLP Light: the Process

The objective of the Quick Wins process is to get results and show value as quickly as possible, while setting yourself up for long-term success. Quick Wins for DLP Light is related to the [Quick Wins for DLP](#) process, heavily modified to deal both with the technical differences and the different organizational goals we see in DLP Light projects.

Keep this process in perspective – many of you will already be pretty far down your DLP Light path and might not need all these steps. Take what you need and ignore the rest.

Prepare

There are two preparatory steps before kicking off the project:

Establish Your Process

Nearly every DLP customer we talk with discovers 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.

Even if you aren't planning on moving straight to enforcement, you need a process in place to manage the issues that will crop up once you activate your tool. You should set up two different processes to handle the three common incident categories:

- **Business Process Failures:** DLP violations often result from poor business processes, such as retaining sensitive customer data and emailing unencrypted healthcare information to insurance providers. This process is about working with the business unit to fix the problem.
- **Employee Violations:** These are often accidental, but most DLP deployments identify some malicious activity. Your process should focus on education to avoid future accidents; and work with business unit managers, HR, and legal to handle malicious activity.
- **Security Incidents:** Traditional security incidents — usually from external sources — require response and investigation.

Determine Existing DLP Capabilities

The next step is to determine which DLP Light capabilities you have in-house, even if the project is driven by a particular tool. You might find you already have more capability than you realize.

Check for existing DLP features in the main technology areas covered above. It's also worth reviewing whether you are current on product versions, as DLP features might be cheap or even free if you upgrade (discounting upgrade costs, of course). Build a list of the DLP Light tools and features you have available, with the following information:

- **The tool/feature**
- **Where it's deployed**
- **Protected 'channels':** Network protocols, storage locations, endpoints, etc.
- **Content analysis capabilities/categories**
- **Workflow capabilities:** DLP-specific vs. general-purpose; ability to integrate with SIEM and other management tools

This shouldn't take long and helps choose the best path for implementation.

Determine Objective

The next step is to determine your goal. Are you primarily concerned with protecting a specific type of data? Or do you want to look more broadly at overall information usage? While the full-DLP Quick Wins process is always focused on information gathering vs. enforcement, this isn't necessarily the case in a DLP Light project. No matter your specific motivation, we find that individual projects sift into three main categories:

- **Focused Monitoring:** The goal is to track usage of, and generate alerts on, a specific kind of information. This is most often credit card numbers, healthcare data, or other personally identifiable information.
- **Focused Enforcement:** You concentrate on the same limited data types as above, but instead of merely alerting you plan to enforce policies and block activity.
- **General Information Gathering:** Rather than focusing on a single type of data, you use tools to get a better sense of information usage throughout the organization. You turn on as many policies as possible to monitor information of interest.

Choose Deployment Type

This is a three-step process for making the final decisions required to deploy:

- **Map desired coverage channels:** Determine *where* you want to monitor and/or enforce – email, endpoints (USB), etc. List every place you want to cover vs. what you know you already *can* cover with your existing capabilities. This also needs to map to your objective and content analysis requirements.
- **Match desired to existing coverage:** Now figure out what you have and where the gaps are.
- **Fill the gaps:** Obtain any additional products or licenses so that your project can accomplish your objectives.

Your entire project might be as simple as, "We want to catch credit card numbers in email using our existing tool," in which case this entire process up to now probably took about 10 seconds. But if you need a little more guidance, this should help.

Implement and Monitor

Now it's time to integrate the product (if needed), turn it on, and collect results. The steps are:

- **Select content analysis policies:** For a focused deployment, this will only include the policy that targets the specific data you want to protect, although if you use multiple products that aren't integrated you will use the most appropriate policies in each tool. For a general deployment turn on every policy of interest (without wrecking performance – check with your vendor).
- **Install (if needed)**

- **Integrate with other tools/workflow:** If you need to integrate multiple components, or with a central workflow or incident management tool, do that now.
- **Turn on monitoring**

We have a few hints to improve your chance of success:

- **Don't enable enforcement yet** – even if enforcement is your immediate goal, start with monitoring. Understand how the tool will impact workflow first, as we will discuss next.
- **Don't try to handle every incident at first.** You will likely need to tune policies and educate users over time before you have the capacity to handle every incident – depending on your focus. Handle the most egregious events now and accept that you will handle the rest later.
- **Leverage user education.** Users often don't know they are violating policies. One excellent way to reduce your incident volume is to send them automated notifications based on policy violations. This has the added advantage of helping to identify the egregious violators later on.

Analyze

At this point you have focused your project, picked your tools, set your policies, and started monitoring. Now it's time to evaluate your results and decide what's next. You might start by looking for the following:

- 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 more violations or unusual usage patterns.
- False positive patterns, for tuning policies later.

Then make two important decisions:

- **Is it time to enforce?** If you know you want to start blocking activity, determine the potential impact on business activities, then reduce violations to a manageable level with education and manual enforcement. The fastest way to kill a DLP project is to jump into enforcement too quickly and interfere with important operations.
- **Should we stay or grow?** If you are happy with your results, stop here. You may also decide to either enable additional policies or consider expanding your deployment through additional DLP Light tools/features, or even by migrating to full DLP. If you aren't ready to make these decisions now, put a date on the calendar to revisit them.

What Did We Achieve?

If you followed this process, by now you have a firm foundation for your ongoing DLP Light usage, ready to achieve useful short-term goals. In a short amount of time you have:

- Established a flexible incident management process.
- Integrated with major infrastructure components.
- Assessed information usage and risk exposure.
- Established a foundation for additional efforts and long-term management.

By following the Quick Wins process you can show immediate results while establishing the foundation of your program.

Who We Are

About the Author

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.