Securing Your Enterprise Applications with the BIG-IP

Using the BIG-IP System with the Application Security Module for Comprehensive Application and Network Security

Overview

The Internet has become increasingly complex, leaving many enterprises vulnerable to malicious attacks. Organizations are faced with trying to protect their infrastructure against network security attacks, as well as attacks that are specific to the application layer. Every year, security breaches cost companies millions of dollars in revenue, productivity, and lost reputations.

Organizations have traditionally responded to these threats by enhancing network security through the use of firewalls. However, this relatively narrow scope has proven insufficient. Although traditional firewalls may protect an organization against network attacks, they are incapable of defending against the new breed of attacks targeting the application layer. Organizations are looking for more dependable, scalable solutions to broaden their security reach and increase their protection levels. F5's BIG-IP® Local Traffic Manager, an application delivery networking solution, gives organizations the tools they need to achieve comprehensive security both at the network and at the application layers.

This white paper describes how the BIG-IP provides a holistic and integrated approach to securing your network and applications against potential application-level and network-level threats and attacks. With the BIG-IP and/or the new Application Security Module, you can realize the following benefits:

- **Robust Application Security** – Enforce, fortify, accelerate, and secure the delivery of your applications and web services. With features like Application Security Module's positive security model, content scrubbing, cookie and session management and protection, application and content filtering, and powerful encryption, organizations can implement comprehensive application security, providing a coordinated and unified line of defense that lowers TCO and improves ROI.

- **Powerful Network Security** – Enforce, fortify, and implement security policies for your networking infrastructure. With features like DoS and SYN attack prevention, packet filtering, and protocol sanitization, organizations can protect themselves against the heaviest of attacks and control the information traversing in and out of their site.

- **Increased ROI** – Maximize application availability for trouble free maintenance and reduced administration overhead. By offloading SSL and critical security functions (processor and server intensive operations), you do not have to buy expensive hardware to support your applications. The result is up to 30% savings on hardware costs with increased application performance. Rate Shaping enables you to allocate bandwidth according to your business policies, saving money and realizing a better ROI.

- **High availability** – Through the use of advanced health-checking capabilities, the BIG-IP can recognize when a resource is unavailable or under performing and redirect traffic to another resource. With the BIG-IP, all of your applications can achieve mission-critical availability (99.999% uptime), while reducing operational complexity and costs.

- **Extensible integration** – iControl is the industry's first open Application Programming Interface (API) for a comprehensive suite of application traffic management products. Made available as a free SDK, iControl overcomes integration challenges, making it quick and easy to create and automate intercommunication between 3rd party applications and the network via F5's products.

Challenges

Applications have become a core component in the way today's organizations conduct business. Applications have a direct impact on the revenue stream of an organization, so protecting business-critical information from malicious attacks that focus on application vulnerabilities, in addition to low-level network attacks, is vital. Organizations are faced with many challenges to achieve true network and application security, because:

- **Application vulnerabilities are on the rise** – Security systems and firewalls today are not smart enough nor are they designed to detect the new breed of application-layer attacks, let alone protect against them. These devices are application agnostic and simply lock or unlock an address, port, or resource. They're either all on or all off, depending on a response from signature matching. Even robust IDS/IPS devices have limited visibility to inspect deep within the packet, and do not maintain any session state information to detect and prevent application attacks from occurring. Application attacks often involve injecting and executing restricted commands, cookie tampering, and gaining illegal access to sensitive documents and user information. These attacks are causing massive loss of revenue and productivity, and can adversely affect an organization's reputation.

- **Increasing network vulnerabilities** – Network attacks are becoming more pervasive and sophisticated across all OSI layers. Malicious users are finding new ways to penetrate a site's defenses and hijack valuable information or even bring down an entire web site. Sophisticated network attacks such as DoS, DDoS, out-of-order packet floods, TCP window size tampering, and so on, are putting tremendous pressure on security systems to protect against high-volume attacks without getting overwhelmed. Tried-and-true network attacks have become "old hat" and are easily defeated the current network firewalls. Now, hackers and malicious users have started scanning sites (*profiling*) to retrieve any system or application information from seemly innocuous sources, like server error codes and source code comments, before launching an attack. Application attacks are becoming more intelligent and invasive every day.

- **Internal security breaches and information leaks** – One of the biggest threats today are security attacks from people within the organization. While external attacks are extremely important and critical to today's e-commerce world, internal attacks should not be overlooked. These internal attacks are very hard to detect and prevent because users within the organization are part of the trusted domain. Perimeter security provides no protection if the attack stems from within the perimeter. Using existing perimeter security solutions, organizations are struggling to implement an effective, unified security policy that complies with regulatory security standards (Sarbanes-Oxley, HIPAA, FIPS), while maintaining the level of ubiquitous network and resource access required by today's mobile and remote workforce.

Solution

The BIG-IP system provides a wide variety of security in both the network and application layers that play a significant role in bolstering the security of an organization. From adding powerful network-level security policies to filtering the most sophisticated application attacks, the BIG-IP is deployed as a critical gateway to your

most precious resources: the applications and networks that run your business. As the leader in integrated SSL encryption and application security management, the BIG-IP hardens your site against a wide class of attacks.

Powerful Application Security

The BIG-IP performs deep packet inspection of the entire application payload to provide powerful application-level security. Through its flexible feature set and unmatched capabilities, it provides administrators with powerful tools to manage and control their application traffic. Comprehensive authentication, authorization, auditing, and application traffic inspection features enforce your security policies at the edge of the network, during session initiation, through session persistence and completion. Features include:

- **Application Security Module** – The Application Security Module is a software add-on that extends the BIG-IP system, turning it into an enterprise-class web application firewall, providing comprehensive, proactive, application-layer protection against both generalized and targeted attacks. Utilizing a positive security model (deny all unless allowed), the Application Security Module permits only valid and authorized application transactions, while automatically protecting critical web applications from attacks. The BIG-IP and the Application Security Module protect against application, infrastructure, and network attacks, such as cross-site scripting, SQL injection, cookie/session poisoning, parameter tampering, forceful browsing, application platform exploits, and zero-day attacks. The Application Security Module protects against entire classes of HTTP and HTTPS-based threats (both known and unknown) rather than only guarding against a limited list of known attacks.

- **Identity Theft Protection** – The Application Security Module gives you the ability to scrub social security numbers, credit card numbers, account numbers, or any other identifiable text patterns from HTTP responses, protecting sensitive information from disclosure and preventing subsequent damage to a business' reputation. Employing the reverse proxy features of the BIG-IP system and bi-directional TCP control, the Application Security Module can inspect and protect information coming into the web servers through HTTP requests and going out to users through HTTP responses. This enables the Application Security Module (and the BIG-IP via iRules) to analyze all data flowing through an HTTP session, down to the content, variable and value data, session management information (such as uniquely generated user IDs and cookies), and basic server response data.

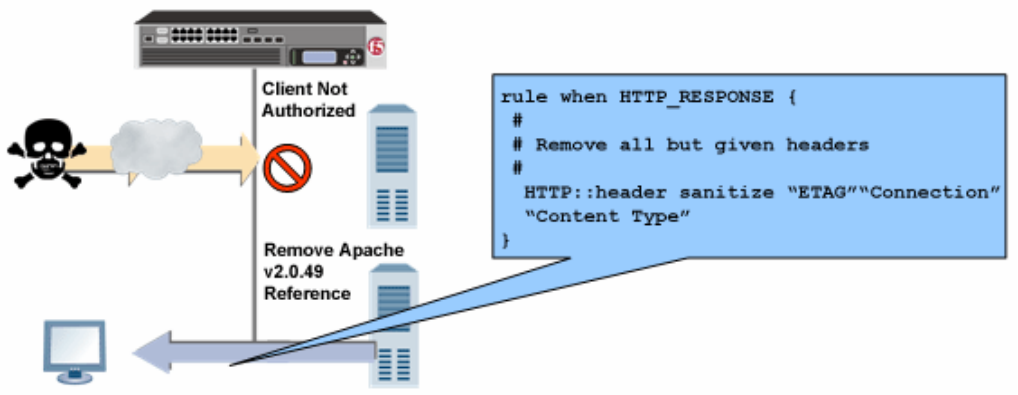

Figure 1: *The iRule scrubs sensitive header information.*

- **Hardened Application Protection** – The BIG-IP with the Application Security Module is a hardened appliance that sits in front of your servers, protecting them from attacks and ensuring that only valid requests get through and valid responses are returned. This scheme protects your site from the standard request-based attacks for hijacking and defacement. Even the most sophisticated attacks can be efficiently identified, isolated, and eliminated without producing any negative effect on your site's performance or impacting legitimate application transactions. Once you define a positive security model, legitimate traffic can pass to and from the network while illegitimate and unknown application traffic is blocked before it's injected into your application network.

- **iRules™** – You can use the BIG-IP to set up and enforce common application-level security policies. In addition, you can use iRules to:

  – Examine application traffic (HTTP, HTTPS, web services)
  – Filter applicable application traffic through the Application Security Module
  – Block application-level attacks and threats

  iRules have the capability to make traffic decisions and modify content on both inbound and outbound traffic. The Universal Inspection Engine enables the BIG-IP to inspect the entire application payload and make decisions about switch, persist, or deny, based on consecutive flows. You can take advantage of this powerful capability using iRules, F5's TCL-based programming language, to create a policy that aligns with your corporate security guidelines. Once created, you can assign iRules to profiles for easy implementation and repeatability.

- **Authentication and Authorization** – The BIG-IP supports Pluggable Authentication Managers that let you choose the scheme (LDAP, RADIUS, TACACS+, SSL Client certification LDAP, OCSP) to authenticate and authorize a remote system when making application requests that pass through the BIG-IP. The BIG-IP routes remote authentication traffic through a Traffic Management Microkernel switch interface that is associated with a VLAN and a self IP address rather than routing traffic through the management interface. This way if the Traffic Management Microkernel service stops for any reason, remote authentication is not available until the service is running again.

- **Cookie Encryption and Authentication** – Inspect, alter, encrypt, and authenticate cookies used in application traffic to prevent hackers from exploiting cookies to launch application attacks. With cookie encryption and authentication enabled, hackers cannot read cookies to access information like JSessionIDs and user IDs that can be used later to modify a cookie and establish an illegal session. Stateful applications are protected from session hijacking and cookie tampering, which exploit critical application vulnerabilities by rewriting the content of a cookie.

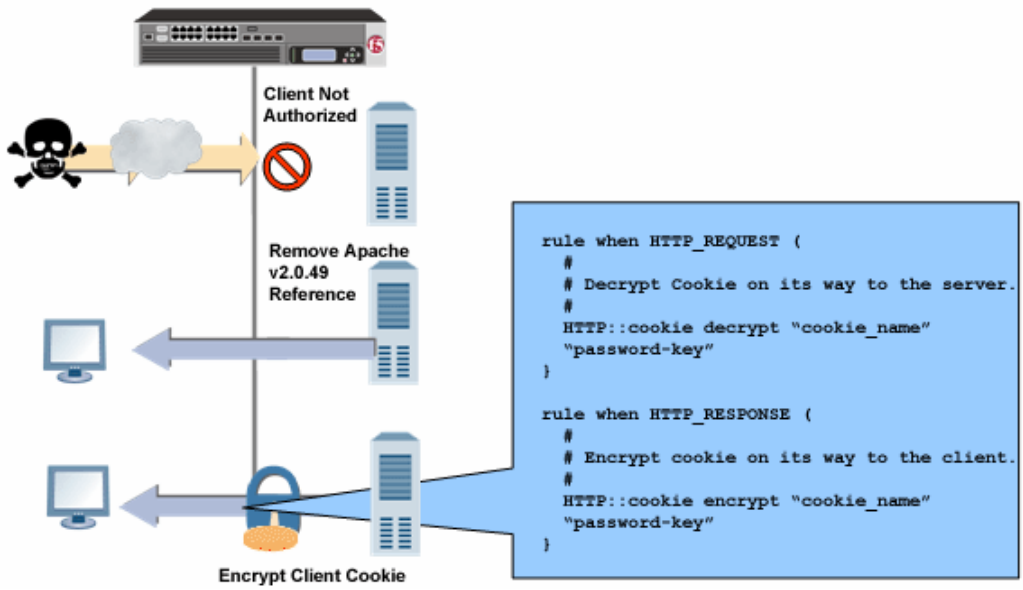

Figure 2: *Encrypt/Decrypt Cookies*

Increasing Network and Infrastructure Security

The BIG-IP further enhances an organization's security with powerful network-layer security to protect against the heaviest of attacks. The BIG-IP system, with its Universal Inspection Engine and iRules, gives you complete visibility into the network payload so that you can intelligently manage and implement your security policies. Protection against common network attacks, DoS and DDoS attacks, and protocol tampering attacks, combined with the BIG-IP's packet filtering capabilities, give organizations unprecedented security to boost productivity and revenue and lower TCO. The BIG-IP does this with the following features:

- **Deny-by-default** – The BIG-IP is a deny-by-default device. All traffic is denied, except for those traffic types you identify. This gives you extremely tight security because you control the traffic that is allowed to pass through the BIG-IP.

- **Automatic defense** – There are numerous built-in processes that enable the BIG-IP to protect your network against common attack types. The BIG-IP ignores directed subnet broadcasts and does not respond to broadcast ICMP echoes used to initiate Smurf and Fraggle attacks. The BIG-IP's connection table matches existing connections so that spoofed connections, such as in a LAN attack, are not passed to the servers. The BIG-IP checks for proper frame alignment to protect against common fragmentation attacks such as Teardrop, Boink, Bonk, and Nestea. Other threats, such as WinNuke, Sub7, and Back Orifice, are denied through the default blockage of ports. With BIG-IP's capability to reassemble overlapping TCP segments and IP fragments, organizations can thwart unknown attacks that are becoming more prevalent these days.

- **SYN CHECK™** – One type of DoS attack is known as a SYN flood, where the attack is launched for the purpose of exhausting a system's resources, leaving it unable to establish legitimate connections. The BIG-IP's SYN CHECK feature works to alleviate SYN floods by sending cookies to the requesting client on the server's behalf, and by not recording state information for connections that have not completed the initial TCP handshake. This unique feature ensures that servers

only process legitimate connections and the BIG-IP SYN queue is not exhausted, enabling normal TCP communications to continue. The SYN CHECK feature complements the BIG-IP's Dynamic Reaping feature that handles established connection flooding. SYN CHECK addresses embryonic connection flooding to prevent the SYN queue from becoming exhausted. Working in conjunction with a high-performance syn-cache, SYN CHECK enables you to use syncookies without the loss of TCP options.

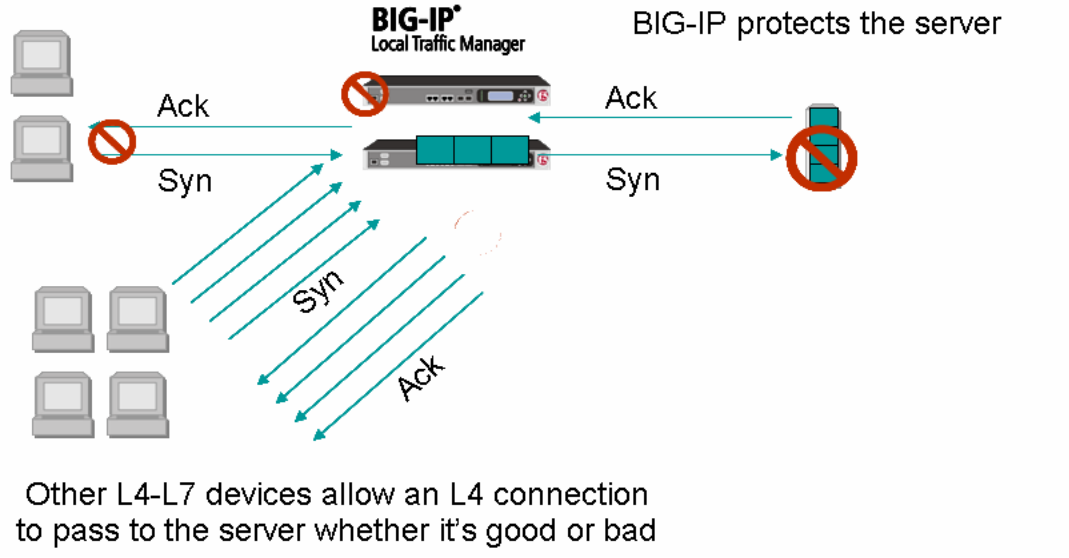

Figure 3: *SYN CHECK example*

- **DoS and Dynamic Reaping** – The BIG-IP has two global settings that you can use to adaptively reap connections. To prevent DoS attacks, you can specify a low-watermark threshold and a high-watermark threshold for reaping connections. The low-watermark threshold determines at what point adaptive reaping becomes more aggressive at reaping connections that are close to their defined timeouts. The high-watermark threshold determines when non-established connections through the BIG-IP are no longer allowed. The value of this variable represents a percentage of memory utilization. Once memory utilization has reached this percentage, connections are not allowed until the available memory has been reduced to the low-watermark threshold.

- **Connection Limits on Virtual Servers** – With BIG-IP, you can limit the maximum number of concurrent connections to a virtual server. This provides another layer of defense against usage attacks, such as DoS.

- **Protocol Sanitization** – With this feature, you can protect your network from hackers that use IP protocol tampering to launch attacks that can overwhelm server resources and bring down a site. By being first in the line of defense and terminating all TCP connections between client and server, the BIG-IP can thwart attacks like out-of-order packet floods, MSS tiny packet floods, and TCP window tampering. Coupled with the Application Security Module, the BIG-IP extends that sanitization to include HTTP header evaluation, RFC violation matches, protocol enforcement, and URI meta character and character set evaluation. The BIG-IP sanitizes client/server communications, looking for attack patterns and exceptions, and cleans the traffic for server and application consumption.

- **Packet Filtering** – The BIG-IP's enhanced packet filter engine provides deep packet inspection and enables you to accept, discard, or reject (send back to the source with various codes like "administratively prohibited") traffic based on advanced packet filter rules. Packet filter rules enforce L4 filtering to allow trusted traffic and handle other specific traffic types according to security policies. You can also use the packet filter with IPv4 or IPv6 addresses to provide basic firewall capabilities, adding another layer of security. The filtering is based on the source or destination IP address of the packet, the source or destination port number (for protocols that support ports), or packet type (UDP, TCP, ICMP). Packet filtering also protects against IP spoofing and bogus TCP flag attacks.

- **Auditing and Logging** – The BIG-IP logs events related to packets discarded due to exceptional circumstances or invalid parameters such as Land attacks, Smurf attacks, bad checksums, unhandled IP protocol numbers or versions, etc. The BIG-IP's security report identifies services and ports that receive unauthorized access attempts by monitoring the source IP address of the attempt, the port used, and the frequency of the attempt. This information is useful for identifying holes in the security network, and determining the source of attacks. In addition to new iRule functions and variables designed for universal content switching, the syntax has been expanded to include two new iRule statements: log and accumulate. Through the use of these capabilities, you can use an iRule to invoke log entries and send real-time alerts.

- **Rate Shaping** – This new feature gives you a powerful and flexible way to defend against bandwidth-abusing attacks. Using rate classes in conjunction with a rate filter, you can protect your network and applications from traffic spikes and regulate abusive users or network attacks from overwhelming network resources. With Rate Shaping, you can specify traffic and application bandwidth limits, and control usage on a per-connection basis at which those resources are allowed to spike or burst to thwart security attacks that attempt to overwhelm network resources.

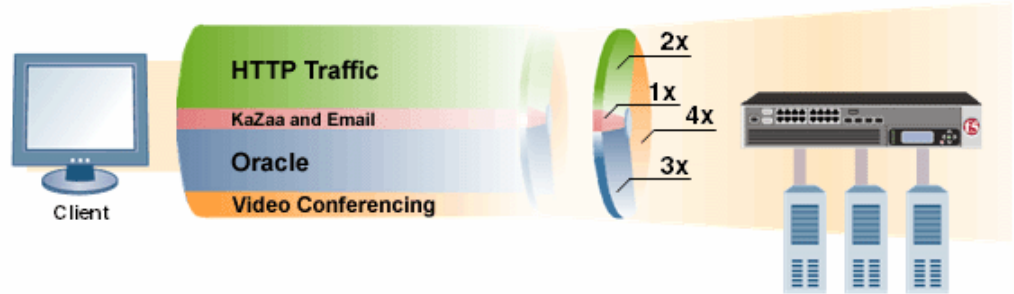

Figure 4: *Rate Shaping controls bandwidth usage*

Protecting Information Leaks

The BIG-IP with the Application Security Module protects an organization from losing valuable information that can be acquired by exploiting security vulnerabilities. Hackers are notorious for scanning or profiling web sites for clues about the infrastructure. These users analyze traffic patterns and examine error codes for vulnerabilities they can use to launch an attack. Internal security breaches are also becoming a common problem, as malicious users on the inside of the network try to get access to sensitive data for illegal use. The BIG-IP blocks access to sensitive data and critical information, and can selectively encrypt traffic when needed with the following features:

- **Resource Cloaking** – The BIG-IP protects your network against hackers that scan sites or applications, looking for clues about possible vulnerabilities. With Resource

Cloaking, the BIG-IP can remove sensitive information about servers and applications in error codes, source code comments on web pages, and server headers. Using the Universal Inspection Engine and iRules, the BIG-IP can block/filter any sensitive information about a site to give you unparalleled protection from malicious users.

- **Secure Network Address Translation (NAT) and Port Mapping** – The BIG-IP can translate the addresses and ports used by the devices behind it, to addresses and ports that are advertised to the outside world. By translating these addresses, you never expose the resources behind the BIG-IP to the outside world. This reduces the risk of an attacker gaining access to your servers. The BIG-IP uses Intelligent Secure Network Address Translation (Intelligent SNAT), which, like NAT, gives servers that have non-public IP addresses a new public IP address for secure, outbound connectivity to the Internet. However, intelligent SNAT differs from NAT by allowing you to assign or map a single IP address to a group of nodes, an entire subnet, or VLAN. Intelligent SNAT can also map an IP address based on any part of the IP packet data. Using the BIG-IP's Universal Inspection Engine, you can map IP addresses based on any part of the IP packet data. By default, SNAT addresses can only be used to initiate outbound connections; whereas, the BIG-IP automatically blocks inbound-initiated connections directed toward SNAT addresses to provide an additional layer of security.

- **Selective Content Encryption** – The BIG-IP can selectively encrypt data based on your application security policies and regulatory standards. With the BIG-IP, you can protect sensitive application data, construct a unified security policy, and bridge together policies from different constituencies in one centralized location. The BIG-IP gives you the ability to pass non-critical traffic and selectively encrypt sensitive traffic such as account numbers and passwords, to comply with regulatory security standards like Sarbanes-Oxley, HIPAA, and FIPS.

Extending Existing Security Solutions

Firewalls, Intrusion Detection Systems (IDS), and VPN devices are the first line of defense against security threats both inside and outside an organization. The significant role these devices play in network security demands that they are available, functioning properly, and responding at all times. Adding the BIG-IP to your existing security infrastructure enables you to extend your security range with scalability and high availability. This is achieved through the use of the BIG-IP system's advanced features, which were developed to support the requirements of an organization's security infrastructure. These features include:

- **Advanced Load Balancing Algorithms** – The BIG-IP offers a wide variety of load balancing algorithms, which are especially useful for load balancing security devices such as Firewalls, VPNs, or IDS systems. Advanced algorithms, such as Predictive, Observed, and Dynamic Ratio, take into account one or more dynamic factors, such as current connection count. For load balancing devices that significantly differ in processing speed, memory, and connection types, these algorithms provide a better and more uniform utilization of resources.

- **Advanced Health Checking** –Transparent health checking capabilities enable you to check an aliased destination through a transparent node. In transparent mode, the monitor checks through the node with which it is associated (usually a firewall) to a destination node. For example, if there are two firewalls in a load balancing pool, the origin server or the BIG-IP device is used as a destination node through the specified firewall. If there is no response from the destination node, the firewall, not the origin server, is marked down and traffic is redirected to a healthy resource.

- **Extended Application Verification** – This feature is used to increase the accuracy of transparent health checking. The Extended Application Verification feature verifies an application on a node by running that application remotely. If the application does not respond as expected in a predetermined amount of time, the BIG-IP directs requests to a healthy resource.

- **Advanced Persistence** – When load balancing client connections through an array of security devices, it is extremely important that all packets with the same session ID get directed to the same device. The BIG-IP provides a variety of persistence mechanisms to ensure that all connections that have the same session ID persist to the same node, and are not load balanced. Universal Persistence can persist on any part of the application payload.

- **Any IP Traffic** – This feature enables the BIG-IP to load balance protocols other than TCP and UDP. For example, you can use this feature when defining a virtual server that is associated with a pool of VPN devices to load balance IPSec traffic.

- **Dynamic Connection Rebind** – The BIG-IP offers dynamic connection rebind for security devices that share their session tables with other devices that are a part of a cluster. If one member of the cluster fails, dynamic connection rebind moves all connections from the failed node to a healthy node within the same pool. Because the other nodes share their session table, the newly-selected node is able to authenticate and permit the existing connection without interruption or intervention by the user.

- **Last Hop Pools** – Using last hop pools when load balancing security devices ensure that the path a response connection takes (from the resource to the client) is the same as the path it took for the original request (from the client to the resource). The BIG-IP lets you manually specify the last hop pool member or allows the system to automatically determine the last hop.

- **Clone Pools** – The BIG-IP's enhanced clone pool feature replicates all traffic being handled by a pool to a clone pool that contains an IDS or a probe device. You can configure a clone pool for any standard load balancing pool. When a standard load balancing pool receives a connection, it picks a node for the regular connection using the regular pool, and then picks a clone node from the clone pool. The clone node receives a copy of all the traffic going through the regular pool.

- **Client-side SSL Proxy** – BIG-IP's client-side SSL proxy feature terminates SSL connections, decrypts the request, and sends the request in clear text to its final destination. During the process of terminating an SSL connection, the proxy performs all of the certificate verification functions normally handled by the target web server, as well as encryption and decryption functions. The BIG-IP includes hardware that accelerates these operations, enabling it to offload this task for large volumes of traffic in an efficient manner. When used in combination with clone pools, this feature extends the effectiveness of IDS devices that would otherwise not be able to process encrypted data.

- **Integration control** – With iControl and the Universal Inspection Engine, other security devices can inject their knowledge into the BIG-IP system by creating, deleting, or editing iRules. Changes are instantaneously applied via the iControl API for fast, proactive action. You can use this functionality to secure web services, mobile applications, and nearly any IP-based application. Coupled with iRules, iControl, and the Application Security Module, the BIG-IP provides a single, integrated solution to protect your network and application infrastructure.

About F5

F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast, and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 9,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information, go to www.f5.com.