

Next-Generation Secure Web Gateways : The Case and Criteria for Embedded Data Loss Prevention

Author: Mark Bouchard

AimPoint Group
keeping IT on target

Executive Summary

Web 2.0 technologies have transformed the Web into an extremely viable and increasingly popular platform for business communications. At the same time, however, associated rich applications featuring real-time interaction and supporting user-generated content have also elevated its potential as a conduit for sensitive information and made the Web a highly attractive target/vehicle for hackers.

As a result, in addition to bolstering their formerly static Web defenses with real-time scanning, analysis, and classification capabilities, today's chief information, security, and compliance officers should also be considering how to address data loss over the Web channel. The secure web gateway (SWG) is a logical consolidation point in this regard, offering the potential for reduced infrastructure, complexity, and cost of ownership. Unless the data loss prevention (DLP) functionality obtained in this manner is a full-strength, enterprise-class set of capabilities, however, it will inevitably provide an inadequate level of protection and fail to deliver any of the other promised benefits.

Why Web DLP is Necessary

First things first, the term Web DLP is used to refer to the application of data loss prevention technology specifically to the Web channel of communications – that is, HTTP, HTTPS, and FTP traffic – as opposed to its use for other channels, such as email/SNTP, print, and other general-purpose traffic. As for why Web DLP is necessary, the first reason is that the Web has become a core platform for business communications and applications.

The introduction and widespread use of Web 2.0 technologies in particular, has dramatically changed the nature of the Web and how we use it. Instead of serving primarily as a repository for static information, it has now become a highly collaborative and extremely dynamic medium for communication. And the benefits of these characteristics apply not just to individual users in their personal lives, but also to businesses. Self-improving applications, the ability to efficiently harness the network effect of the Web, and rich, real-time interactions with customers are enabling enterprises to streamline numerous business processes while simultaneously improving customer satisfaction and retention. For instance, customer support, research and development, and payroll services are delivered over the Web, while social networking tools and sites are routinely used for recruitment, lead generation, and marketing.

The second, inherently related reason Web DLP is important is that the Web is now a bi-directional communications channel. Instead of traditional one-way transactions, new content can be generated and shared in real time, effectively becoming part of the application for other/future users. The downside, of course, is that it's now significantly easier not only to accidentally disclose important data, but also to do so intentionally, if one is so inclined. Not only that, but there also happen to be several feedback loops that are reinforcing users' tendencies in this regard:

Key Criteria for Secure Web Gateway DLP Capabilities

To ensure a superior level of effectiveness and otherwise maximize an organization's return on investment, data loss prevention capabilities incorporated as part of next-generation secure web gateways should match up well with the following, key characteristics and capabilities:

Comprehensive Visibility and Control

- *High accuracy detection*
- *Insight beyond identification*
- *Granular control*

Enterprise-Class Ease of Use & Management

- *Streamlined implementation*
- *High-efficiency operations*
- *Unified administration*

Enterprise-Class Architecture

- *Native DLP functionality*
- *Flexible deployment options*
- *Extensibility to full DLP*

- Software-as-a-service (SaaS) offerings are increasingly validating the outbound flow of data over the Web and, in the absence of contravening guidance, can lead users to believe it's acceptable to employ a wide range of other data export and sharing services to help get their jobs done (e.g., YouSendIt and Google Docs).
- Consumerization of IT is blurring the lines between personal and work lives, causing carry-over of what constitutes accepted practice – in this case widespread sharing of information – from one domain to the other.
- The availability of new sites and tools, such as Dropbox, SugarSync, and Syncplicity, are making it easier than ever to share and/or transfer data, as Web protocols and services rapidly displace other methods historically used to support these functions.

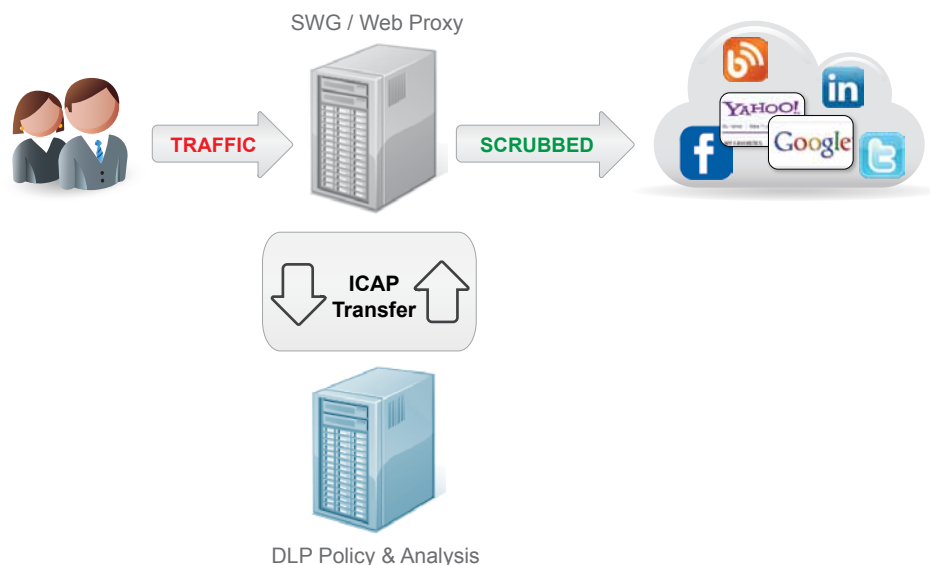
The net result of all these factors is: (a) that the Web has rapidly become the leading channel not just for data distribution, but for data loss as well, and (b) that when it comes to Web security, organizations cannot stop after upgrading their defenses to include real-time scanning and analysis capabilities, but must also invest in Web DLP.

Why Web DLP Should Be Part of the Web Security Gateway

So Web DLP is necessary. But what approach makes the best sense?

One possibility is to implement two distinct solutions, requiring at least a pair of devices to be deployed at each enterprise location to separately meet an organization's Web security and Web DLP requirements. A slight variation to this is an arrangement where the two solutions are somehow integrated. Typical in this case is having an inline Web security platform/proxy communicate – for example, using the Internet Content Adaptation Protocol (ICAP) – with an out-of-band DLP policy store and analysis engine.

The Integrated (multi-box) Approach to Web DLP



A far superior approach, however, is one that simultaneously protects users, data, and the corporate computing environment with a single, consolidated device. Indeed, some of the more significant, comparative strengths of a next generation SWG that natively incorporates Web DLP include the following:

- **It consolidates infrastructure and reduces costs.** Having fewer devices to purchase, deploy, operate and maintain reduces capital expenditures, complexity, and operational effort.
- **It simplifies and unifies policy management.** Policies for both Web security and Web DLP can be set in the same way and at the same time, promoting efficiency and reducing the likelihood of conflicts or omissions.

- **It minimizes introduced latency.** Packet handoffs and processing, even at the network and transport layers, take time. The need to conduct higher-layer operations, such as layer 7-protocol parsing, session reconstruction, and content analysis introduces considerably more delay. So why make matters even worse by requiring much of this processing to be done multiple times?
- **It facilitates inspection of encrypted traffic.** These days, a substantial portion of Web traffic is encrypted – a condition that is particularly true for organizations in regulated industries. SWGs generally include the capability to inspect such traffic. Furthermore, a combined solution can process it against associated DLP policies without having to send it, unencrypted, over the network to a separate device.

The Embedded (single-box) Approach to Web DLP



What Enterprises Should Look for in a Solution

Given that Web DLP is necessary and that it should ideally be incorporated as part of the SWG, the next task is to find a solution that provides the right set of functionality. Although many vendors now include a basic set of DLP capabilities with their SWG products, that really isn't good enough. Support for keyword matching, regular expressions, and so-called "custom signatures" is marginally useful, at best. In most cases, the effort required to configure such features and then sift through the resulting collection of mostly false positive events is simply not worth it – not to mention the high potential for false negatives and the inability to automatically respond to legitimate loss events in a meaningful manner (e.g., by blocking, quarantining, or encrypting associated data/traffic). The only thing worse than altogether missing an incipient data loss is detecting the occurrence of the event but not being able to do anything about it.

What organizations need instead is a SWG that natively incorporates a full-featured, enterprise-class DLP capability. Only with high detection accuracy, granular control mechanisms, and robust features that deliver ease of use and a high degree of flexibility will enterprises obtain a solution that is truly effective, yet remains affordable and efficient to operate. In this regard, specific criteria and capabilities that enterprises should use to evaluate SWG solutions offering DLP include those described in the following sections.

Comprehensive Visibility and Control

The fundamental basis for any DLP solution is the ability to accurately detect selected data elements and collections thereof, and subsequently to enable and/or directly take meaningful action in response to associated findings. Specific technical capabilities and criteria that can be used to gauge the potential of candidate solutions to achieve these objectives can be grouped into three categories as follows.

High accuracy detection. The effectiveness of DLP is dependent first and foremost on the presence of multiple detection mechanisms and content classifiers. In addition to basic keyword and regular expression matching methods, leading solutions will incorporate a broad range of additional techniques, such as the use of topic-specific dictionaries, file matching, digital fingerprinting, statistical analysis, correlation, custom/enterprise-developed classifiers, and natural language processing. Using these mechanisms, it should be possible to detect not just individual elements of data and entire documents, but also excerpts of documents that contain sensitive

information. So-called “smart” detection should also be supported, whereby the solution is able to detect small amounts of related data sent over multiple communication sessions, or, alternately, quantities of data that fall within a specified threshold – such as one or two credit card numbers, which is indicative of a legitimate transaction, as opposed to hundreds or thousands going by in matter of seconds.

Insight beyond identification. One of the greatest fears of organizations embarking on a DLP initiative is the expected frequency of false positives – that is, indications of data loss events that are either completely erroneous or, more likely, legitimate transmissions of the data elements in question. Inaccurate detection mechanisms and sub-par correlation capabilities are undoubtedly part of the problem. The greater issue, though, is that most solutions only identify what pieces of data are being lost and how the loss is occurring – for example, via which specific protocol and IP address. This minimal amount of context is simply not enough to quickly triangulate on the true nature of an event – a situation that leads to several negative consequences, such as: administrators having to spend an inordinate amount of time manually researching and validating or dismissing each detection event; not being able to reliably utilize automated response capabilities (e.g., actual prevention); and having reports that are cluttered with all sorts of meaningless information.

What organizations need to look for in this regard are solutions that provide a more complete picture of an event. The goal is to have ready access to (and optional automated correlation of) additional contextual information that can be used to further resolve the true nature of a detection event, such as the identity of the user involved and the actual destination, not just IP address, where the data is being sent.

“Only with high detection accuracy, granular control mechanisms, and robust features that deliver ease of use and a high degree of flexibility will enterprises obtain a solution that is truly effective, affordable, and efficient to operate.”

Granular control. This aspect of DLP involves two components. The more obvious of these is actually the back half of the pair, which entails being able to respond to a data detection event. In particular, the solution should support a wide range of actions, such as the ability: to block, quarantine, or encrypt corresponding sessions; to strip the offending content from a session without otherwise impeding communication; to obtain user/manager confirmation prior to sending; to run a script; and, of course, to log the event and all applicable details.

Somewhat less obvious – at least until one gets into the actual implementation process – is the front half of the pair. Specifically, there also needs to be a powerful policy framework. This is the mechanism that allows enterprise policies to be translated into granular rules that then lead to a specific action being invoked. Administrators should be able to delineate not just a specific element of data to look for, but also the originator, destination, and channel being used. This way financial records sent by the finance department to the company’s external auditor are not impeded, while an attempt by someone in engineering to send the same material to a location in Asia can be blocked and flagged for further investigation.

Elements of Powerful Policy Framework



Overall, this is another aspect of DLP that is central to the goal of minimizing not just false positives, but false negative ones as well.

Enterprise-Class Ease of Use and Management

Another of the great fears of organizations considering DLP technology is the complexity involved. Accordingly, it is imperative that any combined SWG-DLP solution incorporate features and capabilities to help simplify not just initial implementation, but ongoing operations too.

Streamlined implementation. The previously mentioned policy framework is instrumental in this area. But so too is the presence of a library of selectable, pre-defined data patterns and policy rules, as well as configuration wizards that automatically employ them. Indeed, wizards should be available for common scenarios, especially compliance as it pertains to financial regulations (e.g., PCI), electronic personal healthcare information (ePHI), or other personally identifiable information (PII). Implementing DLP for common scenarios should require little more than engaging a default configuration based on accepted best practices or, alternately, tabbing through a series of well-documented screens and making changes and selections to account for organization-specific requirements. For that matter, even configuration for completely unique situations should be relatively straightforward, perhaps based on the presence of a generic wizard.

High-efficiency operations. Initial setup and configuration is only one piece of the puzzle. Complexity and resulting inefficiencies must also be wrung out of subsequent stages of operation as well. In this regard, “insight beyond identification” definitely has a major role to play. The time savings enabled by providing additional contextual information quickly adds up – easily to the tune of at least 10 to 15 minutes per detection event. Additional savings can also be had if the solution incorporates workflows for common processes, such as: (a) the creation/preservation of a comprehensive audit trail for incidents involving highly sensitive information and/or that have potential legal ramifications; (b) when an incident requires review, involvement, or acknowledgment by a series of internal parties (e.g., HR, legal, business-unit management); and (c) combining all records applicable to a group of related incidents into “a case”.

Unified administration. Following in the same vein, unified administration is also primarily about reducing complexity – in this case, by not requiring multiple management consoles. Instead, the goal should be to have combined and sensibly inter-woven administration capabilities that not only span all lifecycle management functions – such as configuration, monitoring, analysis, troubleshooting, and reporting – but each domain of coverage (i.e., both Web security and Web DLP together) and all deployment options as well. Given the scope and power of such a management solution, of course, enterprises should also look for granular, role-based administration capabilities. In any case, the net result should be the ability to easily track, analyze, and respond to all content and data security events in one place, while also being able to generate and distribute detailed and executive-level reports that accurately illustrate the state of affairs.

Enterprise-Class Architecture

Having a robust set of DLP capabilities that are easy and efficient to use is one thing. Ensuring they are affordable and widely applicable and, therefore, represent a good investment, however, is quite another matter – and one that also needs to be addressed.

Native DLP functionality. This item almost goes without saying at this point, but is important enough that it deserves repeating. Simply put, preference should be given to DLP capabilities that are natively part of the SWG, rather than being integrated or completely standalone. All else remaining equal, embedded DLP is the most effective and least costly to own and operate.

Flexible deployment options. Affordable, best-fit coverage should be available for all enterprise scenarios, ranging from mobile users, to branch offices, and large central site implementations to different operating models and network configurations. As a result, an ideal solution should come in multiple form factors, including software, virtual appliance, hardware appliance, and software-as-a-service. It should also support both inline and pass-by arrangements and be configurable, down to the individual policy level, to optionally allow monitoring-only, prevention, and any combination in between. Overall, a high degree of flexibility is essential to ensuring a high degree of utilization and, consequently, an excellent return on investment.

Extensibility to full DLP. Another major way to provide a significant measure of investment protection is to have an extensible solution. In this regard, CIOs/CISOs need to recognize that although it's an excellent place to start given the prominence of the Web channel and its growing capacity to enable the outflow of sensitive data, Web DLP is only one facet of full, enterprise DLP. Accordingly, an ideal solution is one that not only tackles what constitutes most enterprises' largest and most immediate problem, but also allows them, at their own pace and discretion, to efficiently expand their DLP implementation to also cover data at rest (e.g., in file stores or databases), data in use on endpoints, and other communication channels (e.g., email) ... all without having to modify existing DLP rules and reports or implement yet another management system.

What Enterprises Stand to Gain

The bottom line is that Web 2.0 technologies have transformed the Web, in turn changing what constitutes adequate Web security and dramatically elevating the need for Web DLP. In this regard, enterprises will be served best by purchasing and implementing a combined solution – a secure web gateway with embedded DLP capabilities. This conclusion only holds true, however, to the extent these capabilities represent full-strength, enterprise-class DLP functionality as described by the criteria outlined herein. For enterprises that embrace next-generation secure web gateways that do in fact meet these criteria, benefits include:

- Having the ability to safely and securely use transformative Web 2.0 services and technologies;
- Having the ability to minimize corporate and compliance risks due to data loss;
- Reduced infrastructure cost and complexity based on having a unified Web and data security solution;
- Reduced operational effort based on having an intelligent, enterprise-class solution; and
- The peace of mind and investment protection that comes from having an extensible solution that enables a practical, progressive approach to DLP.

About the Author

Mark Bouchard, CISSP, is the founder of AimPoint Group, an IT research and analysis firm specializing in information security, compliance management, application delivery, and infrastructure optimization strategies. A former META Group analyst, Mark has assessed and projected the business and technology trends pertaining to a wide range of information security and networking topics for more than 14 years. During this time, he has assisted hundreds of organizations worldwide with strategic and tactical initiatives alike, from the development of multi-year strategies and high-level architectures to the justification, selection, and deployment of security and networking solutions. A veteran of the U.S. Navy, Mark is passionate about helping enterprises address their IT challenges.