



The Application Delivery Firewall Paradigm

The increasing sophistication, frequency, and diversity of today's network attacks are overwhelming conventional stateful security devices at the edge of the data center. A new data center architecture based on the security services of the F5 application delivery firewall solution effectively combats modern attacks while providing significant CapEx savings.

White Paper
by Lori MacVittie



Introduction

In most organizations, firewalls are the first line of defense for web and application services. The firewall is, and has been, the primary foundation around which conventional network security architectures are built. The conventional architecture has matured so that many security standards require the deployment of certified firewalls. For example, any data center that processes credit card numbers must comply with the Payment Card Industry (PCI) standard, which requires a certified network firewall. The de facto standard referenced by PCI auditors is an International Computer Security Association (ICSA) Labs–certified network firewall. ICSA defines a short list of firewalls that can be used for card processing purposes.

But the conventional firewall is beginning to show its limitations in detecting and repelling modern attacks. Attacks targeted at the application or network layers are causing failures of these stateful—and often expensive—firewalls, and the number of such attacks is growing.

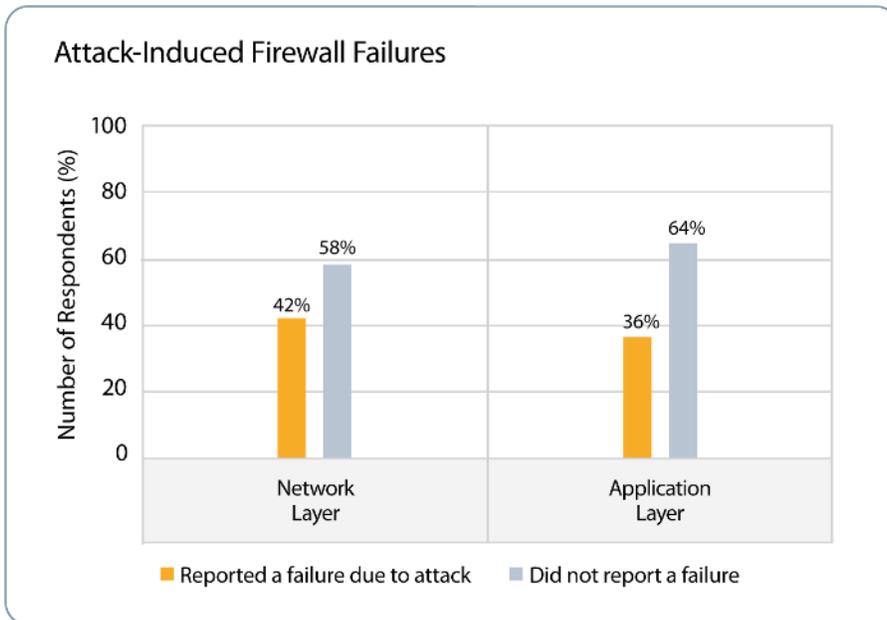


Figure 1: In 2011, significant numbers of respondents to an Applied Research survey experienced a stateful firewall failure due to an application or network layer attack¹.



WHITE PAPER

The Application Delivery Firewall Paradigm

These firewall failures are particularly disconcerting in light of the circumstances enabling attackers. While the Anonymous and LulzSec attacks were tightly focused and required planning, the bulk of attacks today require no such preparation, thanks to the creation of a vast pool of resources upon which attackers can draw to overwhelm their chosen targets. The lack of legal oversight in emerging technology powers, such as China and India, has enabled the establishment of massive botnets that can be leased on a moment's notice.

These increasingly diverse attacks involving multiple layers of the network stack are causing firewall failures with alarming frequency. As a result, traditional firewall services alone are insufficient for detecting attacks and subsequently preventing business disruption. It is necessary to also employ capabilities at the application layer to halt attacks that take advantage of application-layer protocols and behaviors.

Firewall Limitations

The selection of the conventional firewall has traditionally been informed by certification, expenditure, and performance. Certification criteria may require specific firewall products for compliance, limiting the selection of devices that can be chosen. From that set, purchasing personnel plot the remaining two factors: price versus performance. A new analysis of these parameters, however, suggests a new paradigm.

Firewalls are rated by data throughput—for example, 1 Gbps or 4 Gbps—and this makes it easy to match a purchase to the size of the ingress pipe. But bulk throughput is not the real story. During a distributed denial-of-service (DDoS) attack, it's not just bulk throughput that matters; it's how the device can handle concurrent connections and connections per second. For example, a typical \$50,000 conventional firewall claims a throughput of 10 Gbps, which should be sufficient for a small- to medium- sized attack. But this class of firewalls can only handle between 1 and 2 million concurrent connections. It is no secret that the WikiLeaks attackers of 2010, using only a single botnet, easily generated more than 2 million concurrent connections, tipping over firewalls across the United States. A conventional firewall with the next level of concurrent connection performance (4 to 10 million connections per second) requires a step up the price ladder into the \$100,000 to \$150,000 tier.

The story is the same for connections per second. When a conventional firewall performs stateful inspection, it incurs a performance penalty for each TCP session set up. This limits the performance of the firewall for incoming connections. In the \$50,000 price bracket, a typical conventional firewall can handle 50 to 100,000 new connections per second.

WHITE PAPER

The Application Delivery Firewall Paradigm

Attackers are aware of these firewall limitations, and modern attacks are designed to exploit them. The resulting firewall failures are not uncommon, unfortunately, as industry analysts have noted. Indeed, these failures may be the reason a mere 8 percent of respondents in a September 2011 security survey indicated that traditional secure measures such as firewalls were sufficient to provide network security.²

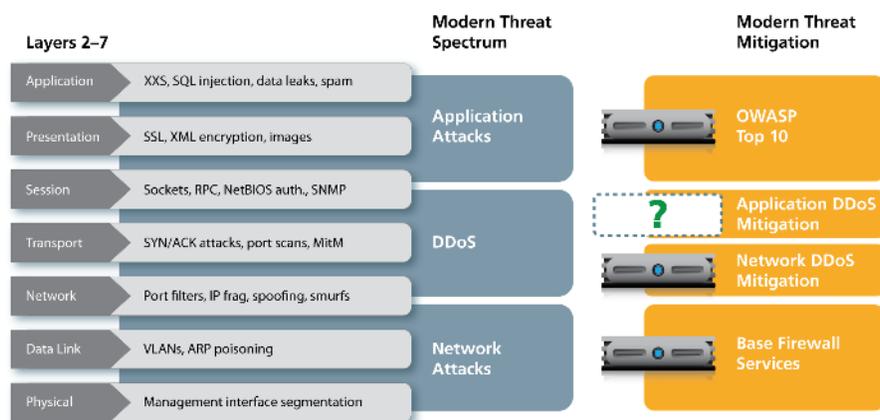


Figure 2: Modern threat mitigation techniques fail to completely address all threats, particularly those aimed at DDoS.

Another limitation of the conventional firewall deployment architecture is its ability to handle the breadth of today's threat spectrum, which encompasses the entire network and application ecosystem. Solutions intended to mitigate these wide-ranging threats have traditionally been deployed separately, with specific technology to address attacks in logical groupings, such as application, network, and DDoS attacks. These disconnected solutions from multiple vendors increase overall management complexity and, of course, incur significant capital and operational expenditures.

When considering the edge of the modern data center, customers are questioning paying any price for a conventional firewall that does little more than pass traffic through port 80, add latency, and incur expense and exposure. Nimble firms, especially startups and those sites without PCI requirements, have been able to run without conventional firewalls for some time. Companies reliant on Web 2.0 and other data center transactions increasingly benefit from a new data center architecture fronted by an integrated security device.



WHITE PAPER

The Application Delivery Firewall Paradigm

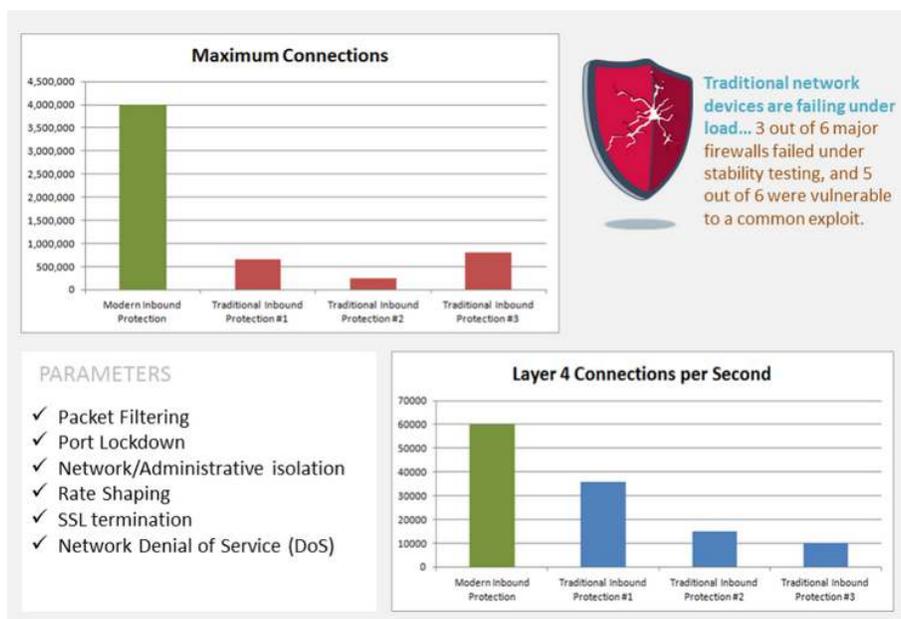


Figure 3: Connection capacity, both maximum and per second, underlies the most important new firewall performance metrics.

A New Data Center Architecture

The F5 approach to the firewall problem—the application delivery firewall solution—converges security services into a single set of Application Delivery Controllers (ADCs) at the edge of the data center.

With version 11.3, F5 breaks new ground by introducing a new firewall product as well as integrated firewall management services into its flagship BIG-IP product family. BIG-IP Advanced Firewall Manager (AFM) is a high-performance, stateful, full-proxy network firewall designed to guard data centers against incoming threats that enter the network on the most widely deployed protocols—including HTTP/S, SMTP, DNS, and FTP. Organizations can combine BIG-IP AFM with F5's other security services to build a new security architecture based on the application intelligence of F5's application delivery firewall solution.

The integration of firewall services in the BIG-IP family means that BIG-IP Local Traffic Manager (LTM), BIG-IP Global Traffic Manager (GTM), and BIG-IP Application Security Manager (ASM) can be placed strategically at the edge of the data center while still maintaining a proper security posture and compliance for the organization.



WHITE PAPER

The Application Delivery Firewall Paradigm

The F5 application delivery firewall solution provides network-layer protection with a much higher connection capacity than traditional firewalls. The application delivery firewall solution can handle hundreds of millions of connections, managing them with various timeout behaviors, buffer sizes, and other security-focused options when under attack. This capacity enables BIG-IP LTM to manage the volume of a traffic onslaught while performing the port and IP-based access control services typically provided by a stateful firewall.

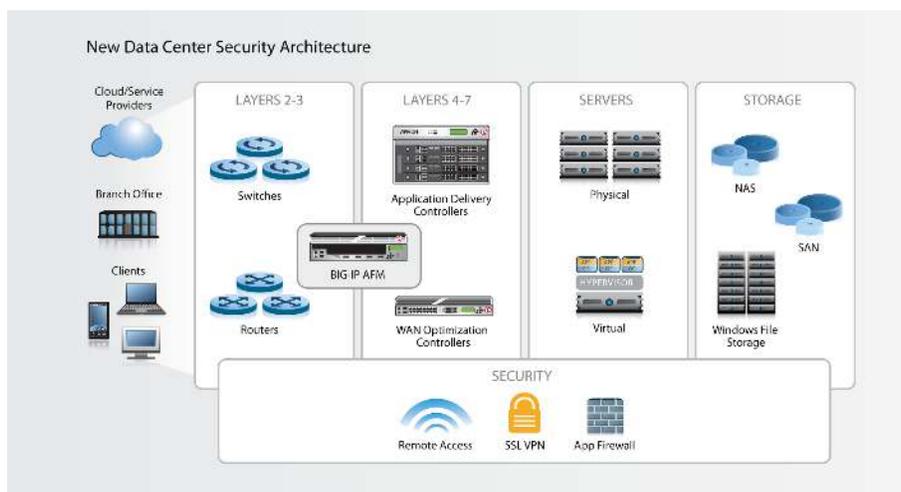


Figure 4: The new paradigm replaces stateful firewall services with BIG-IP LTM in the data center architecture.

Native Application Protocol Fluency

The F5 application delivery firewall solution can help halt attacks that take advantage of application layer protocols and behaviors. Because the application delivery firewall is fluent in application protocols, it can monitor and act on behavior, not just specifications and standards. The application delivery firewall solution decodes IPv4, IPv6, TCP, HTTP, SIP, DNS, SMTP, FTP, Diameter, and RADIUS communications, enabling more sophisticated analysis based on protocol as well as payload. This allows the application delivery firewall to detect anomalies indicating an attack in progress and to take appropriate action. For example, the application delivery firewall can detect the number of layer 7 connections per second, per client, and impose various rate-limiting schemes that have proven effective in mitigating layer 7 attacks.

This native protocol fluency also ensures enforcement of protocol compliance, mitigating attacks that seek to leverage vulnerabilities introduced by lax interpretation of the protocol. The combination of protocol compliance and F5's full-proxy architecture results in a unique DDoS mitigation solution.



WHITE PAPER

The Application Delivery Firewall Paradigm

The native aspect of the protocol compliance enforcement is significant. The programmatic ability of the F5 iRules scripting language provides a flexible means of enforcing protocol functions on both standard and emerging or custom protocols. Using iRules, organizations can direct BIG-IP LTM to enforce protocol compliance and perform rate limiting, response injection, and traffic steering and related actions. Security teams are finding that the flexibility of iRules enables them to mitigate a broad range of security solutions:

- Using iRules in BIG-IP LTM, organizations can build a fingerprint-cloaking profile for application servers by obfuscating server and OS headers and rewriting outbound HTTP response codes (such as 301, 401, and 501 errors).
- For transport layer security, iRules can reach deep into the SSL/TLS protocol stack, mitigating protocol attacks such as 2010's SSL renegotiation vulnerability, in which a single handset can attack a secure server.
- Using iRules, organizations can react quickly to application vulnerabilities for which a patch has not yet been released. Mitigating iRules can be developed in-house, sourced from F5's global DevCentral development community, or even published from F5 product development. The Apache Killer³ vulnerability, for example, was addressed by an iRule from the F5 security team weeks before an official solution was published by the Apache Server Foundation.

Advanced Network Protection

The BIG-IP Advanced Firewall Manager (AFM) module, new in version 11.3, integrates with BIG-IP LTM to provide three primary security services. First, BIG-IP AFM offers high-speed, application-aware firewall rules that allow security personnel to manage their layer 4 network. Second, BIG-IP AFM is aware of 38 types of DDoS attacks and automatically alerts and mitigates them (organizations can further define their own DDoS scenarios). Finally, BIG-IP AFM offers, for the first time, granular visibility and logging at the application level, allowing organizations to slowly begin to deprecate their zone abstractions and provide instrumentation directly to the individual application teams.



Advanced DNS Protection

For advanced DNS protection beyond that natively provided by version 11.1 of BIG-IP LTM, BIG-IP GTM adds iRules support, increasing the native fluency and compliance protection for the DNS protocol. BIG-IP GTM was the first commercial global traffic manager to support the Domain Name System Security Extensions (DNSSEC), thereby providing protection from cache-poisoning and man-in-the-middle attacks. Add the DNS Express feature of BIG-IP GTM to protect vital DNS services from DoS attacks.

Advanced Web Application Protection

For advanced web application security, the integrated BIG-IP ASM module provides web application firewall (WAF) services to secure individual applications against the Open Web Application Security Project (OWASP) Top 10 attacks, such as cross-site scripting (XSS), cross-site request forgeries (CSRF), and SQL injections. BIG-IP ASM is the only web application firewall with a learning mode that enables it to become aware of an application's normal input parameters and reject attacks that don't fit the normal traffic pattern. BIG-IP ASM also fulfills the critical WAF requirement in the PCI 2.0 specification.

Web Access Management

BIG-IP Access Policy Manager (APM) is the final component of the application delivery firewall. Many web applications need to limit certain users' access, and BIG-IP APM supports this requirement with multi-factor authentication, authorization, and single sign-on (SSO) services. Dynamic access control into the data center is enforced using level 4 and level 7 access control lists (ACLs) derived from contextual information such as user identity, endpoint inspection results, geolocation, and any attribute pulled from a directory store. BIG-IP APM performs exceptionally by enforcing ACLs at forwarding speeds of up to 72 Gbps, supporting thousands of logins per second, and scaling to 100,000 concurrent users on a single platform.

Cumulative Benefits

The cumulative effect of these benefits (performance, protocol compliance, full-proxy architecture, access control, and iRules flexibility) is reduction of the overall threat surface. Having fewer, higher-capacity devices means fewer configurations and, ultimately, fewer battles to fight during an attack. IT staff can concentrate their defenses at a single point of control instead of doing the reboot dance as different devices fail up and down the security stack. This reduction of the threat surface works to mitigate the breadth of today's threat spectrum, which encompasses not only traditional network attacks, but also complex DDoS attacks and layer 7 vulnerabilities.



An application-proxy gateway is a feature of advanced firewalls that combines lower-layer access control with upper-layer functionality. These firewalls contain a proxy agent that acts as an intermediary between two hosts that wish to communicate with each other, and never allows a direct connection between them.

—NIST, "[Guidelines on Firewalls and Firewall Policy](#)"



WHITE PAPER

The Application Delivery Firewall Paradigm

An approach using the F5 application delivery firewall solution consolidates security services for a superior defense against all three attack types (network, DDoS, and application) within a full-proxy architecture, something no conventional stateful firewall can provide.

Conclusion

Stateful firewalls have been the centerpiece of security at the edge of data centers for the past 25 years. But cracks are starting to appear in conventional firewall-based architectures as attackers use new techniques and global botnets to turn what used to be the defensive shield into a liability, exactly when that shield is needed most. The threat spectrum has significantly changed over time; traditional firewalls can mitigate simple network attacks, and so-called "next-generation" firewalls can address the outbound vulnerabilities of the enterprise data center.

But only a new data center firewall architecture—one based on F5 products and featuring full-proxy, high connection-capacity ADCs at the edge of the network—can ensure standards-based compliance while significantly lowering CapEx by eliminating firewall devices and upgrades and maximizing other data center resources.

Agile organizations are converging security services to address the three primary vectors of the modern data center threat spectrum:

- Traditional network attacks.
- Complex DDoS attacks on HTTP and DNS.
- Application-level vulnerabilities.

The new application delivery firewall paradigm addresses each of these vectors in a complete, integrated solution. Traffic management and network firewall services are managed by BIG-IP LTM and BIG-IP AFM. Deploy BIG-IP GTM to implement DNSSEC and DNS Express that protect vital DNS services from DoS and hijacking attacks, and deploy BIG-IP ASM to provide web application firewall services for the OWASP Top 10. Complete the solution with BIG-IP APM to provide secure web access management and SSO for applications. This F5 application delivery firewall solution is a modern threat mitigation platform that provides complete protection from the bottom to the top of the network stack.

A security solution centered on F5 products enables organizations to implement a holistic, scalable security strategy that can mitigate today's most challenging attacks while remaining flexible enough to address those that will undoubtedly appear tomorrow.



WHITE PAPER

The Application Delivery Firewall Paradigm

- ¹ Applied Research 2011 ADC Security Study
- ² Applied Research 2011 ADC Security Study
- ³ F5 Friday: Zero-Day Apache Exploit? Zero-Problem

F5 Networks, Inc.
401 Elliott Avenue West, Seattle, WA 98119
888-882-4447 www.f5.com

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

Europe/Middle-East/Africa
emeainfo@f5.com

Japan
f5j-info@f5.com