Distributing Security and Optimization in the Application Delivery 2.0 Era

By Dr. Jim Metzler

Sponsored by:

Produced by:

Introduction

While ensuring acceptable application delivery has always been important, it historically was not a top of mind issue for the majority of IT organizations. That changed several years ago when IT organizations began to develop a focus on it. As part of this focus, most IT organizations deployed a first generation of solutions that were intended to protect the organization from a growing number of security attacks, to mitigate the impact of chatty protocols such as CIFS (Common Internet File System) and to offload computationally intensive processing (e.g., TCP termination and multiplexing) from servers. Throughout this white paper, the application delivery challenges and solutions of this era will be referred to as Application Delivery 1.0.

While many IT organizations are making progress relative to implementing solutions that respond to the challenges of the Application Delivery 1.0 era, a new generation of challenges and solutions is emerging. Throughout this white paper, this new generation of application delivery challenges and solutions will be referred to as Application Delivery 2.0.

One of the challenges associated with Application Delivery 2.0 is the shifting emphasis and growing sophistication of cyber crime. For example, at the beginning of the Application Delivery 1.0 era the majority of security attacks were caused by individual hackers whose primary goal was to gain notoriety for themselves. While those types of attacks still occur, either organized crime or a rogue nation is behind today's biggest security attacks and their motivation is to make money or to cause significant harm to an organization for political or ideological reasons.

The ongoing adoption of cloud computing on the part of most IT organizations results in an increased amount of traffic transiting the Internet to access highly visible sites that are attractive to hackers. The adoption of cloud computing also causes some performance and security problems that are not addressed by Application Delivery 1.0 solutions such as WAN optimization controllers (WOCs), application delivery controllers (ADCs) and network firewalls.

Relative to application delivery, cloud computing is a double-edged sword. The performance and security issues associated with cloud computing are some of the Application Delivery 2.0 era's most significant challenges. However, the judicious use of cloud computing services to distribute critical optimization and security functionality throughout the Internet is one of the key characteristics of the solutions that are associated with Application Delivery 2.0.

As will be discussed in this white paper, an example of a cloud computing service that solves some of the Application Delivery 2.0 performance-related challenges is an application delivery service that optimizes the performance of the Internet. As will also be discussed, the use of a Web application firewall service is an example of a cloud computing service that solves some of the Application Delivery 2.0 security-related challenges. One of the advantages of using a Web application firewall service is that it

distributes security functionality out close to the source of security attacks and hence can prevent many attacks from reaching the organization.

Application Delivery 1.0

This section of the white paper will discuss some of the key challenges and solutions that are associated with Application Delivery 1.0. This section will also discuss the limitations of a perimeter-based approach to security.

Challenges and Solutions

Several years ago most IT organizations began to develop a systematic focus on ensuring acceptable application delivery. The goal of this new focus was to ensure that the applications that an enterprise uses:

- Incorporate appropriate levels of security
- Exhibit acceptable performance
- Provide high availability
- Are cost effective

Some of the first generation of applications delivery challenges included the fact that:

- The typical application development process did not allow for adequate code review and vulnerability testing. This resulted in a growing number of security vulnerabilities.
- Many IT organizations had begun to implement Web-based applications that utilize protocols such as HTTP. As explained below, the use of HTTP creates security vulnerabilities.
- The damages that result from cyber crime had been growing on an annual basis.
- IT organizations were significantly increasing their use of chatty protocols such as Common Internet File System (CIFS) as well as computationally intensive protocols such as Secure Socket Layer (SSL).

To respond to these application delivery challenges, IT organizations deployed a first generation of application delivery solutions. These solutions included:

- Traditional network firewalls that examine network traffic and block unauthorized access while permitting authorized communications.

- Intrusion detection systems (IDS) that passively watch packets of data from a monitoring port and compare the traffic to configured rules, and set off an alarm if they detects anything suspicious.

- Intrusion protection systems (IPS) that sit inline with traffic flows on a network and which can actively shut down attempted attacks.

- WAN optimization controllers (WOCs) that were designed in part to reduce the impact of chatty protocols.

- Application Delivery Controllers (ADCs) that were designed in part to offload from server farms the processing of protocols such as SSL.

Role of a Traditional Firewall: Protect the Perimeter

Roughly twenty years ago IT organizations began to implement the first generation of network firewalls, which were referred to as packet filters. These devices were placed at the perimeter of the organization with the hope that they would prevent malicious activities from causing harm to the organization.

Today most network firewalls are based on stateful inspection. A stateful firewall holds in memory attributes of each connection. These attributes include such details as the IP addresses and ports involved in the connection and the sequence numbers of the packets traversing the connection. One of the weaknesses associated with network firewalls is that they are typically configured to open up ports 80 and 443 in order to allow passage of all HTTP and SSL traffic. Given that ports 80 and 443 are generally configured to be open, this form of perimeter defense is porous at best.

Whereas network firewalls are focused on parameters such as IP address and port numbers, a more recent class of firewall, referred to as a Web application firewall, analyzes messages at layer 7 of the OSI model. Web application firewalls are typically deployed as a hardware appliance and they sit behind the network firewall and in front of the Web servers. They look for violations in the organization's established security policy. For example, the firewall may look for abnormal behavior, or signs of a known attack. It may also be configured to block specified content, such as certain websites or attempts to exploit known security vulnerabilities. Because of their ability to perform deep packet inspection at layer 7 of the OSI model, a Web application firewall provides a level of security that cannot be provided by a network firewall.

One of the challenges associated with a Web application firewall is capacity planning. Providing too much capacity wastes money, while providing too little capacity results in either degraded performance or failure. In addition, a Web application firewall shares a critical weakness with a network firewall. That weakness is that both devices try to protect the perimeter of the organization. That approach is fundamentally flawed for two reasons. One reason is that as companies do more joint development and implement supply chain management, customer relationship management, and other collaborative

systems to reach out to customers, suppliers and partners, the perimeter of the typical organization is becoming increasingly ill-defined. These shifting business models are enabled by the ongoing adoption of new distributed application architectures (e.g., Web-enabled applications, SOA/Web Services, and Cloud Computing) that result in an increasing amount of traffic that traverses multiple enterprises.
The second reason that attempting to protect the perimeter of the organization is a flawed approach to security is that, as will be explained below, even if an organization could deploy a non-porous firewall at every WAN boundary point, that would not protect them from the current generation of distributed denial of service (DDoS) attacks.

Application Delivery 2.0

This section of the white paper will discuss some of challenges associated with Application Delivery 2.0. This includes the shifting emphasis and growing sophistication of cyber crime, as well as the performance and security challenges that result from the adoption of cloud computing.

This section will also discuss some of the solutions that are associated with Application Delivery 2.0. In many cases these solutions are acquired as a service and are based on distributing optimization and security functionality around the Internet. In most cases, these solutions are complimentary to the onsite solutions that are associated with Application Delivery 1.0. Examples of Application Delivery 2.0 solutions include the use of Internet-based application delivery services and the use of Web application firewall services.

Challenges

The Changing Face of Cyber Crime

As noted, preventing security breaches was a key component of Application Delivery 1.0. However, over the last couple of years the financial impact of security attacks has increased dramatically. For example, McAfee recently published a report[1] based on a survey of 800 CIOs that was performed by Purdue University's Center for Education and Research in Information Assurance and Security. The report stated that, "Based on the survey findings McAfee conservatively estimates that the global damage from data loss to top one trillion dollars".

The McAfee Report also discussed some of the ways that cyber crime was maturing and becoming more sophisticated. The report stated that, "Credit card fraud and identity theft have moved into the so-called 'cash cow' phase of criminal strategy. In other words, it's a source of revenue, but there's not much room for growth, so criminals are looking for the new stars of their portfolios. And intellectual property has emerged as a favorite." Also included in the report was the observation that malware writers now have R&D and test departments.

[1] http://www.mcafee.com/us/about/press/corporate/2009/20090129_063500_j.html

The Center for Strategic and International Studies (CSIS) recently released a report[2] entitled "In the Crossfire – Critical Infrastructure in the Age of Cyber-War" that provided further evidence that cyber crime is becoming more sophisticated and costlier. That report included the results of a survey that was completed by six hundred IT and security executives in 14 countries. According to the report, "More than half of the executives surveyed (54 percent) said they had experienced 'Large-scale denial of service attacks by high level adversary like organized crime, terrorists or nation-state (e.g. like in Estonia and Georgia).' The same proportion said they had been subject to 'stealthy infiltration" of their network by such a high-level adversary "e.g. like GhostNet'—a large-scale spy ring featuring individualized malware attacks that enabled hackers to infiltrate, control, and download large amounts of data from computer networks belonging to non-profits, government departments and international organizations in dozens of countries." The executives that were surveyed estimated that 24 hours of down time from a major attack would cost their organization on average U.S. $6.3 million. Apart from cost, the executives were also concerned that a major attack would harm their organization's reputation and could result in the loss of personal information about their customers.

Another aspect of the changing nature of cyber crime is the sheer scale of the attacks. In January 2010 Arbor networks published their Fifth Annual Infrastructure Security Report[3]. According to that report, the peak rate of DDoS attacks has grown from 400 Mbps in 2001 to 49 Gbps in 2009. This is a growth of more than a factor of one hundred in less than a decade. The report also stated that, "We expect DDoS attack rates to continue to grow, but given that most enterprises are still connected to the Internet at speeds of one gigabit per second (Gbps) or less, any attack over one Gbps will be typically effective, and often trigger collateral damage to adjacent network or customer service elements as well."

Unfortunately, pulling together an attack of that scale is not that difficult in the current environment. For example, a botnet (also known as a zombie army) is a number of Internet computers that, although their owners are unaware of it, have been set up to forward transmissions to other computers on the Internet. In 2009 it was estimated that as many as 34 million computers in the US may currently be part of a botnet[4]. A hacker does not need to be very sophisticated in order to carry out an attack using a botnet. Yet another example of the changing nature of cyber crime is that a hacker that lacks the ability to create a botnet can simply rent one[5].

An Increase in Security Vulnerabilities

One of the reasons that the global damage due to cyber crime is so high is that organizations of all types are increasingly vulnerable to being attacked. Some of the

[2] http://csis.org/event/crossfire-critical-infrastructure-age-cyber-war
[3] http://www.marketwire.com/press-release/Arbor-Networks-Fifth-Annual-Infrastructure-Security-Report-Finds-Service-Application-1103590.htm
[4] http://blogs.zdnet.com/emergingtech/?p=1706
[5] http://news.cnet.com/8301-1009_3-10223716-83.html

reasons for this increase in security vulnerability are an extension of the security challenges that were part of Application Delivery 1.0. For example, IT organizations are continually increasing their deployment of Web-based applications that rely on protocols such as HTTP and SSL. As previously noted, the network firewalls that are deployed at the enterprise's perimeter with the intention of protecting the enterprise from a cyber attack are typically configured to allow passage of all HTTP and SSL traffic. Hence, a growing amount of traffic is passing through network firewalls without being analyzed.

The current movement on the part of most IT organizations to implement some form of cloud computing has also created security vulnerabilities. A recent report[6] on the risks associated with cloud computing highlighted the fact that by a wide margin, the primary concern that IT organizations have relative to cloud computing is the security and confidentiality of their data. This concern stems in part from the fact that the use of cloud computing in general, and public cloud computing in particular, results in more of an organization's data being accessible via a wide area network (WAN). Data that is accessible via a WAN is more vulnerable to security attacks than is data that is only accessible via a LAN.

The previously referenced Arbor Networks' report[7] alluded to the fact that an organization can be severely damaged even if they are not the intended target of an attack. What the report was referencing is the fact that many organizations suffer collateral damage when a network or organization in close proximity is attacked. This type of vulnerability is becoming increasingly common as IT organizations increase their utilization of public cloud computing solutions such as software as a service and infrastructure as a service. In almost all cases, these solutions are based on a multi-tenant environment. A side affect of having a multi-tenant environment is that if one tenant is being attacked, it can impact all of the other tenants. This vulnerability is magnified by the fact that as cloud computing service providers become better known, they attract an increasing amount of attention from cyber criminals.

Figure 1 depicts the typical n-tier application architecture that is commonly used by most IT organizations. Figure 1 gives the impression that the LAN and WAN are each a single entity. They are not. They are comprised of myriad switches and routers as well as devices such as firewalls, intrusion detection systems, intrusion protection systems and WAN optimization controllers. The resulting system is quite complex. Unfortunately, the probability of a security intrusion is proportional to the square of the complexity of the system. That follows because the more complex the system is, the more difficult it is for the IT organization to exert appropriate controls and the more places that a hacker can attack. As a result, even a small increase in the complexity of a system significantly increases the security risk.

[6] http://www.webtorials.com/content/2009/12/cloud-computing-a-reality-check-guide-to-risk-mitigation.html
[7] http://www.marketwire.com/press-release/Arbor-Networks-Fifth-Annual-Infrastructure-Security-Report-Finds-Service-Application-1103590.htm

Figure 1: The Typical n-Tier Application

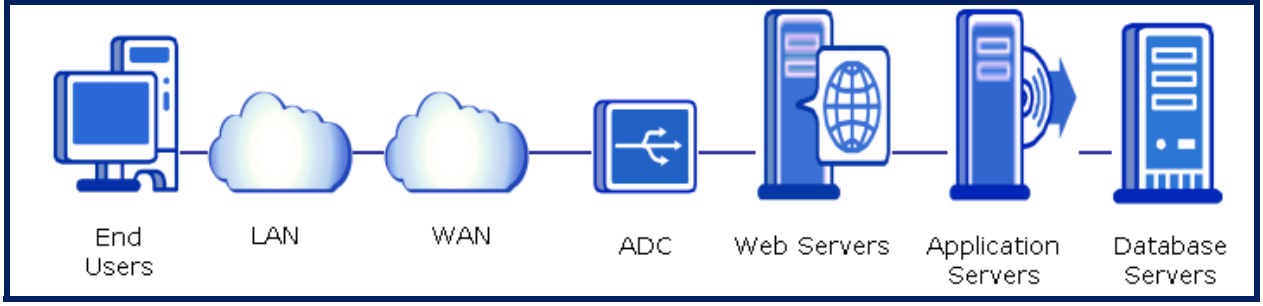

One example of how a relatively small change in the composition and complexity of an application can significantly increase the security risk is the growing use of AJAX, which is shorthand for asynchronous JavaScript and XML. AJAX is actually a group of interrelated web development techniques used on the client-side to create interactive web applications. While the interactive nature of AJAX adds significant value, it also creates some major security vulnerabilities. For example, if they are not properly validated, user inputs and user-generated content in an application can be leveraged to access sensitive data or inject malicious code into a site. According to the AJAX Resource Center[8] the growth in AJAX applications has been accompanied by a significant growth in security flaws and that this growth in security flaws "has the potential to turn AJAX-enabled sites into a time bomb."

Another reason that the global damage due to cyber crime is so high is that the IT organizations that lacked the time and the resources to do an adequate code review and vulnerability testing a few years ago, have even less time and resources today to perform these tasks. One measure of the impact of this can be found in a report[9] recently published by the Web Application Security Consortium. According to that report, almost 97% of Web sites contain a severe security vulnerability with the most wide spread security vulnerabilities being:

Cross-Site Scripting

This is a type of computer security vulnerability that enables malicious attackers to inject client-side script into web pages. This enables an attacker to gain elevated access privileges to sensitive page content, session cookies, and a variety of other information maintained by the browser on behalf of the user.

Information Leakage

This is an application weakness whereby an application reveals sensitive data, such as technical details of the web application, the environment (e.g., the framework, languages, or pre-built functions being utilized by a web application), or user-specific data. An attacker may use this sensitive data to exploit the targeted web application, its hosting network, or its users.

[8] http://www.ajaxtopics.com/security.html
[9] http://projects.webappsec.org/Web-Application-Security-Statistics#Summary

SQL Injection

This is an attack in which malicious code is inserted into strings that are later passed to an instance of SQL Server for parsing and execution. The primary form of SQL injection consists of direct insertion of code into user-input variables that are concatenated with SQL commands and executed.

The Web Application Security Consortium's report also stated that 99% of web applications are not compliant with the PCI Data Security Standard (DSS). Among the PCI DSS requirements is a mandate for IT organizations to implement Web and application security in order to minimize the risk due to security vulnerabilities such as those listed above. Requirement 6 of the PCI DSS requirements specifically calls for merchants and credit card issuers to develop and maintain secure systems and applications. The minimum vulnerabilities that must be accounted for are described in PCI DSS section 6.5.

The goal of PCI DSS section 6.6 is to ensure that web applications that are exposed to the public Internet are protected against at least this minimum set of vulnerabilities. Section 6.6 identifies two approaches that IT organizations can use to satisfy the PCI DSS requirements[10]. One approach is to utilize a Web application firewall and the other approach is to perform application code reviews. These reviews must be performed whenever changes are made to the code. If no changes have been made to the code, code reviews still need to be performed annually. While not required by section 6.6, some IT organizations implement a Web application firewall and they also perform application code reviews. However, many other IT organizations prefer to just implement a Web application firewall because of the excessive amount of resources that are required to perform an effective application code review.

Solutions

Most solutions in the Application Delivery 1.0 era, such as WOCs, ADCs and firewalls, were located onsite. While these solutions still provide value, many of solutions in the Application Delivery 2.0 era provide additional value because they are distributed across the Internet.

Internet-Based Application Delivery Services

One of the main drivers of Application Delivery 2.0 is cloud computing. There are two primary forms of cloud computing – public and private cloud[11]. *Public cloud computing* refers to organizations acquiring services from third parties such as Salesforce.com, Amazon and RackSpace. These third parties are often referred to as cloud computing

[10] http://www.google.com/url?q=https://www.pcisecuritystandards.org/minisite/en/docs/information_supplement_6.6.pdf&ei=AT-ES-vhJ43gsQPTmdnuAg&sa=X&oi=nshc&resnum=1&ct=result&cd=1&ved=0CA4QzgQoAA&usg=AFQjCNGCott-kObLqQpoehI4YGlJlur90w

[11] http://www.webtorials.com/content/2009/11/a-guide-for-understanding-cloud-computing.html

service providers (CCSPs). *Private cloud computing* refers to IT organizations implementing inside of their environment the same techniques as CCSPs implement, the goal of which is an order of magnitude increase in the cost effective, elastic provisioning of IT services.

As IT organizations continue to adopt cloud computing, this results in an increasing amount of traffic transiting a WAN in order to access IT resources that are housed in a remote data center. The conventional wisdom is that the vast majority of this traffic will transit the Internet. However, a recent report[12] highlighted the fact that many IT organizations have concerns about the use of the Internet to access cloud services. According to the report, when accessing public cloud services, IT organizations are only somewhat more likely to use the Internet than they are to use other WAN services. The report also pointed out that one of the WAN services that IT organizations are interested in using to access cloud services is an Internet overlay service, sometimes referred to as an Internet-based application delivery service (ADS). An ADS is an example of a type of solution that is closely associated with Application Delivery 2.0.

Both WOCs and ADCs are premise-based first generation application delivery solutions that continue to add value in the current environment. However, these solutions make the assumption that performance characteristics within the WAN itself are not optimizable because they are determined by the relatively static service parameters controlled by the WAN service provider. This assumption is reasonable in the case of WAN services such as MPLS. However, this assumption does not apply to enterprise application traffic that transits the Internet because there are significant opportunities to optimize performance within the Internet itself based on Application Delivery Services (ADSs). An ADS leverages service provider resources that are distributed throughout the Internet in order to optimize the performance, security, reliability, and visibility of the enterprise's Internet traffic. As shown in **Figure 2**, all client requests to the application's origin server in the data center are redirected via DNS to an ADS server in a nearby point of presence (PoP) close to application users, typically within a single network hop. This edge server then optimizes the traffic flow to the ADS server closest to the data center's origin server.

An ADS provides a variety of optimization functions that generally complement Application Delivery 1.0 solutions such as an ADC rather than overlap or compete with them. One such function is content offload. This calls for taking static content out of a data-center and placing it in caches in ADS servers and in replicated in-cloud storage facilities. IT organizations that offload content and storage improve response time and simultaneously reduce both their server utilization as well as the bandwidth utilization of their data center access links.

[12] http://www.webtorials.com/content/2009/12/cloud-computing-a-reality-check-guide-to-risk-mitigation.html

Figure 2: An Internet-based ADS

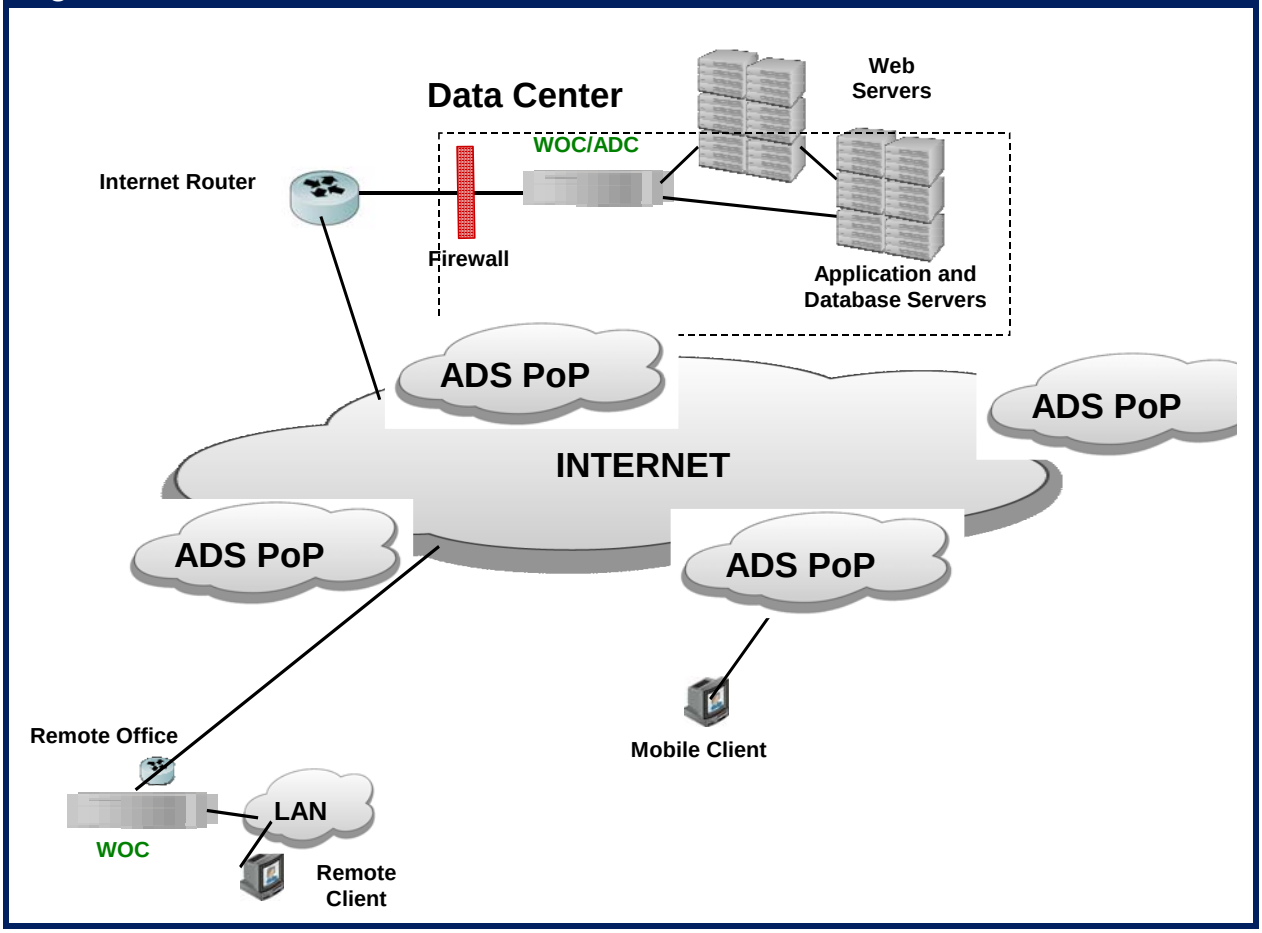

Some of the other common ADS functions include:

- Route optimization
- Transport optimization
- HTTP protocol optimization
- Visibility

In addition to the functionality listed above, many ADCs now incorporate Web application firewall functionality. As is discussed in the following section, an ADC that supports Web application firewall functionality is complimentary to an ADS that supports a Web application firewall service.

Defense in Depth: The Role of a Web Application Firewall Service

As was previously mentioned, there are fundamental flaws with an approach to security that focuses only on the perimeter of the organization. To overcome these flaws, most IT organizations have moved to an approach to security that is typically referred to as *defense in depth*. The idea behind defense in depth, an approach to security that was first conceived by the National Security Agency (NSA), is that defending an organization

against sophisticated attacks requires using several security methods that are deployed in layers. The advantage of this approach is that if an intruder manages to penetrate one layer of security, there are additional layers of security separating the intruder from the organization's assets.

The concept of defense in depth is not new. This approach was widely used during the Application Delivery 1.0 era as IT organizations often deployed multiple layers of security functionality including virus scanning, authentication, firewalls, intrusion detection systems and intrusion protection systems. In the Application Delivery 1.0 era, however, all of the layers of security functionality were typically deployed onsite. What is new in the Application Delivery 2.0 era is the deployment of a layer of security, such as a Web application firewall service, that is distributed throughout the Internet so that it is close to the source of security attacks and hence can prevent many security attacks from reaching the organization. The distribution of security functionality on the part of a Web application firewall service is analogous to the distribution of optimization functionality on the part of an application delivery service that was discussed in the preceding section. Because IT organizations are acquiring a service, and not installing additional hardware, this eliminates the need for IT organizations to expend capital or get involved with complex tasks such as capacity planning for a Web application firewall.

One of the advantages of preventing a security attack from reaching an organization was highlighted in the Arbor networks report. As previously noted, that report stated that peak DDoS attacks generate 49 Gbps of traffic and that since most enterprises are connected to the Internet at speeds of 1 Gbps or less, any attack over 1 Gbps will be effective. Attacks of this magnitude cannot be prevented by onsite solutions. They can, however, be prevented by implementing security functionality such as a Web application firewall service that identifies and mitigates the DDoS-related traffic close to attack traffic origin.

The way that DDoS attacks are often described, it may appear as if all DDoS attacks are similar in nature. The reality is that organizations are vulnerable to a very broad diversity of DDoS attacks[13]. For example, a DDoS attack can originate from a specific region or more globally from around the world. The attack can target an organization's entire IT environment or just the environment at a single site. In any case, a DDoS attack can cause harm to an organization in a number of ways, including the:

- Consumption of computational resources, such as bandwidth, disk space, or processor time.
- Disruption of configuration information, such as routing information.
- Disruption of state information, such as the unsolicited resetting of TCP sessions.

[13] http://en.wikipedia.org/wiki/Denial-of-service_attack

- Disruption of physical network components.

- Obstructing the communication media between the intended users and the victim so that they can no longer communicate adequately.

Some of the most common forms of a DDoS attack include an ICMP flood attack, teardrop attacks, and reflected attacks.

Because there are a variety of possible DDoS attacks, IT organizations need to implement a variety of defense in depth techniques. This includes:

- Minimizing the points of vulnerability

  If an organization has most or all of its important assets in a small number of locations, this makes the organization more vulnerable to successfully being attacked as the attacker has fewer sites on which to concentrate their attack. One way to avoid this vulnerability is to offload content to a cloud computing platform. As previously mentioned in the discussion of application delivery services, this approach also improves application performance.

- Protecting DNS

  Many IT organizations implement just two or three DNS servers. As such, DNS is an example of what was discussed in the preceding bullet – how IT organization are vulnerable because their key assets are located in a small number of locations. As was also discussed in the preceding bullet, one way to minimize this vulnerability is to distribute DNS functionality around a global network.

- Implementing robust, multi-tiered failover

  Many IT organizations have implemented disaster recovery plans that call for there to be a stand-by data center that can support at least some of the organization's key applications if the primary data center fails. Distributing this functionality around a global network increases overall availability in general, and dramatically reduces the chance of an outage due to a DDoS attack in particular. A complimentary approach is to implement techniques to keep a site from failing. These techniques include directing requests to a virtual waiting room or to a site with reduced functionality.

In order to be effective, a Web application firewall service needs to be deployed as broadly as possible, preferably in tens of thousands of locations. This is necessary in order to:

- Be able to intercept malicious traffic as close as possible to the source of the attack.

- Scale to support the massive and growing scale of DDoS attacks.

- Not be vulnerable itself to a DDoS attack.

Since organizations are vulnerable to a broad range of attacks, an effective Web application firewall service must be able to flexibly implement a wide range of security techniques in order to both eliminate attacks before they occur and to respond to them quickly when they do occur. For example, in order to eliminate potential attacks, a Web application firewall service must be able to authorize, deny, or redirect traffic based on characteristics such as user-agent (e.g. browser) or language. The service must also be able to identify when there is an unusually high level of new users, as this may indicate an attack. When faced with a potential attack the Web application firewall service must be able to quarantine suspicious traffic to a small set of servers and to limit the rate at which requests are forwarded to the origin server in order to ensure its availability.

When responding to an attack, a Web application firewall service must also be able to:

- Block or redirect requests based on characteristics such as the originating geographic location and whether or not the originating IP addresses are on either a whitelist or a blacklist.

- Direct traffic away from specific servers or regions under attack.

- Issue slow responses to the machines conducting the attack. The goal of this technique, known as tarpits[14], is to shut down the attacking machines while minimizing the impact on legitimate users.

- Direct the attack traffic back to the requesting machine at the DNS or HTTP level.

As was previously noted, an ADC that supports Web application firewall functionality is complimentary to an ADS that supports a Web application firewall service. That follows because while a Web application firewall service can perform many security functions that cannot be performed by a Web application firewall, there are some security functions that are best performed by a Web application firewall. An example of that is protecting an organization against information leakage by having an onsite Web application firewall perform deep packet inspection to detect if sensitive data such as a social security number or a credit card number is leaving the site. If sensitive data is leaving the site, the onsite Web application firewall, in conjunction with other security devices, can determine if that is authorized and if it is not, prevent the data from leaving the site.

[14] http://en.wikipedia.org/wiki/Tarpit_(networking)

Summary

In order to respond to the ongoing challenges that are associated with the Application Delivery 1.0 era, most IT organizations have either already deployed, or are in the process of deploying, a first generation of solutions. These solutions are intended to protect the organization from a growing number of security attacks, to mitigate the impact of chatty protocols and to offload computationally intensive processing from servers. A characteristic of most Application Delivery 1.0 era solutions such as network firewalls, WOCs and ADCs is that they are deployed onsite and hence they only process traffic as it enters and leaves a site.

Driven by a number of factors, including the shifting emphasis and growing sophistication of cyber crime and the wide spread adoption of cloud computing, IT organizations are now entering the Application Delivery 2.0 era. The fact that IT organizations need to respond to a new generation of challenges doesn't mitigate the need for IT organizations to continue to implement site-based solutions such as ADCs. In most cases, these first generation solutions are complimentary to Application Delivery 2.0 era solutions.

A common characteristic of Application Delivery 2.0 era solutions is that they are decentralized. An example of this is an ADS that leverages service provider resources that are distributed throughout the Internet. An ADS provides a variety of optimization functions that generally complement the functionality provided by an onsite ADC. One such function is content offload. This calls for taking static content out of a data-center and placing it in caches in ADS servers and in replicated in-cloud storage facilities. IT organizations that offload content and storage improve response time and simultaneously reduce both their server utilization as well as the bandwidth utilization of their data center access links.

Another example of an Application Delivery 2.0 era solution is a Web application firewall service. One of the advantages of a Web application firewall service vs. an onsite Web application firewall is that because the service is distributed throughout the Internet, it can stop attacks close to the origin. In order to be effective, a Web application firewall service needs to be deployed as broadly as possible, preferably in tens of thousands of locations. When responding to an attack, a Web application firewall service must be able to perform a number of tasks such as directing traffic away from specific servers or regions under attack. However, just as an ADS is complimentary to an onsite ADC, a Web application firewall service is complimentary to a Web application firewall. For example, in many cases a task such as preventing data leakage is best performed by an onsite Web application firewall.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Webtorials Briefs
Vol 3, Number 1

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer
All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.

Enhancing Web Application Security: Akamai Web Application Firewall Managed Service

Enhancing application security to support your business

As organizations drive business online, valuable corporate and financial assets are now increasingly connected to the Internet. With the pervasiveness of critical Web applications, threats can easily bypass traditional firewall security controls resulting in the majority of attack traffic now using the HTTP protocol to target Web applications. A security breach in an enterprise's Web application should not be taken lightly. Corporations can suffer loss of proprietary and personal information, as well as possible fines, sanctions, and lost revenues. It also impacts customer and brand loyalty.

Key Challenges in Securing Web Applications

In recent years, there has been a dramatic shift in the way applications are delivered from the data center. Instead of being confined to residing within the firewall on a company's private network, applications are increasingly being extended to the Internet. This allows a broad range of users, such as mobile employees, business partners and customers, to have access to Web-enabled applications simply by having last-mile Internet connectivity – anywhere, anytime. This trend is further fueled by the momentum behind public cloud computing services, where core enterprise applications are accessed over the Internet residing within a multi-tenant infrastructure offered as a service by third parties.

Web application layer vulnerabilities are often introduced by developers making mistakes with rapidly changing code base. This is compounded by the increasing sophistication of website and application design, the leveraging of newer techniques such as AJAX, XHTML, and integration with third party services through web services. With cloud computing, Web application security poses a bigger challenge. Attacks directed at other customers can take down one of your own applications.

"Defense in Depth" approach for Enterprise Security

Traditional security products such as appliance-based Web application firewalls are fairly well understood. They are deployed within the data center and provide centralized protection from attack traffic. But a datacenter approach to application security is not effective in combating all types of threats and vulnerabilities. As attack traffic increases in both scale and sophistication, new cloud-based services are needed to offer on-demand scalability and extend the security perimeter to the edge of the Internet.

For a security strategy to be effective, enterprises need to consider a "defense in depth" approach. "Defense in depth" architecture includes the use of a cloud based Web Application Firewall (WAF) service along with traditional datacenter security tools. In addition to providing defensive cover, cloud based WAFs stop the majority of attack traffic at the source. This helps offload origin based security hardware to focus on tasks which are best filtered in the data-center; such as deep packet inspection to ensure private information does not extend outside the firewall without authorization. A cloud based WAF with a distributed architecture at the Internet's edge provides an integrated view of threats on a global basis and is able to effectively protect cloud-based as well as datacenter-based applications. A "defense in depth" approach is therefore critical for today's enterprises to effectively secure their corporate resources.

Akamai Web Application Firewall Service

Figure 1: Akamai WAF

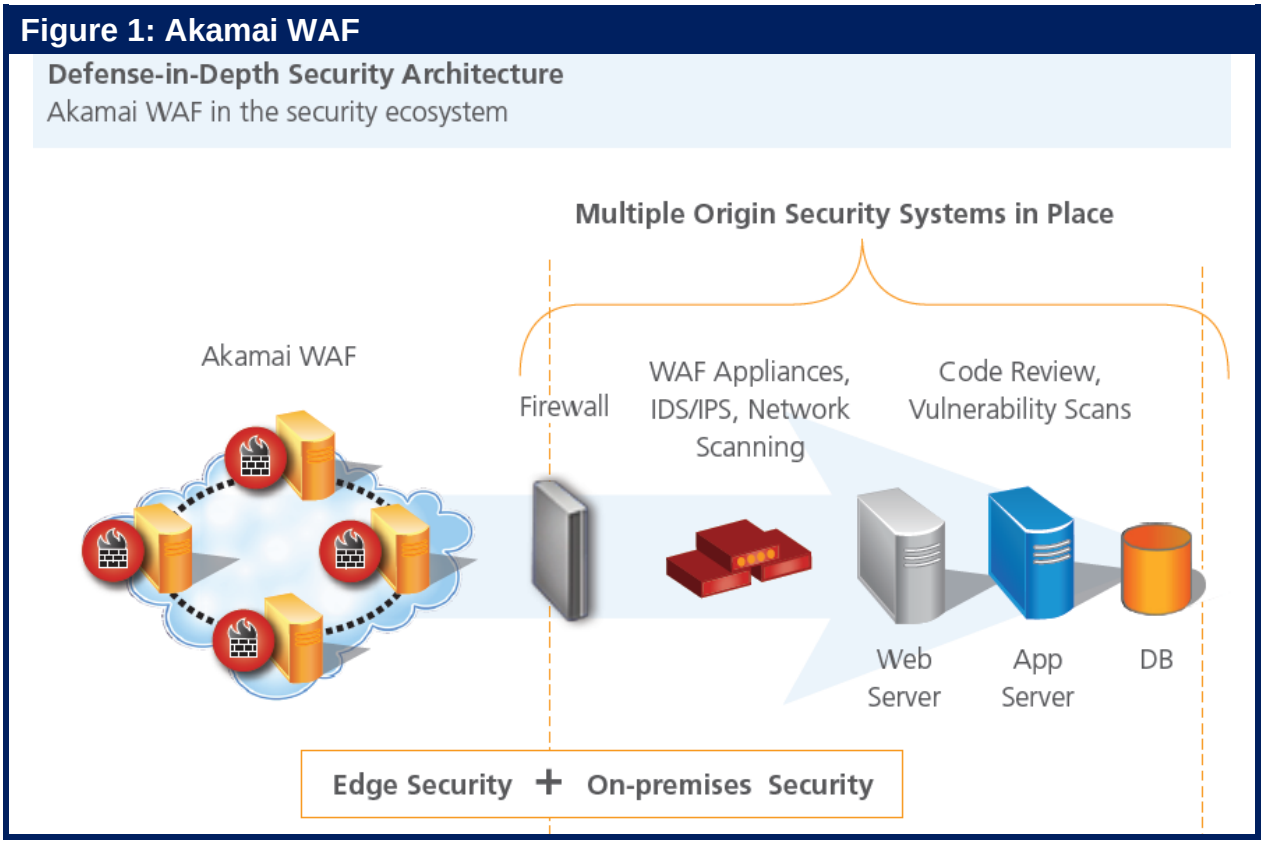

The Akamai WAF (**Figure 1**) provides customers a highly scalable, outer defensive ring of Web application protection. Akamai WAF is a *service* that runs on Akamai's distributed network of tens of thousands of servers and mitigates Web application attacks before they reach origin servers. It detects malicious Web traffic, issuing alerts and/or blocking traffic with multiple, configurable responses.

The Akamai WAF service is based on the ModSecurity Core Rule Set (CRS). CRS is very robust, flexible and provides critical protection against Web attacks. The Akamai WAF architecture allows it to sit inline thereby allowing it greater flexibility in denying traffic. Its distributed architecture provides on demand scalability to defend against really large attacks (**Figure 2**). It protects against OWASP top 10 vulnerabilities like SQL Injection and Cross Site Scripting (XSS), among other Layer 7 attacks. It also allows for custom rules or "virtual patching" to be deployed in the cloud.

Akamai WAF can be deployed stand alone or as part of a "*defense-in-depth"* security architecture. Since it is a managed service, there is no physical deployment and no network changes are required. It is very easy to integrate and is complementary to other origin-based WAF solutions and/or security controls like Intrusion Prevention Systems (IPS). Akamai WAF when deployed with existing WAF appliance enhances security and further mitigates risk. Akamai WAF also provides network layer controls such as IP whitelisting and IP blacklisting. Changes to IP whitelists and blacklists are implemented quickly through the web portal.

Akamai WAF runs concurrently with other services on the globally distributed Akamai platform. As a result, users not only get improved protection from Web application attacks, but also a comprehensive suite of application delivery services to optimize the performance, scale, availability, security and visibility of dynamic websites and business applications.

Figure 2

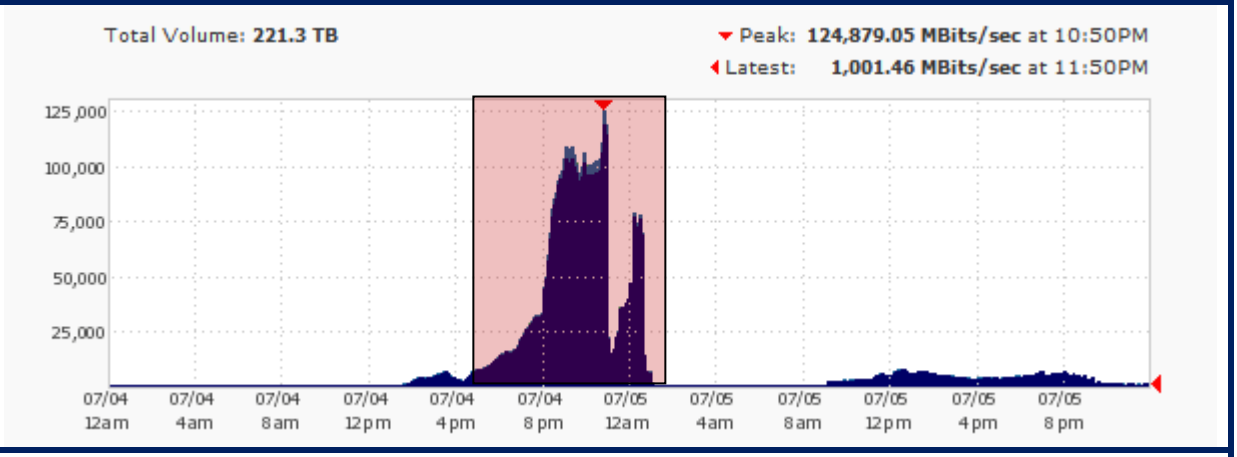

600X spike in traffic caused by a Layer 7 botnet driven DDOS attack on a public website

Customer Benefits

- Protects from information theft and site downtime caused by web application attacks
- Scales on-demand as attack traffic volume increases
- Convenient managed service complements origin based security solutions.
- Helps businesses comply with evolving Payment Card Industry (PCI) standards

About Akamai

Akamai® provides market-leading managed services for powering video, dynamic transactions, and enterprise applications online. Having pioneered the content delivery market one decade ago, Akamai's services have been adopted by the world's most recognized brands across diverse industries. The alternative to centralized Web infrastructure, Akamai's global network of tens of thousands of distributed servers provides the scale, reliability, insight and performance for businesses to succeed online. Akamai has transformed the Internet into a more viable place to inform, entertain, interact, and collaborate. To experience The Akamai Difference, visit www.akamai.com.

For helpful insights and thought leading whitepapers on how Akamai Security Solutions protect your key Web properties, **visit** www.akamai.com/security