# Application Firewalls

As businesses place more applications on the web, they expose more of their sensitive customer data to hackers. Browser-based applications tunnel through the entire security perimeter of an organization, giving users unprecedented access to internal systems. It's little wonder that the majority of attacks today target the application directly.

For most organizations, the web application has in itself become the security perimeter, and the only way to ensure the security of those applications is what's known as an application firewall.

However, application firewalls can only be effective if they are tailored to match closely to an application. Poorly-tailored gateways will inevitably block legitimate user or customer traffic or let in hackers.

The F5 Application Security Manager (ASM), an application firewall, is a new class of device that protects applications from hackers and other malicious attacks. It enforces granular security policies to protect web applications as well as confidential information from both random and targeted application security attacks. And thanks to breakthrough technology that automatically generates an extremely accurate model of all legitimate user interaction with an application, ASM is able to filter all application requests and deny anything that is not legitimate user activity.

## Challenge

Today, every aspect of business is migrating to the web. Unfortunately, every time a new web-based application is created, back-end systems previously sheltered from direct access are now connected to the Internet—and potentially the world. The result is that a company's critical data is exposed to an external attack.

Meanwhile, hackers are finding new ways to penetrate traditional defenses. According to a recent CSI/FBI study[i], 68% of respondents reported 100% of security-related losses came from company system penetration from the outside in 2006, despite the fact that 98% of the respondents had firewalls in place and 69% implement IPS technologies. Reported financial losses from these attacks, including system penetration, misuse of web applications, web site defacement, theft of proprietary information, and Denial of Service totaled more than $52 million among the 313 company respondents.

Traditional firewalls, which have historically done an excellent job preventing outsiders from accessing company networks, are no longer sufficient. Network firewalls traditionally offer little or no protection for data in the application layer because they live in the network and transport layers and are solely focused on ACL and port management. Firewalls have to leave application ports open for services such as HTTP in order for requests to make it on the network. That's where their expertise stops.

Recently, the traditional firewall vendors have begun touting "application-layer security," incorporating functionality from Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS) into their products. Unfortunately these solutions have proven ineffective for two fundamental reasons:

- **They rely on static attack signatures or other patterns of abnormal user behavior.** This leaves systems exposed to new types of attacks ("zero-day attacks"), and threats cloaked as normal traffic. More importantly, it leaves them blind to targeted attacks exploiting the specific vulnerabilities of an application, for which there is no generic pattern.
- **They operate at the network layers, not the application layers.** Despite their claims, many network security products are inherently limited because of the

information they are able to interpret. Firewalls and IPS systems look at packets on the wire, not entire requests and user session data, so they lack the application-specific knowledge to tell a good request from a bad one. In fact, many can't even look into simple SSL encryption.

Accordingly, businesses today are being forced to build a high degree of security into the applications themselves. Popular tools such as application scanners help identify obvious holes, but ultimately it's a labor-intensive job to scan and then patch them all. More problematically, scanning can never reveal all application vulnerabilities; there are simply too many parameters to check and too many possible entry points. Developers can patch thousands of holes, but a hacker only needs to find one to inflict major damage.

The ideal solution, therefore, would be to offload the security function to a network appliance that has enough application-specific knowledge to filter out malicious requests. This protection would work like a firewall, but would be based on the application's specific logic, not generic traffic patterns. This type of device would know exactly what the application's traffic should look like, and block anything else. In response to this need, a new class of security solutions called *application firewalls* has emerged.

## Solution

F5's Application Security Manager is a unique application firewall that provides comprehensive protection for web, SOA, and XML applications against application layer attacks, not only from the exploitation of known web application and infrastructure vulnerabilities, but also more malicious, targeted attacks.

Specifically, ASM can stop attacks that no other solution on the market can stop. Take, for instance, two common hacker techniques that pass straight through today's security solutions, even those solutions that claim to feature "application security" or "content-aware" blocking:

- **Hackers enter as one user, then change their ID or escalate privileges once they are past the authentication "gate."**
  The most complex form of this (dynamic parameter tampering) is invisible to nearly every solution on the market.
- **Hackers change or bookmark the URL of a web application to enter areas which should be restricted.**
  Sometimes this is as simple as changing a URL from .../webapp/user to ...webapp/admin. More often, the path is convoluted or even buried within an application. Users who are familiar with the application, however, can often guess or detect where to go.

Only ASM can protect against these breaches because it is the first to have a comprehensive understanding of the user interaction with the firewall based on real-life traffic. This means it not only has a very granular understanding of legal activities, but also a firm understanding of the user context (or state) at any given time, enabling ASM to build and implement a security policy specific to the application it is protecting.

Additionally, ASM's accurate security policy is based on a proprietary model (called the Real Traffic Policy Builder) which combines automatic analysis of live traffic content (both requests and responses) with iterative adjustments based on automated, real-life traffic analysis.

### Real Traffic Policy Builder

The best way to enforce a security policy is to have a detailed model (or policy) of the ways users interact with the application. Once you have defined what is legal, all other activities can be declared illegal. An accurate model of the user activity, then, is critical to security enforcement. Without this, the policy with either permits attacks or—more likely—blocks users who are attempting to perform legitimate activities.

F5's Real Traffic Policy Builder is a logical representation of the interaction between a legal user and the web application, built from live traffic as it is requested from and delivered to the end-user. For each web page presented to the user, the model describes the structure of the HTTP requests that are generated by the client-side source code of the web page and the authorized transitions to other web pages and the responses served from the web server. The representation of the flow that transfers the user from the browsed (source) web page to another (target) page includes:

- The current web page
- The target web page
- The names of the parameters that can or must appear in the request
- The characteristics of the values allowed for each parameter

The Real Traffic Policy Builder represents a breakthrough in application modeling because previous approaches only created models of user requests based on scanning user traffic. F5's model automatically builds a policy for the entire application based on live traffic and data from the application, mapping the flow or total pattern of user interaction with the web site. This ability to map the application in addition to tracking traffic patterns is far more accurate in modeling user interactions than any other previous methods. The benefits of this model are threefold:

- **Knowledge of state** – Only the Real Traffic Policy Builder tracks which pages a user is coming from and the specific permissions associated with that context. A request which is perfectly "legal" within the context of one page might be inappropriate for a user on another page.
- **Bidirectional** – Only the Real Traffic Policy Builder looks at server responses to the client as well as client requests to the server. This is essential to verify that the user hasn't attempted to tamper with the credentials sent to him in his response.
- **Granularity** – Only the Real Traffic Policy Builder creates a complete logical rendering of the transitions between every page, including every object, every parameter of each object, and every legal value within each object parameter.

Building the security policy on the basis of this model allows ASM to verify that the user interaction with the web application follows the web application design and enables it to block any attempts that vary from it. In other words, the Real Traffic Policy Builder turns the problem of application security into a problem of *session-based Application Delivery Security*, something that firewalls and IPS devices have been, and continue to be, ineffective at doing for so many years.

### Comprehensive Protection Against All External Threats

In order to offer comprehensive protection for the enterprise application delivery infrastructure, ASM combines robust application-layer filtering with best-in-class network and encryption technology for a complete web security solution:

**ASM Features**

**XML/SOA Protection**

- Schema validation
- Parser protection (XML Bombs, Recursion Attacks)
- XPATH injection
- RSS/Atom feed injection
- XML islands

**Validated Application Security Policies**

- Pre-built security policies for the most common applications
- Tuning of existing policy templates
- Building and exporting customer policy templates for any application

**Security Attack Filters Protect Against Random Attacks**

- ASM Live Update signature update service
- Known worms and vulnerabilities
- Requests for restricted object and file types
- Other known exploits

**Policy-based Security Protection Against Targeted Attacks**

- Policy Evasion Detection Engine
- Manipulation of invalidated input
- Broken access control (Forceful Browsing)
- Buffer overflow
- Cross-site scripting
- SQL/OS injection
- Cookie poisoning

**Content Scrubbing**

- Identity theft protection for web servers
- Ensure customer information is never served on a web page
- Configurable to scrub any identifiable information such as:
- Social security numbers
- Credit card numbers
- Account numbers
- Patient health data
- Phone numbers

**Network Security Services**

- SSL accelerator
- IP/Port filtering
- Reverse proxy
- Key management and failover handling
- SSL termination and re-encryption to web servers

**Cloaking**

- Prevent OS and web server fingerprinting
- Conceal any HTTP error messages from users
- Remove application error messages from pages sent to users
- Prevent leakage of server code

Only ASM combines all these functions into one simple-to-manage device for complete Application Delivery Security.

## Deployment Options

When deploying ASM in an enterprise or large government environment, the need for security of course must be managed as part of a broader risk profile. Some applications will require immediate and strict policy enforcement, while others require a more rapid deployment.

ASM can be used in a variety of security postures, from a basic intrusion protection *shield* (requiring just minutes of set-up time) to a complete positive-security *blanket* (utilizing the existing application policy templates). The ability to be configured at these different security levels provides security administrators with the flexibility to ensure enterprise-wide protection immediately, without compromising the security for their most sensitive applications. ASM provides application security from day one, with full flexibility to tune the application policy as needed.

## Enterprise Class System Architecture

ASM's multi-layer system architecture is based on a hardened security appliance designed to meet all of an enterprise's demands for infrastructure security, including:

- Negligible latency (less than 1 ms) and high throughput.
- *Scalable architecture*—additional units can be added to handle larger traffic volumes.
- *High Availability*—units can be configured for hot, stateful failover between yoked pairs of servers (an active server and a standby server). In the unlikely event of a server failure, all session data is preserved and failover to a backup unit is invisible to the user.
- Zero-fault configuration, easy deployment and maintenance—ASM is composed of optimized and pre-configured appliances that effectively address configuration, deployment, and maintenance issues.
- Central and secure management.
- Easy integration with enterprise security information management or management framework systems.

## ASM Business Benefits

### Closing the Door on Web Attacks

Obviously the most important benefit of the ASM solution is that it eliminates the risk of cyber-attacks through a company's web applications. As more systems are opened to web traffic, more and more sensitive customer data is exposed to threats which current security systems cannot prevent. And once hackers are in, the costs to your business can be staggering, costs that do not take into account the increased insurance expenditures and the legal responsibilities that accrue with deficient security.

### Identity Theft and other Regulatory Compliance

Across industries, new regulations such as the Basel Accords, HIPAA, California's SB 1386 and a host of other national and trans-national regulations are making the security of personal customer data an imperative. Currently, web applications are the main entry point for hackers seeking customer data. Application-layer attacks that companies know about (certainly only a small portion of the total) cost them hundreds of millions of dollars per year. ASM is an absolute necessity for any company with sensitive customer information.

### Improved Time to Market

In addition to preventing hacks, ASM can actually improve the development cycle for new applications. Right now, the deployment of new applications is hampered by the “scan-and-fix” cycle of application security scanning tools. Code is written, scanned by one of several off-the-shelf products, and then sent back to the developers to be rewritten. Not only is this costly and time-consuming, it is ineffective, as scanners can only detect a limited set of known security breaches. With a product like ASM, the development team can focus on rapid development of new applications and functionality, knowing that their code will sit behind a powerful security perimeter.

### Plug and Protect

The integrated ASM solution delivers on the promise of a "plug and protect" appliance. ASM is delivered as a network device, which can be easily nested in a company's web infrastructure. Once installed, the proprietary, automated learning mechanism quickly and accurately builds security policies tailored to the unique specifications of the applications it protects. Policy management and configuration are minimal, as ASM automatically generates recommendations, rather than waiting for manual configuration.

## Easily Quantified Return on Investment

The benefits of an application-layer security solution such as ASM can be easily quantified. ASM will slash costs related to security enforcement, attack damage, and damage response. Consider the savings associated with the following:

### No Attack Incidents

In addition to the costs of the attacks themselves (stolen funds, lost revenue) companies have extensive costs associated with responding to the attack and repairing the damage. This response is not limited to the IT department, as it can involve public relations, litigation, and even regulatory costs.

### No Code Rewrites

Without adequate protection around their application, application developers are forced to scour their applications for individual security holes, and hard-code plugs for those holes. Application scanners can detect some of these, but rigorous code review and rewriting is always necessary. Once ASM is in place, developers can focus on the rapid deployment of new applications and new functionality.

**No Reactive Patching**

Knowing that applications are often open to direct attacks, IT managers constantly monitor sites and company announcements to immediately install the latest patches. As mentioned above, a secure application firewall ensures that only legal activity is permitted on an application, reducing the reliance on patches.

**No False Positives**

Blocking customers from their accounts while they are conducting legal transactions is a good way to lose those customers. Without an accurate model of legal activity (such as the Application Flow Model), companies face the tough choice of relaxing security protocols and letting in hackers, or tightening them and blocking customers.

[i] CSI/FBI Computer Crime and Security Survey: http://www.gocsi.com/forms/fbi/csi_fbi_survey.jhtml