Preventing Data Leaks At The Firewall

A Simple, Cost-Effective Way To Stop data leaks within your network

Executive Summary

Numerous examples of accidental and deliberate data leakage continue to make headlines, and data leak prevention (DLP) technologies are being touted as a panacea. Unfortunately, given the scope, size, and distributed nature of most enterprise datasets, just discovering where the data is and who owns it is a challenge, and many DLP projects are proving slow to progress. Furthermore, given the absence of control over applications in most enterprises, it isn't clear that DLP technologies would have helped in several of the recent cases (US Army, Pfizer, etc.). Some organizations, for business model or cultural reasons, will have to go through the effort of large-scale DLP implementations. For most organizations, however, controlling the applications most often used to leak sensitive data and stopping unauthorized transmission of files, credit card and social security numbers and their ilk would be acceptable. Many organisations would like to monitor the files uploaded by their employee and then have controls to block such uploads. Exerting that control at trust boundaries is ideal – whether the demarcation point is between inside and outside the firewall sits on the perfect spot, seeing all traffic.

Existing firewalls not good enough

Unfortunately, legacy port- and protocol-based firewalls can't do anything about any of this – being ignorant of applications, users and content. Existing firewalls would allow you to setup policies based on the port, with some going to an extent to allow based on users. To do this correctly, enterprises should first achieve a degree of control over applications – thus limiting the avenues of data leakage. Second, organizations need to scan the applications they do want on their networks, for confidential data. Third, organizations should be able to understand which users are initiating these application transactions.

GajShield's UPTM firewall belong to the next generation firewalls which bring contextSense into the content flowing over the applications. ContextSense firewall are able to indentify the sensitivity of data based on the context of the information flowing. For example, sales quote sent to customers should be allowed whereas to another id should be blocked. Current firewalls are unable to identify the context of information and hence unable to validate the criticality of information.

GajShield appliances incorporate six senses that enable enterprise customers to incorporate some of the most commonly needed DLP functionality at their network perimeter – easily, and without adding more appliances. GajShield's ApplicationSense, UserSense, ContentSense, NetworkSense, TimeSense and ContextSense, offers immediate relief to the most common data leakage pain allowing enterprises to complete their large scale DLP projects at their leisure.

Data Leakage Continues to Be a Problem For Enterprises

Large scale, public exposures of personal information remain far too common. Practically weekly, headlines declare tens of thousands of credit card numbers leaking out of retailers, or social security numbers leaking out of government agencies, health care organizations, or employers. A recent example (December, 2008) showed a misconfigured and prohibited peer-to-peer file sharing application putting a database of 24,000 US Army soldiers' personal information in the public domain. This is similar to the Walter Reed Medical Center breach in June 2008, or the Pfizer incident from 2007 – all three involved data leaking through the perimeter via an application expressly prohibited by policy.

GajShield Infotech Pvt. Ltd, 4, Peninsula Centre, S.S. Rao Road, Parel, Mumbai – 400012
Tel : 91-22-66607450, Fax : 91-22-66607450, Email : info@gajshield.com, Web : http://www.gajshield.com

DLP Technology is Cumbersome, Incomplete, and For Many, Overkill…And Don't Forget Expensive

Data leak prevention (DLP) technology has captured the attention of many IT organizations, with a promise to help organizations manage their confidential data. Project scope, for these technology providers, however, is a problem. Questions of access control, reporting, data classification, data at-rest vs. data in-transit, data ownership, desktop agents, server agents, and encryption have slowed DLP projects to a crawl in many organizations. Venture capitalists funded many data leakage prevention vendors, many of which have been acquired by larger security companies, who have further expanded the scope of an already unwieldy offering. Some of these vendors are now marketing data *loss* prevention, which incorporates practically the entire information security function (and even includes elements of storage management!). Needless to say, this broadened scope is beneficial, but adds complexity, time, and expense – both in hard costs and in staff time. Oddly enough, many of the recent breaches caused by unauthorized and misconfigured peer-to-peer file sharing applications wouldn't have been prevented by the typical implementation of DLP technologies on the market today – because control of applications isn't addressed.

For 90% of Enterprises, knowing what is uploaded and controls to block it, keeping Social Security and Credit Card Numbers From Leaking Would Be Enough

For a few highly intellectual property-dependent organizations, implementing a long-term, comprehensive DLP project – which should ultimately include data discovery, classification, and cataloging – is appropriate. But for the remaining 90% of enterprises out there, stopping a couple classes of confidential data (e.g., uploaded files, social security numbers and credit card numbers) at the trust boundary would be a great start. This would avoid the expensive and embarrassing public exposure of employee and/or customer personal data.

The Perimeter Is Key, But Legacy Security Technology Can't Help

If enterprises could control the flow of confidential data at the trust boundary, it would stop a large percentage of incidents that regularly make the news. Unfortunately, the legacy security infrastructure at most enterprise perimeters is poorly equipped to offer this functionality. Most firewalls sit in a great position to help – they demarcate the trust boundary, they see all traffic, and they exert policy control (i.e., they can block traffic). But legacy firewalls don't understand content, don't understand applications and context, can't see inside SSL-encrypted traffic, and have no understanding of users. In fact, if it isn't source or destination IP address, source or destination port, or network protocol, firewalls don't understand it. Other firewall "helpers" (e.g., intrusion prevention systems, web proxies, URL filtering devices) only see a portion of the traffic, don't sit in-line, and/or have limited application and content understanding.

First - Block "Bad" Applications

Examining most of the recent incidents, the first thing enterprises need to do is get control over which applications are running on the network. Every organization has a different view of desirable and undesirable applications. Each enterprise needs to look at applications from both benefit and risk perspectives. On the benefit side, an application might help an employee do their job better, faster, or cheaper, or improve customer relations, or make the workplace more pleasant. On the risk side, applications may harbor vulnerabilities, carry malware, be prone to misuse, or – particularly relevant to this discussion – transfer files. In some cases, organizations want to enable social networking applications for cultural reasons, or for business reasons – or block them for security reasons. Either way, the first thing to do with regard to stopping confidential data leakage is to identify which applications are moving across the network – regardless of whatever evasive tactic the application employs, and block undesirable applications (thus limiting the avenues through which confidential data can flow). In order to do perform this control effectively, the device needs to "see" all traffic.

Second - Scan "Good" Applications. Including SSL.

The second thing to do is scan desirable applications for confidential data leakage. Once an organization has settled on the applications it wants on its network, the next step is to scan applications for confidential data leakage – including SSL-encrypted application traffic and compressed content. As an aside, any applications that use proprietary encryption (e.g.,Skype) should be very closely evaluated, because if allowed, they cannot be scanned. More specifically, the scanning capability should be simple to enact in policy, and adjustable in sensitivity, to allow normal appropriate transactions without triggering response – yet still detect abnormalities.

Thirdly - Know Users, not just IP addresses

The third thing to do is to bring users into the picture. Understanding which users are using applications, and which are engaged in moving particular classes of content has two benefits – actionable visibility, and refined policy. The most efficient way to do this is to tie into the enterprise directory (identities and groups are already there). Often, when an organization hears that they've had a leak, the first thing they ask is, "who leaked it?" Having the ability to understand users – not just IP addresses – gives the granularity enterprises need to guide specific users about policy, and take more serious action if warranted. The second benefit understanding users offers is the ability to incorporate that understanding into policy – so certain users might be enabled to use certain classes of applications, and other users might not. This empowers enterprises to further compartmentalize and contain risk.

Finally – Bring Context to information flow

Every information does not work independentily. It has a context. For example, when a sales quote is sent to a customer, the information send (by which user, which application, what data) has a different criticality level than when the same data is sent to different users. Next generation firewalls, like GajShield, would understand this context and help in creating policies to either allow or block them. With its unified approach to UserSense, Application Sense, ContentSense, NetworkSense and TimeSense it is easily able to identify the context and hence the criticality.

In summary – if IT staffs know the application, the user, the content and the context (i.e., whether or not the traffic contains confidential information), they can act – block or alert – quickly, and archive appropriately, without sifting through dozens of log files.

GajShield UPTM Includes DLP Functionality In The Firewall

GajShield offers enterprises a unique approach – visibility and control over applications, the ability to scan application content, and build a context and control of users and groups. GajShield's UPTM appliances incorporates 6 key senses. ApplicationSense, ContentSense, UserSense, TimeSense, NetworkSense and more importantly ContextSense give organizations business-relevant control over applications.

ApplicationSense Classifies Applications. ApplicationSense technology identifies applications regardless of port, protocol, encryption, or evasive tactic. It gives enterprises visibility and policy control over actual applications, not just ports.
ContentSense Identifies Content – Including Confidential Content. ContentSense technology incorporates 3 key content security elements – confidential data (DLP functionality), threat prevention, and a URL filtering capability. The data filtering feature in GajShield makes implementing DLP functionality in the firewall simple. Adding a policy object that scans application traffic is a matter of assigning the data filtering profile to the policy, determining what sort of data to scan for. Enterprises can also use the regular expression capability built into the data filtering feature to create custom patterns. More importantly, ContentSense keeps track of all uploaded data and archives it. This gives an enterprise insight of what is being uploaded even for content where no policy has been set.

UserSense Integrates With Enterprise Directories. UserSense technology integrates GajShield's UPTM with enterprises' Active Directory implementations. Meaning that the single policy engine governing application and content security also has the ability to refine that policy with the user and group definitions already used in theenterprise.

TimeSense – identifies the time when any information is sent. Some information may have time sensitivity. For example, you may not want your audited reports to be published or sent before it is publically announced.

NetworkSense – Organisation would not want that critical data travel through public networks. For example, when a mobile device is uploading data within the network would need to go through multiple checks of malware detection than data coming from within private networks.

ContextSense – Every information has criticality based on the context over which it has been sent. GajShield's ContextSense helps in indentifying the context with the help of the above five senses to categories the criticality of data. For example, when a user is allowed to connect to a POP3 server, he could possibily download mails for any user, but a contextsensitive firewall will identify the user using UserSense, the application (POP3) using ApplicationSense, the id it is using to identify to the POP3 server using ContentSense and create a context using the ContextSense engine which enables an organisation to either block or allow the id to connect thus ensuring that only valid users connect to their POP3 ids.

It's Time To Fix The Firewall

Comprehensive DLP is a worthwhile pursuit, but is complex, expensive, and takes a while to implement. In the meantime, the legacy firewall is sitting in a prime spot to help out – but due to its inability to see applications, users, content and context – can't do a thing. Firewalls should be able to:

- First, block undesirable applications.
- Then, scan allowed applications for confidential information using context analysis.
- See and manage policy by users and groups, not IP address

Fortunately, GajShield, with its identification technologies, delivers this in a highperformance firewall platform. The application visibility and control, coupled with the data filtering feature found in GajShield's UPTM can enable simple, high-performance DLP controls at the enterprise perimeter – which would have stopped the data leaks in several recent, highly publicized incidents. And that should free up staff to work on that data *loss* prevention project.