



The Power to Prevent:
It's the Next Generation

FIDELIS SECURITY SYSTEMS, INC. 4416 EAST WEST HIGHWAY, SUITE 310, BETHESDA, MD 20814



A Buyer's Guide to Network DLP

Fidelis Security Systems
February, 2009
Version 6.0

Table of Contents

Executive Summary	Page 3
Defining Network DLP	Page 3
Key Requirements of Network DLP	Page 4
Value of Network DLP	Page 5
Key Features for Network DLP	Page 5
Control to stop network data leakage	Page 6
<i>Prevention via a network sniffer</i>	Page 8
<i>Prevention via proxy/gateway servers</i>	Page 8
<i>Combining the sniffer/proxy/gateway</i>	Page 9
Accurate inspection of enterprise content	Page 10
<i>Defining accuracy</i>	Page 10
<i>Visibility</i>	Page 11
<i>Content analysis methods</i>	Page 11
<i>Registration</i>	Page 12
<i>Profiling</i>	Page 12
Enterprise scale	Page 13
<i>Network performance</i>	Page 13
<i>Supported zones of control</i>	Page 14
Enterprise architecture-friendly	Page 15
<i>Minimizing operational impact</i>	Page 15
<i>Adding value to existing IT assets</i>	Page 15
Robust architecture	Page 16
Appendix A: RFI Checklist	Page 17

Executive Summary

Data leakage can have an incredibly negative impact on an organization's brand, customer trust, and resources, potentially resulting in legal liability, lawsuits, fines, penalties, or incarceration. Because of the potential effects of a data leakage, simply reporting a leak is not sufficient. Prevention is required to ensure digital assets are secured and to demonstrate that an organization is exercising adequate care to avert the loss of their digital assets.

Network data leakage prevention (network DLP) is the process of stopping the unauthorized disclosure of digital assets out of computer networks, regardless of the channel of communication. This process is typically implemented to protect digital assets—intellectual property, personal identity information, financial information, sensitive/protected data—and enables an organization to detect and prevent potentially harmful leaks.

To ensure an organization's digital assets are well protected, it is important to choose a full-featured network DLP solution. Many solutions in the data leakage prevention marketplace today are more akin to monitoring and alerting solutions, able to alert when data has been leaked, but can not do anything about it or can only take action in a small number of circumstances. To actually protect digital assets, it is imperative to implement a solution that not only identifies but prevents a leak *before* it occurs. The following key attributes are critical to the solution's success:

Control to stop network data leakage—by stopping leaks before they occur without dependency on proxy servers, which cannot be depended upon to see all network traffic. Prevention capabilities should address all channels, not just e-mail and/or web traffic—including HTTP, FTP, webmail accounts, instant messaging (IM), internet relay chat (IRC), and peer-to-peer (P2P) channels.

Accurate inspection of enterprise content to ensure data leakage prevention. Content recognition technologies must go beyond exact matching for ease of deployment and solution scalability. Network DLP solutions must analyze all traffic and not sample or drop traffic on high-performance networks.

Enterprise network scale on multi-gigabit networks, to ensure all leaks can be prevented in real time and have the ability to support zones of control beyond the internet perimeter at key internal points including data centers and organizations.

Enterprise architecture-friendly with no impact on network performance, enterprise architecture, or server/desktop configuration. With these key attributes, an organization will have the ability to stop the unauthorized transfer of protected data in real time on all channels—the necessary requirements to prevent leakage and the detrimental financial, legal, and business consequences that accompany it.

Robust architecture that is able to adapt to the changing threat landscape and IT infrastructure, while meeting requirements both today and into the future.

Defining Network DLP

Network Data Leakage Prevention (DLP) is the process of stopping the unauthorized disclosure, whether unintentional, accidental, or intentional, of digital assets out of computer networks, regardless of the channel (a channel is the transport mechanism for data, and includes e-mail, instant messaging, webmail, or FTP), of communication. Network DLP enables an organization to detect and prevent potentially harmful leaks of digital assets, which includes information such as: personal identity information about customers, employees, or citizens; credit card cardholder data; health care records; classified information; product designs; source code; intellectual property; financial information; and any other information deemed protected and/or sensitive.

Network DLP is a paradigm shift from traditional security tools, which have previously focused on threat mitigation (firewalls and intrusion prevention) and determining roles for information access (authentication and access control). These controls are necessary but insufficient to protect against the risks presented by the unauthorized transfer of protected information.

Key Requirements of Network DLP

In order to be successful, a network DLP solution must be able to understand both the information being communicated, as well as how the information flows across the network. A network session could violate content disclosure policies by containing protected or sensitive information, violate network use policies by using channels or applications disallowed on the network, or both!

In order to stop data leaks across multiple communication channels on the network, at a minimum a network DLP solution must meet certain base requirements:

Requirements	Benefit
<ol style="list-style-type: none"> 1. Conduct session-level (not packet-level) inspection of network traffic across all 65,535 network ports. 2. Provide visibility into the protocols, channels and applications in use on the network 3. Be able to extract enterprise human-readable content and related meta-data contained in the session and any attachments and compressed files for analysis. 4. Provide multiple sophisticated content analysis technologies to detect sensitive and/or protected information. 5. Provide a policy engine to implement rules to determine network sessions that violate policy. 6. Ability to prevent an individual network sessions that violates policy across all 65,535 ports. 	<ol style="list-style-type: none"> 1. The majority of enterprise human-readable data is not in plain text formats and requires knowledge beyond an individual packet for DLP decisions. 2. The network DLP solution needs to understand a wide variety of the network traffic to determine how people are communicating and extract information for analysis. 3. Network DLP needs to be able to see all the content even in nested attachments and compressed files. 4. Accurate content recognition. 5. Policy-based enforcement mapping rules to an organizations content disclosure or network use policies. 6. To prevent data leaks across all network traffic.

When a network DLP solution detects a network session transferring information, it checks to see whether the channel is authorized whether sensitive information is included, and if so, whether the destination is authorized to receive that information. If the transfer is unauthorized, the solution can stop the transfer of the information and prevent the leak from occurring.

Because of the potential grave consequences ranging from the financial and regulatory risks of information leakage, it is critical that organizations choose a network DLP solution that meets the requirements of their organization. This white paper will explain the key requirements and decisions to choose a network DLP solution for your organization.

The Value of Network DLP

Today, the risks of data leakage are very real—it seems as if a data leakage incident appears in the newspaper headlines every week, if not daily. These leaks can have a profoundly negative impact on an organization, potentially resulting in legal liability, lawsuits, fines, penalties, incarceration, and significant damage to an organization's reputation and brand.

DLP, much like intrusion prevention and threat mitigation, is available in network, endpoint, and scanning solutions. Much like intrusion prevention, most organizations choose to start with the network because the maximum amount of risk reduction can be achieved with minimal resources and operational impact. Network DLP also delivers quicker time-to-value and lower total cost-of-ownership when compared with deploying or scanning tens or hundreds of thousands of endpoints. Organizations' IT budgets are continually strained, so the ability to mitigate the risk of data leakage by deploying technology at key network egress points, often with no or minimal incremental resources is a very attractive prospect.

**The direct method of securing digital assets against theft is inherently simpler:
If 10,000 people have access to a digital asset, it is much easier to focus on securing the asset
than on controlling the people.**

Most organizations start their network DLP deployment at their organization's main internet connections. However, it is important to note that network DLP deployments can cover other network egress points. Potential locations where network DLP can be deployed include:

- Organization's internet connections;
- Extranets or trading partner networks;
- Data centers;
- VPN concentrators;
- Divisions;
- Departments;
- Campus/Remote offices;
- As well as any other network egress point where sensitive information flow should be restricted.

Key Features for Network DLP

Beyond the base requirements outlined above, which are required to even be considered a network DLP solution, organizations should look for the following key features when choosing a network DLP solution:

- Control to stop network data leakage;
- Accurate inspection of enterprise content;
- Enterprise network scale;
- Enterprise architecture-friendly;
- Robust architecture.

This white paper will discuss each of these requirements in detail to provide you with the information required to select the right network DLP solution for your organization.

Control to stop network data leakage

The "P" in network DLP is for *prevention*. Leaking protected or sensitive information can have a profoundly negative impact on an organization. Because of this potential effect, it is important to examine a solution's ability to actually prevent the leak, rather than just alert that it happened.

Would you rather have a solution that alerts you that your sensitive information has been leaked, or one that actually prevents it?

Unfortunately, many solutions that purport to be DLP solutions are actually data leakage detection or alerting solutions with no prevention capabilities, and others have very limited prevention capabilities depending on integration with third-party products. Detection of a leak is necessary, but insufficient to protect an organization against the harmful consequences of data leakage. Many organizations and technology solutions lose sight of the goal: Detection enables a report on the state of the organization's compliance position and describes when policies were breached—it doesn't stop the breach from occurring, only prevention can.

A comprehensive ability to prevent data leakage is required to stop data leakage, ensuring digital assets are secured and demonstrating that an organization is exercising adequate care to avoid leaks of sensitive or protected information. While one can be compliant with certain regulations by solely reporting a breach after it occurs, that does nothing to protect the organization from the ensuing customer mistrust, potential civil and regulatory penalties, and the massive distraction of the investigation and reporting of a breach. In fact, one could argue that data leakage detection, versus prevention, automates data breach notification instead of stopping data breaches.

When evaluating network data leakage prevention solutions, it is important to understand what type of traffic can be prevented by the solution. Internet Protocol (IP) networking defines 65,535 ports, which are sub-addresses or logical locations allowing two computers to connect simultaneously over a variety of protocols. Most people are familiar with these port numbers from firewall rule set configuration. These 65,535 ports can be grouped into three ranges:

- Well known ports, which were reserved for particular services run by privileged users of a system, cover ports 0 through 1,023;
- Registered ports, also defined by IANA and now managed by ICANN, cover ports 1,024 through 49,151 and are designed for use by user applications;
- Finally, dynamic or private ports which can be used for any purpose cover ports 49,152 through 65,535.

While some protocols still observe their official defined port, the port/protocol paradigm is unfortunately no longer dependable for security controls. Some network services have been deployed on non-standard ports to enable multiple services to be offered on the same machine. However, a majority of the traffic running on non-standard ports is likely to be traffic attempting to evade controls. As consumer applications have evolved, they have implemented port hopping to find ways to work through enterprise firewalls. This is particularly true of instant messaging and peer-to-peer clients—both present significant risk for data leakage! As an example, AOL Instant Messenger (AIM) can run via its proprietary protocol, OSCAR, on any open port or can be tunneled via HTTP. A solution only looking for AIM via OSCAR on 5190 would miss a large amount of the instant messaging traffic in most enterprises.

Applications that port hop can easily evade analysis tied to a single port number.

Protocol:	Defined Port:	Easily available on:
SMTP	25	21, 23, 587, 2525, 5521, 5525, 7721, 7725 or ANY PORT with root access to a system
AOL Instant Messenger	5190	ANY PORT
HTTP Proxy	80	808, 2301, 3124, 10201, 12323, 30003, 37474, 48725, or ANY PORT with root access to a system
ICQ	5190	All ports above 1024

In order to inspect content everywhere information flows, the solution must be port and protocol independent. Many vendors claim to view all network traffic, but are still tied to the port and protocol paradigm. Because of this, many network DLP solutions can only control traffic on a small number of ports, versus the 65,535 where network traffic can flow. This is because many require the use of an application gateway or proxy server, which operate at layer 7 of the OSI model, in order to implement prevention capabilities. This means the solutions that depend on these devices can only prevent for traffic that is passing through the device on the port(s) it controls. That typically means providing prevention capabilities for one to three ports, or 0.005% of ports on the network! These solutions typically cover port 25 for the organization's standard outbound SMTP-based e-mail using an SMTP relay, port 80 for the organization's standard HTTP web browsing traffic and perhaps port 443 for the organization's standard HTTPs secure web browsing traffic when the proxy supports decrypting SSL traffic. Any traffic that uses a non-standard port or bypasses the proxy server leaves the network uncontrolled.

Next-generation network DLP solutions, like the Fidelis Extrusion Prevention System® (Fidelis XPS™), can provide prevention capabilities for all 65,535 ports unlike proxy-based solutions. All ports coverage is required to prevent data leakage for all outbound network traffic.

It is also important to evaluate the architecture used for prevention. There are two main categories of prevention capabilities:

1. Network sniffer, Data Link Layer of the OSI Model, capable of preventing traffic on all 65,535 ports.
 - Out-of-band prevention via session poisoning with TCP RST (reset) packets.
 - Inline dropping of sessions.
2. Proxy/Gateway, Application Layer of the OSI model that views traffic on a specific protocol on a specific port.
 - Proprietary gateway with dedicated proxy/gateway.
 - Integration into existing proxy/gateway solution.

Each of these architectures provides different prevention capabilities, impact on the network infrastructure, and user experiences. Any solution should have both capabilities to provide the benefits of each approach.

Prevention via a network sniffer

A network sniffer, unlike a proxy server, operates at the Data Link Layer of the network and therefore can see all of the traffic flowing on the network. This architecture also can be port and protocol independent enabling visibility into traffic where channels run on non-standard ports or protocols.

However, a sniffer alone is not a prevention device—there are many monitoring products on the market that use a sniffer-based architecture. To actually provide prevention, a sniffer must be able to analyze the content of the network traffic in real-time in memory and be able to take preventative action before the session completes. If any of these requirements are not met, the sniffer will only be able to monitor the network and report after a leak had occurred.

A prevention sniffer can be deployed in two configurations in the network, either inline or out-of-band. An inline sniffer can actually terminate a network connection by dropping the packets flowing in the session, but it is also another “hop” in the network and a potential failure point. In an out-of-band configuration, the prevention sniffer is not inline of network traffic; however it receives a copy of all network traffic via a network tap or Switched Port Analyzer (SPAN) port. When out-of-band, the prevention sniffer can terminate a session by instructing the sender and recipient to end the session. This is accomplished with a TCP Reset (TCP RST) packet, which instructs the sender and recipient to terminate the network connection.

Fidelis XPS is the only network DLP solution available that does not depend on a proxy server for prevention, and therefore can prevent on all ports, including real-time protocols. Fidelis XPS is also the only solution that can prevent out-of-band, enabling all-channels prevention without slowing down, or interrupting real-time protocols. This is in direct contrast to proxy-based prevention solutions which do not see all traffic flowing on the network leaving significant risk unmitigated.

Prevention via proxy/gateway servers

A proxy is a server that accepts client connections and then requests the requested resource from another system, enabling a client to connect to the other system without a direct connection. Proxy servers operate at the application layer of the OSI model, and most proxy servers deployed in the enterprise are for the HTTP protocol to control how clients browse web servers on the internet and to cache content locally to accelerate client requests. Proxy servers can prevent a session routed through by ending the connection with the client.

Gateways, which are very similar to proxy server but designed for store and forward protocols, are typically deployed for the SMTP protocol for the organization's standard e-mail traffic, and are typically referred to as a mail transfer agent or MTA. The gateway accepts e-mail from mail servers, inspects them (if configured to do so), and forwards them on the appropriate mail server. Gateways, and MTAs in particular, can prevent a session by either deleting the appropriate message or by not forwarding it on pending further review. This process is referred to as quarantining.

Proxy and gateway solutions are available in the two configurations mentioned above, either proprietary solutions or via integration into existing proxy or gateways.

Proprietary solutions include both the proxy and the DLP solution bundled in a single solution. This provides the benefit of integrating the DLP solution into the proxy/gateway, and is the recommended solution for an organization that has no proxy/gateway solutions in place. If an organization has proxy/gateway solutions, it must integrate or chain these solutions together, which often requires reconfiguration of desktops, servers, or mail routing.

Solutions that integrate into existing proxy or gateway solutions take advantage of the existing deployed proxy or gateway solutions, avoiding reconfiguration of desktops, servers, or mail routing. This integration is typically via supported standards. The standard for inspecting HTTP (and HTTPS) traffic via a web proxy server is the Internet Content Adaptation Protocol (ICAP), which enables the web proxy to send session content to a third party for analysis, in this case the network DLP solution. For e-mail, the Milter standard, which originated as the Sendmail Content Management API, but has now been adopted by other products as the interface is open source. Integrating into existing solutions tends to be easier than chaining proxies or gateways together, so this is typically the recommendation approach in most organizations.

The benefit of a proxy or gateway server is that they typically provide for a graceful user experience in the prevention process. As they control the individual session or message at the application layer, they have the ability to interact with the end-user. For web browsing traffic, that means the user can be redirected to a policy page explaining the policy violation when they attempt to send out information that violates disclosure policies. For e-mail traffic, the user can be notified of the policy violation with an automated e-mail response, and also can be notified of the quarantine status of the message.

Proxy servers have a few drawbacks. First, since they integrate at the application layer, they require terminating and reinitiating the session. This causes a performance impact on network traffic. Beyond performance, however, is a bigger drawback of reduced visibility. Since proxies and gateways are bound to specific protocols on specific ports they do not see all the traffic flowing on the network. Proxy servers also require endpoint changes (typically in the web browser or other application settings) to ensure traffic is routed to the proxy server. When a proxy server is circumvented, which can be easily accomplished when a firewall has one or more ports open without protocol compliance, the proxy cannot see and therefore can not alert or prevent on any traffic it does not see.

Prevention only via proxy servers is much like putting toll booths in only a few lanes in a highway—people will quickly figure out how to evade the toll.

Combining the Sniffer, Proxy, and Gateway for a Comprehensive Solution

An ideal network DLP solution will combine the coverage and risk mitigation benefits of a prevention sniffer with the graceful handling of traffic of proxy and gateway solutions. In order to best understand the value of the combined solution, it is important to understand the three varieties of communications protocols and applications in use on most computer networks: official, unofficial, and rogue.

Official	Protocols and applications endorsed for official business communications over the internet. This typically included SMTP e-mail, web browsing via HTTP and HTTPS, and file transfers via FTP for system-to-system transfers. All official communications typically occur on the standard defined well-known port numbers.
Unofficial	Unofficial protocols and applications are the applications that are not formerly covered by formal network use policies, but in some cases may have positive or negative views of their use in the organization. Unofficial protocols and applications are typical consumer applications like instant messaging, social networking, and webmail. Some of these applications, particularly instant messaging, will port hop or tunnel inside of other protocols to avoid traditional port-based security inspection.
Rogue	Protocols and applications that are explicitly denied on the network. Peer-to-peer (P2P) services are designed for illegal file sharing and are almost uniformly prohibited. On some networks, instant messaging, webmail, and other consumer-based services are considered rogue. Many of these applications, particularly instant messaging and P2P services will port hop or tunnel inside of other protocols to avoid traditional port-based security inspection.

A comprehensive network DLP solution needs to be able to analyze all of the official and unofficial communications in the organization, and block rogue communications (whether analyzed or not)—all regardless of port, to ensure sensitive information does not leak.

The prevention sniffer is the only network DLP architecture to provide prevention for all types of traffic, including official, unofficial, and rogue applications. The prevention sniffer provides this robust coverage because it is the only approach to cover all 65,535 network ports. As the prevention sniffer is the only solution that covers all ports, it is also the only approach that can control port-hopping or tunneling, which is prevalent in the consumer applications typical of unofficial and rogue communications.

Organizations also typically start their prevention implementations with the prevention sniffer, as it enables comprehensive detection paired with prevention for unofficial and rogue traffic, where organizations are most comfortable preventing. Protected or sensitive information shouldn't be leaving the organization at all, however there may be official communications that contain mishandled sensitive information. However, unofficial and rogue protocols should not be used for business and therefore should never contain digital assets, so the likelihood of prevention interrupting business process is very low.

Proxy and gateway solutions are a good fit for official communications traffic. As official communications occur on standard ports, the drawback of weak port coverage of a proxy server is minimized. Additionally, as official communications are how you choose for your organization to communicate, user feedback is more important. The ability to provide users with policy feedback, and in the case of e-mail, quarantine a message for release following review, provide the controls to not only prevent data leakage, but to also educate the user on their actions.

Any network DLP deployment should start with a prevention sniffer to gain all ports coverage, and, if required, follow with targeted proxy or gateway solutions to cover the organization's official channels of communication in a graceful manner.

Fidelis XPS provides the only prevention sniffer available in the Fidelis XPS Direct sensor. Additionally, Fidelis provides Fidelis XPS Proxy, which integrates via ICAP to any ICAP-compliant proxy server. Fidelis XPS also provides Fidelis XPS Mail, the only network DLP mail solution that supports both SMTP gateway, as well as Milter-based integration to existing mail gateways. The combination of Fidelis XPS Direct, Fidelis XPS Proxy, and Fidelis XPS Mail provides an organization the most comprehensive network data leakage prevention solution available for both comprehensive prevention and options for managing the prevention experience for the end-user.

Accurate inspection of enterprise content

Defining accuracy

While the architecture of the solution must be capable of preventing a data leak across all ports on the network, the network DLP solution must also be able to accurately detect sensitive or protected information to prevent leaks of digital assets. If the system cannot detect the information, it cannot alert on or prevent the unauthorized disclosure from occurring.

When evaluating content analysis, both the potential for false positives and false negatives should be analyzed. Significant false positives can make a system difficult to manage. Hence, much discussion is focused on false positives. But false negatives present a much greater risk for data leakage and compliance than false positives. False negatives mean that protected information is disclosed without the system generating an alert, thus bypassing the prevention and remediation processes required under internal policy and external regulation. A false negative could expose your organization to finding out about a leakage of digital assets from the news media or regulators—what the DLP solution was deployed to avoid in the first place.

Additionally, it is important to talk about the cost of ownership or maintenance required to manage the content inspection. Many vendors will attempt to discuss false positives, or the incorrect generation of an alert, without discussing false negatives and cost of ownership. All three are of significant importance when choosing a network DLP solution.

	Definition	Error Type	Impact
False Positive	“When a test incorrectly reports that it has found a positive result where none really exists.” (Wikipedia)	Type 1 Error	An alert is generated for valid network traffic that does NOT violate policy.
False Negative	“When a test incorrectly reports that a result was not detected, when it was really present.” (Wikipedia)	Type 2 Error, Miss	An alert is NOT generated for unauthorized network traffic that does violate policy.

Visibility is a key component of accuracy

Before discussing content analysis methods, it is important to understand what visibility the network DLP solution has into the network traffic. If the solution cannot analyze or understand the network session because it doesn't understand the protocol/application or isn't looking for it outside of standard network ports—it will trigger a false negative anytime sensitive information is sent in that fashion!

First, it is important to understand what traffic the network DLP solution understands. If a DLP solution advertises "all channels", "all protocols", or "all applications", it is important to understand what that means. Most network DLP solutions have an unknown protocol decoder that attempts to find information. However, these unknown decoders are highly inaccurate, often suffering from both significant false positives from noise on the network and false negatives from the inability to recognize the information in the session. The list of actual decoded protocols provides a much more realistic view into the information actually understood by the network DLP solution.

Second, it is important to understand where the network DLP solution looks for the protocols it understands. The network DLP solution, to avoid risks of false negatives from port-hopping, must look for all of the channels on all 65,535 ports. Unfortunately, many network DLP solutions only look for the protocols on well-known ports (e.g., SMTP on port 25) or require the customer to tell the solution where to look without providing the ability to cover all ports.

A comprehensive network DLP solution needs to be able to analyze all of the communications, both official and unofficial, in the organization to avoid false negatives.

Content analysis methods

Also, a variety of content recognition methods exist. However, they almost all fall into two general categories: profiling and registration.

Profiling uses rules that describe information, typically statistical, pattern, and/or key attributes that the system uses to evaluate information. It does not require that the actual protected information be provided.

Registration requires that protected content be enrolled in the system. This system then generates algorithms to detect an exact match or fingerprint of the actual content that has been registered with the system.

Registration

The first generation of network DLP solutions invested heavily in registration technologies, specifically exact matching technologies. Exact matching has an attractive message of incredibly low false positives. However, in order for exact matching to be successful, *all* protected information must be enrolled. Any information that would violate policy, but was not enrolled, would leave without being detected—a false negative and a resulting data leak!

The initial enrollment process for sensitive information is a daunting, if not impossible, task which makes the deployment feel more like that of an ERP solution rather than that a network security solution. Many organizations have spent years trying to determine what digital assets they have and still have not completed the task. While this inventory is difficult, imagine the challenges of attempting to enroll or register all of this information with a content monitoring solution. This alone could delay the implementation, and therefore value, by months or years.

Could your organization actually create a comprehensive list of all of its digital assets and keep it up to date on a daily basis in order to depend on matching technologies?

Even if that initial implementation were accomplished, any information that was modified since its registration would also not be detected, which creates a total cost-of-ownership challenge. The continuous integration of new digital assets into the system creates significant cost of maintenance over the entire life of the solution, drastically raising the total cost-of-ownership of exact matching-based solutions.

To attempt to deal with the daunting task of keeping registrations current as data changes, many first generation technologies then moved to partial matching algorithms, which could deal with some level of changes to the protected information, as well as derivative works created by "cut-and-paste". However, not only does this still have the initial registration challenges, it also creates an accuracy issue—how to accurately determine how much information needs to be detected to pose a risk?

As registration presents challenges, particularly with scalability, it should only be used for small information sets like a critical marketing plan or a CEO's letter. However, other methods are typically required to scale to the enterprise requirements for protecting sensitive information.

Profiling

To be successful, a network DLP solution must go beyond registration and matching to ease-of-deployment and solution scalability. Profiling of information, or describing digital assets, is inherently more scalable. Rather than requiring matching of the exact information to be protected, it enables an organization to create policies that detect digital assets without the intensive registration and maintenance processes common to registration-based solutions. Because of this, profiling also greatly shortens the time-to-value of the network DLP solution, as it can deliver value immediately rather than being dependent on slow and time-consuming registering of information.

Profiling technologies can range from simple keywords to complex statistical analysis. Examples of profiling-based content analysis:

- Keywords, key phrases;
- Pattern matching and regular expressions;
- Algorithms, including checksums and validators;
- File type, file name;
- Statistical analysis.

Profiling, as it doesn't require data registration, typically has very few, if any false negatives. The amount of false positives generated can vary, based on the individual content analysis technology and how it is applied.

For example:

- Keywords are very effective with low false positives for unique terms, but can present a false positive challenge for more general use in common words;
- Pattern matching and regular expressions are very effective for unique, predictable alphanumeric sequences, but present a false positive challenge for small numeric-only values;
- Checksums and validation algorithms typically have very low false positives, trending to acceptable levels, while avoiding false negatives. Checksums and validations are available for many common attributes including credit card numbers, Social Security numbers, ZIP codes, states, and ABA routing numbers.

With profiling it is also possible to reduce false positives to nearly zero by combining different profiling content analysis. For example, a combination of keywords, combined with statistical analysis and validators. The ability to combine multiple profiling technologies will typically provide the incredibly low false positives of matching with the incredibly low false negatives of profiling—all without the requirement for registering information. A network DLP solution that allows for combining profiling analyzers not only provides the lowest total cost-of-ownership, it also gives the organization the highest levels of confidence in identifying a digital asset.

In order to be successful, a network DLP solution must go beyond registration-based content analysis and also include granular profiling-based content analysis. Fidelis XPS goes beyond matching technologies to provide the most robust content analysis covering registration and profiling available in a network DLP solution. For registration, Fidelis XPS supports both exact matching, as well as the most robust partial matching engine. More importantly, Fidelis XPS provides a wide variety of profiling-based analysis, including Smart Identity Profiling™ which delivers the most accurate analysis for lists of personal identity information without requiring data registration. Additionally, all Fidelis XPS content analyzers can be logically combined in a single rule to provide the most granular rules providing organizations the accuracy they require in a network DLP solution.

Enterprise scale

In order to be successful, the network DLP solution must be able to scale to keep up with the network and computing infrastructure of the organization. When evaluating scale, there are two critical areas to consider. First, network performance—or the ability for the network DLP solution to analyze the traffic at the wire-speed of the network. Second, supported zones of control—defined as the locations in the network the network DLP solution can be deployed.

Network performance

In order to be able to prevent data leakage on the network, the network DLP solution must be able to capture and analyze information in the network session in real time at wire-speed, which is not a trivial feat to accomplish. It must be efficient enough to analyze all network traffic, all without introducing a delay into the flow of packets across the network.

As organizations further leverage networking and the internet, the speed of network connections is increasing. Additionally, when moving beyond the network perimeter to positions inside the enterprise network, performance becomes even more critical. It is crucial to select a network DLP solution that can meet an organization's connection requirements today—one that is built upon an architecture that will scale to future bandwidth requirements that are likely to be seen as network collaboration continues.

Unfortunately, most first generation network DLP solutions on the market struggle above 100 Mbps of bandwidth. It is important to note that many solution providers will attempt to disguise their poor performance by providing the performance of the fastest supported network interface card, versus the supported throughput their analysis engine is capable of delivering. Many solutions claim to support a gigabit Ethernet interface, however their analysis engine only supports 100Mbps or less of bandwidth. It is critical to know the performance capabilities of the analysis engine, as once that capability is met the solution will either begin to fail, drop traffic, or sample. Failure tends to be easy to detect, as the system no longer functions. However, dropped traffic and sampling are often not advertised, but always lead to critical issues. Dropping traffic is when a system discards, and therefore does not analyze, any traffic over the threshold the analysis engine can handle. Sampling, is similar, except the system attempts to be systematic about selecting which sessions it analyzes and which sessions it does not. In any case, if the network DLP solution fails, drops traffic, or samples, the organization is incredibly exposed to the risk of data leaving the network undetected (a false negative) and the risk of finding out about an unauthorized disclosure of sensitive information from the news media or regulators becomes reality.

Supported zones of control

Organizations have deployed network security controls beyond the internet perimeter across the enterprise network, creating internal zones of control within the network. These internal perimeters, where technologies like firewalls and intrusion prevention have been deployed, are also likely locations where network DLP also needs to be deployed.

Examples of internal zones of control typically include:

- Edge of the data center controlling information leaving to endpoints and departmental servers;
- Inside extranet connections controlling access from business partners;
- Inside network connection to outsourcing providers controlling information extracted (versus systems access);
- Inside the VPN concentrator controlling information leaving to remote endpoints and employee home computers;
- Between divisions of companies (e.g., manufacturing division to financing division) or between division and enterprise backbone;
- Protecting secure networks (e.g., HR or engineering network).

Support for internal zones of control requires additional capabilities beyond the requirements at the internet gateway. First, performance is critical. Internal networks are typically significantly higher performance, further stressing the importance of wire-speed performance (above). Additionally, different protocols are typically seen inside the network than at the internet gateway. File sharing between servers and database traffic are examples of network protocols seen inside the enterprise and typically not at the internet gateway. In order to control this important internal traffic, the network DLP solution must understand these internal protocols.

Fidelis XPS was designed from the beginning to support multi-gigabit-speed networks without sampling or dropping traffic. Unlike the first generation of DLP solutions that only support 100Mbps of analysis, Fidelis XPS provides appliances with support for up to 10Gbps networks with up to 2.5Gbps of analysis in a single appliance *without sampling or dropping traffic*.

Fidelis XPS is also the only network DLP solution to include a sensor designed for internal networks. Fidelis XPS Internal provides the only support in a network DLP solution for internal protocols like database traffic and file sharing. The ability to provide these critical internal zones of control enables organizations to get more value with Fidelis XPS than other network DLP solutions.

Enterprise architecture-friendly

Practically every information technology (IT) purchase has impact, either positive or negative, on other assets in the enterprise. Minimizing or eliminating negative impact, while creating synergies with existing deployed assets, enables any solution to add more value than it does as a standalone system—so the network DLP solution should support that same goal.

Minimizing operational impact

The network DLP solution must be able to be deployed in a manner consistent with enterprise architecture standards and be minimally invasive to the enterprise architecture in order to lower the solution's operational impact and therefore maintain the lowest possible total cost-of-ownership. The implementation should not negatively impact network performance, add additional points of failure, and/or require desktop or server reconfiguration.

Adding value to existing IT assets

While minimizing impact to exist assets is often enough for deployment approval, adding value to existing IT assets—particularly other security assets—should be a goal of the network DLP solution. The most impactful way to add value to other systems is integration and information sharing in order to simplify management of risk. Network DLP, while critical to protecting an organization's sensitive information, is only one component of the security infrastructure. It is important to integrate the knowledge gained from network DLP solutions into other security and risk management solutions. Unfortunately, many DLP providers aren't interested in sharing information with third-party products, as many of the DLP "suite" providers are motivated by the sale of their own security solutions, regardless of the quality of the solution.

Data leakage prevention should not be a silo. Network DLP provides the content-aware mechanism to understand what information is leaving the enterprise or crossing internal zones of control. However, to maximize risk reduction this information should be correlated with other information present in the security and risk management infrastructure including network security/intrusion prevention and identity and access management systems. Security Information and Event Management (SIEM) tools are a popular place for aggregating this information.

Additionally, network DLP provides information on how the organization communicates and where sensitive information is flowing. This is an important intersection with other security technologies, especially encryption and systems management. DLP helps an organization understand compliance with their encryption policies and encryption can provide the capability to send protected information securely to approved parties.

Fidelis XPS provides high-performance network appliances designed to deploy within your infrastructure with no impact. As the only network DLP to both prevent without a proxy server and prevent when installed out-of-band, Fidelis XPS can stop data leaks without enterprise architecture impact.

Fidelis XPS also provides robust support for integration into other security and IT management technologies. The Fidelis XPS CommandPost™ management console supports a variety of open standards to externalize alert information including syslog, SNMP, and SMTP. Additionally, specific integration has been developed with: ArcSight, for security information and event management; Verdasys, for best-of-breed endpoint data leakage prevention; and IBM ISS SiteProtector™, for correlating threat mitigation and information protection events. Finally, Fidelis XPS Mail provides robust integration into e-mail encryption solutions.

Robust architecture

The last aspect of a network DLP solution that should be analyzed is the robustness of the architecture and its ability to provide additional value in the future. The risks associated with unauthorized disclosure of sensitive information are likely to continue to be present for decades to come. It is critical that the network DLP solution be able to adapt to the changing threat landscape and IT infrastructure. While predicting the future is often challenging, you want to purchase the network DLP solution that will meet your requirements both today and into the future.

To help, here are some key areas of the underlying technology to evaluate. A review of the solution's underlying technology should explore:

- ***Is a platform technology owned by the network DLP provider or is it an integration of third party components?*** An integration of third party components is less likely to be able to meet future requirements, as the network DLP provider is beholden to other technology companies for key components of their infrastructure, giving them less control over their own product's future. Determine what components are not owned by the provider and how critical they are to the future of the product. Areas like network traffic and content analysis are critical to the network DLP solution's success. Lack of control of the core technology architecture components presents significant future risk.
- ***Is the network DLP system likely to be able to meet future performance requirements?*** As mentioned before, future performance requirements are also growing each year as networks become more and more valuable and bandwidth lowers in price. It is likely that network performance will grow that fast, or even faster over the next few years. A solution that doesn't have the headroom today isn't likely to have it in the future either.

Fidelis XPS is built on Fidelis Security Systems' patented deep session inspection™ platform. As Fidelis Security Systems owns all of the product's intellectual property, from network sniffing through content analysis, customers are ensured that Fidelis XPS has the ability to continue to meet their requirements for network DLP. Unlike other providers who purchase document cracking and content analysis components from third parties, often designed for database indexing or search applications, Fidelis XPS is built from the ground-up to provide a robust network security solution to the network data leakage problem. This wholly-owned platform processes all network traffic in real-time in memory, delivering the highest performance of any network DLP solution with support for multi-gigabit networks—giving it the ability to scale for today's network demands and even further in the future.

Appendix A: RFI Checklist

This checklist identifies baseline requirements for a network DLP solution. It describes the general requirements and objectives for a technology solution capable of preventing a leakage of classified/sensitive/protected information from leaving the network. For a network DLP solution to be effective, it must provide the following critical capabilities:

Visibility

- Conducts session-level (not packet-level) inspection of network traffic across all 65,535 network ports in a single layer 2 network appliance.
- Provides visibility into the protocols, channels and applications in use on the network including, at a minimum, the following channels: SMTP, POP, IMAP, FTP, HTTP, P2P, IRC, AIM, Jabber, Yahoo Messenger, MSN Messenger, Telnet, CIFS/SAMBA/SMB, DB2, Oracle, LDAP, web services, webmail providers, and webmail platforms.
- Performs port-independent protocol inspection that inspects all 65,535 ports for all supported protocols.
- Extracts enterprise human-readable content and related meta-data contained in the session and any attachments and compressed files for analysis.
- Inspects SSL encrypted sessions via integration to an ICAP-enabled proxy server.
- Provides appliance that understand database and file sharing applications to establish internal zones of control.

Performance

- Operates at network wire-speed in a single appliance, even on multi-gigabit networks, without sampling or dropping sessions.
- Supports at least twice current forecasted enterprise bandwidth requirements.

Control

- Prevents unauthorized transfers across all 65,535 ports.
- Prevents out-of-band via TCP reset packets when installed out-of-band.
- Prevents by dropping packets when deployed inline.
- Able to prevent without third party proxy or gateway products.
- Quarantines SMTP e-mails that violate policy.
- Redirects SMTP e-mails to an encryption server for secure delivery.

Content Analysis

- Provides multiple sophisticated content analysis technologies to detect sensitive and/or protected information.
- Provides ability to combine multiple content recognition methods in a single rule.
- Able to go beyond exact matching for ease-of-deployment and solution scalability and provides analyzers to profile or describe digital assets without the need for intensive registration and maintenance processes to identify sensitive information.

Architecture

- Changes to configuration of servers or desktops not required.
- Analyzes traffic on all ports without introducing a delay into the flow of packets across the network.
- Able to analyze network traffic, decode documents, or perform content analysis without the use of major third party components.
- Built on an architecture that will scale to future bandwidth requirements.
- Externalizes alert information via standard including syslog, SNMP and SMTP.



About Fidelis Security Systems

Since 2002, Fidelis Security Systems has given global enterprises the power to prevent data leakage while solving their biggest data leakage challenges—safeguarding intellectual property and identity information, complying with government and industry privacy regulations, and enabling visibility into and control over their networks. Built on a patented deep session inspection™ platform, the Fidelis Extrusion Prevention System®, Fidelis XPS™, is the industry's only next-generation data leakage prevention solution with the power to deliver comprehensive prevention, complete control, and the lowest total cost-of-ownership to stop data leakage on multi gigabit-speed networks. To learn more about why organizations choose Fidelis XPS to protect their brand, intellectual property, and resources, visit www.FidelisSecurity.com.

FIDELIS SECURITY SYSTEMS, FIDELIS EXTRUSION PREVENTION SYSTEM, FIDELIS XPS, the FIDELIS SECURITY SYSTEMS logo, and/or other FIDELIS SECURITY SYSTEMS products referenced herein are trademarks of Fidelis Security Systems, Inc. Copying, use or distribution of any material contained herein is expressly prohibited. *Copyright © 2009 Fidelis Security Systems, Inc. All rights reserved.*