



Mitigating DDoS Attacks with F5 Technology

Distributed denial-of-service attacks may be organized by type into a taxonomy that includes network attacks (layers 3 and 4), session attacks (layers 5 and 6), application attacks (layer 7), and business logic attacks. Each type may be matched with the best F5 technology for mitigating that attack. Taken together, the F5 BIG-IP portfolio of products provides effective anti-attack technology for each layer of the taxonomy and can also defend against specific attack tools, network reconnaissance, and low-bandwidth asymmetric attacks.



Introduction

Distributed denial-of-service (DDoS) attack types have moved up the OSI network model over time, climbing from network attacks in the 1990s to session attacks and application layer attacks today. Network attacks include DDoS variants such as SYN floods, connection floods, or ICMP fragmentation. Session attacks, which target layers 5 and 6, include DNS and SSL attacks. Application attacks at layer 7 represent approximately half of all attacks today. Finally, though layer 7 tops the OSI model, attacks are now moving into business logic, which often exists as a layer above the OSI model. But even with these changes in the current threat spectrum, organizations must continue to defend against network and session attacks, too.

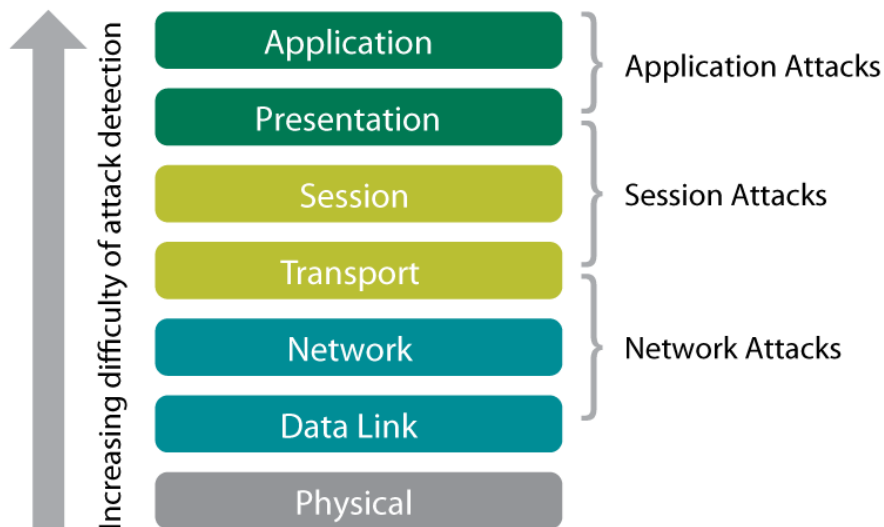


Figure 1: DDoS attacks target many layers of the OSI network model.

Meanwhile, the Application Delivery Controller (ADC) has become a strategic point of control in the network. ADCs can be both network- and application-aware and can be managed by network, application, and security teams. Over time, ADCs have evolved into flexible, high-performance components of the network that can offload services such as load-balancing, caching, and acceleration to save organizations both CapEx and OpEx. In addition to consolidating performance, scalability, and flexibility solutions into a single platform immediately in front of web services, the ADC becomes a logical defensive position against both DDoS attacks and targeted application-layer attacks.




WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

Having occupied this position in many large enterprises and data centers for years, F5 ADC technologies have evolved to mitigate attacks targeting not only the network but also the application and business logic levels.

F5 solutions can securely deliver applications while protecting the network, the session, and the user. Specific F5 mitigation technologies map directly to individual DDoS attacks for the network, the session, the application, and business logic. Effective security solutions involve core F5 products such as F5 BIG-IP Local Traffic Manager (LTM) and BIG-IP Global Traffic Manager (GTM) as well as the new BIG-IP Advanced Firewall Manager (AFM). F5 iRules, a programmatic scripting language, can be easily adapted as a final, flexible security defense. Lastly, security products like BIG-IP Application Security Manager (ASM), F5's web application firewall module, can block the most sophisticated attacks in the DDoS threat spectrum.



	Attack	F5 Mitigation Technology
Application	OWASP Top 10 (SQL injection, XSS, CSRF, etc.), Slowloris, Slow POST, HashDos, GET floods	BIG-IP ASM: Positive and negative policy reinforcement, iRules, full proxy for HTTP, server performance anomaly detection
Session	DNS UDP floods, DNS query floods, DNS NXDOMAIN floods, SSL floods, SSL renegotiation	BIG-IP LTM and BIG-IP GTM: High scale performance, DNS Express, SSL termination, iRules, SSL renegotiation validation
Network	SYN floods, connection floods, UDP floods, PUSH and ACK floods, teardrop, ICMP floods, ping floods, and smurf attacks	BIG-IP AFM: SYN Check, default-deny posture, high-capacity connection table, full proxy traffic visibility, rate limiting, strict TCP forwarding

Figure 2: Today's attacks are moving up the OSI stack.



Mitigating Network Attacks

The most basic network attacks attempt to overwhelm a defensive device with sheer volumes of traffic. Sometimes these volumetric attacks are designed to overload the connections-per-second (CPS) capacity (e.g., the ramp-up rate). Another, slightly more sophisticated attack method is to establish many legitimate connections (a connection flood) to overwhelm the memory of any stateful defensive devices so they lose the ability to accept legitimate connections. Attacks of both kinds are mitigated by the full-proxy position of BIG-IP LTM and its underlying F5 TMOS architecture, which deliver the intelligence to distinguish between legitimate and malicious connections plus the capability to either absorb or drop the malicious ones before they consume network resources behind the device.

There are three key technologies within BIG-IP LTM that deliver its network defense functionality: F5 Packet Velocity Accelerator (PVA) processing, a full-proxy architecture, and protocol validation. Memory management and custom configuration complement this trio of technologies to help organizations repel attacks.

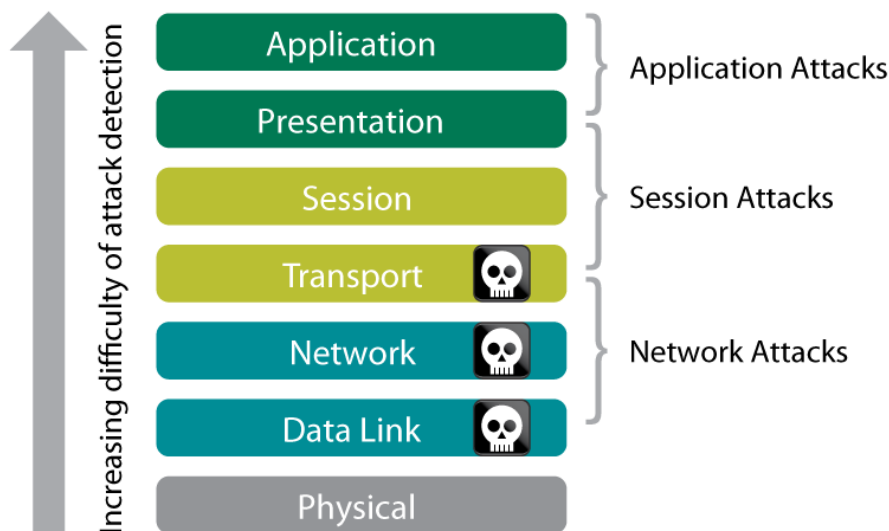


Figure 3: Network attacks target layers 2 through 4.

PVA Processing

The PVA is a purpose-built, customized hardware processor that assists BIG-IP LTM to scale by an order of magnitude above software-only solutions. PVA technology is fully session-aware and contains mitigation code for common network attacks such as SYN floods.



Full-Proxy Architecture

Solutions built atop a full-proxy architecture can be active security agents because their architecture makes them part of the flow of traffic, not simply devices sampling that traffic. Products that are full proxies provide inherently better security because they actively terminate the flow of data, essentially creating an “air gap” security model inside the product.

With full proxies like BIG-IP LTM, traffic coming from the client can be examined before it is sent on its way to the application tier, ensuring that malicious traffic never passes the proxy barrier. Traffic returning from the server can be fully examined before it is deemed acceptable to pass back to the client, thereby ensuring that sensitive data such as credit card or Social Security numbers are never passed across the proxy barrier.

Protocol Validation

A third method of network attack involves sending malformed data, such as packets with invalid combinations of flags or incomplete fragments. These attacks can be very effective because they tie up the CPU or memory of devices that examine them. Often the number of CPU cycles spent defending the packet dwarfs the processing that it takes to launch the packet, leading this method to be known as an asymmetric attack.

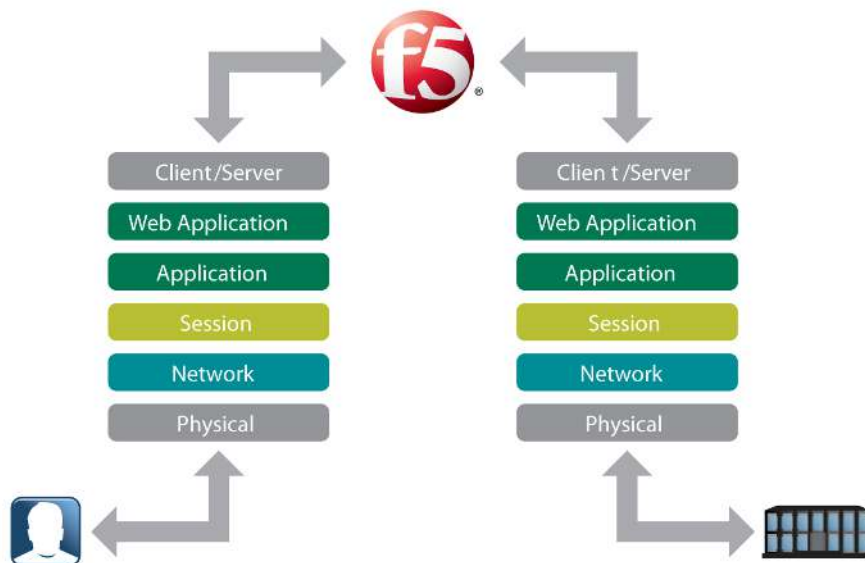


Figure 4: Because they actively terminate the flow of data, F5 full-proxy solutions provide inherently better security.



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

Such invalid data and asymmetric attacks are mitigated by the protocol validation technology of BIG-IP products. In protocol validation, the ADC understands the expected network protocol of traffic destined for each application and can discard malformed traffic before it penetrates deeper into the data center.

Repelling Specific Network Attacks

Network attacks, which have been around a long time, have evolved with impressive longevity and variety. The BIG-IP product family mitigates a long list of network attack types, most through built-in technologies or in default configurations.

SYN floods

An old attack and the most common network DDoS attack, the SYN flood exploits the three-way handshake of the TCP setup. Any device, including a firewall, that terminates TCP is susceptible to the SYN flood attack unless specific measures are taken to defend against it. Conventional firewalls mitigate this attack using different technologies and with varying rates of success.

Over time, three main mitigation techniques have evolved to combat SYN floods. The SYN proxy defense, found on many modern pass-through firewalls, stalls TCP connections to filter out invalid ones. The drawback to this approach is that it only forestalls the problem and makes the firewall itself vulnerable to larger SYN floods.

A second mitigation approach is the SYN cache. This technology, found mostly on server platforms, relies on optimized memory tables to scale more connections. Results have been mixed, and the SYN cache approach is losing market traction.

The third and best mitigation technique is called the SYN cookie approach. SYN cookies are encrypted sequence numbers that allow a defending device to filter out invalid sessions without consuming any state information.

The SYN cookie approach underlies the F5 SYN Check feature. The majority of F5 devices include the PVA technology, either as an ASIC chip or set of field-programmable gate arrays (FPGAs). For hardware-accelerated virtual servers, the PVA is the first line of defense against SYN floods. When a SYN flood is detected, the PVA turns on its SYN Check feature to prevent invalid sessions from getting past the PVA to the servers behind it.

BIG-IP virtual editions (or any configuration that cannot take advantage of the hardware-assisted PVA technology) also benefit from the SYN Check feature. The high-performance traffic management microkernel in the TMOS platform contains a software version of SYN Check that uses high- and low-water marks to control the encrypted-cookie gating mechanism.



Connection floods

Another old, yet still common, attack is the TCP connection flood. This DDoS variant consumes connection table resources for any stateful device between the perimeter and the target servers. The full-proxy nature of BIG-IP LTM protects data center resources by accepting the DDoS connections and then using memory management, via its high-capacity connection table and aggressive connection reaping, to soak up connection floods before they reach server resources.

UDP floods

The key to fast denial of UDP floods historically has been the default-deny security posture. BIG-IP LTM provides this posture for the data plane. Any packets that do not match a defined virtual server are dropped as quickly as possible, thus mitigating UDP floods. No UDP packets ever reach HTTP-based applications behind a BIG-IP device.

Fake sessions

The fake TCP session is a clever attack that often passes through conventional firewalls. It contains not just a proper-looking SYN packet but also a series of fake TCP payload packets and even a closing packet. When BIG-IP LTM is in place, fake session packets sent at a high volume are filtered out by its built-in SYN Check feature, while fake session attacks sent at low volume have their connections dropped when the ADC rejects the invalid sequence numbers.

PUSH floods and ACK floods

A full-proxy ADC can mitigate PUSH and ACK floods. Because BIG-IP LTM is part of every conversation between every client and every server, it can recognize packets that do not belong to any valid flow, such as typical PUSH and ACK flood packets. These are dropped quickly and never pass beyond the ADC.

ICMP floods, ping floods, and smurf attacks

One of the few layer 3 attacks still in use today is the ICMP flood. Often these floods are triggered by amplifying ICMP echo replies from a separate (and often innocent) network to a target host. BIG-IP LTM mitigates ICMP floods by limiting the rates of all ICMP traffic and then dropping all ICMP packets beyond the limit. The limit is adjustable by the operator.



Ping of death ICMP attacks

The ping of death attack uses overly large ICMP packets to reboot vulnerable servers. These packets are denied at the BIG-IP LTM ADC in its default configuration. Only if an operator enables the device's ANY IP feature for the target virtual server will the ADC allow these fragmented ICMP packets into the enterprise or data center.

Christmas tree attacks

A Christmas tree packet is one that is “gifted” with all of the possible TCP flags enabled (such as SYN and RST, which is an illegal combination). Older devices become confused by the packets, which leads to unpredictable behavior. When BIG-IP LTM's strict TCP forwarding option is configured, it rejects Christmas tree packets.

LAND attacks

Local Area Network Denial (LAND) attacks use incoming packets whose source address is spoofed to match the ADC itself. BIG-IP LTM checks specifically for LAND attack packets and quickly drops them.

Teardrop attacks

The teardrop attack exploits an overlapping IP fragment problem in some common operating systems. It causes the TCP reassembly code to improperly handle overlapping IP fragments. In its default configuration, the BIG-IP system handles these attacks by correctly checking frame alignment and discarding improperly aligned fragments. Teardrop packets are dropped and the attacks are mitigated before the packets can pass into the data center.

Layer 4 Security Management and Visibility

In addition to readily defeating these common network attacks, the BIG-IP product family includes BIG-IP AFM, which enables security teams to manage security rule sets in the same way they might manage conventional firewall rules. Security administrators can use the BIG-IP AFM point-and-click interface to drop, allow, or log incoming traffic using at the network level. BIG-IP AFM tracks 38 different types of network DDoS attacks (including all of those mentioned above) and reports on each. Organizations can also define the parameters of their own attack detection signatures and be alerted when thresholds for these are passed.



Mitigating Session Attacks

Session attacks, which take place at layers 5 and 6, include DNS and SSL attacks. Conventional firewalls have no ability to mitigate SSL attacks and offer only limited defensive value for DNS attacks.




Figure 5: Session attacks typically defeat conventional firewalls.

F5 products can defend DNS and SSL resources against session- and presentation-level attacks. For both, the key to the defense is the high-performance, full-proxy F5 functionality that validates and shapes every DNS and SSL connection between the Internet and the data center.

The security services offered in BIG-IP GTM provide protection against DDoS attacks at the DNS security perimeter. BIG-IP LTM protects SSL resources by offloading SSL processing onto its high-performance, high-capacity hardware and through judicious use of iRules.

Mitigating Specific Session Attacks

Session attacks can be defeated through a combination of hardware capacity and technologies such as the proprietary features of F5 ADCs.



DNS UDP floods

Normal DNS servers cannot withstand a typical distributed UDP flood. BIG-IP GTM mitigates UDP floods by scaling performance far beyond that of a normal DNS server. Since version 11.0, the full-proxy BIG-IP GTM validates each and every DNS request packet and discards those that are invalid (such as packets from a UDP flood).

DNS query floods

A more advanced DNS attack is the query attack, in which multiple clients flood the target with valid DNS requests, attempting to overload it. The F5 DNS Express feature in BIG-IP GTM can mitigate these attacks by enabling multi-core, linear scaling. DNS Express further protects the perimeter by handling all valid and invalid DNS requests itself, at a capacity up to an order of magnitude greater than a typical DNS server.

DNS NXDOMAIN floods

One of the most sophisticated DNS attacks is the NXDOMAIN query flood, which is designed to foil DNS caches and bring down DNS servers. It works by causing DNS servers to spend their time looking for thousands or millions of nonexistent host records. DNS Express is ideally suited to help an organization survive an NXDOMAIN flood because it retains all the valid organization zone information even during the flood.

SSL floods

Organizations are starting to see more malicious floods of SSL connections coming into their data centers. These SSL floods bypass firewalls, intrusion prevention system (IPS) perimeters, and cloud scrubbers, and they can take down server resources or overflow stateful firewalls. By terminating SSL at a capable ADC, an organization can stop SSL floods. A full proxy for SSL processing, such as BIG-IP LTM, simply drops malicious or empty SSL connections, protecting the resources behind it.



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

SSL renegotiation attacks

The notorious SSL renegotiation attack was discovered when it was initially launched against an F5 customer, who was then assisted by the F5 field services team. What makes this attack so effective is that it exploits the asymmetric encryption property of SSL, so the attacker needs only one-tenth of the computational power of the unprotected server. Still, the high capacity and performance of F5 hardware for SSL cryptographic offloading means that an SSL renegotiation attack has to be extremely strong to overcome a BIG-IP device.

The original attack was repelled by a simple iRule now published on the F5 DevCentral online community. F5 still has one of the only solutions to this thorny protocol attack. The premise of the iRule is that if a client connection attempts to renegotiate more than five times in any 60-second period, that client connection is silently dropped.

One of the benefits of this iRule and its silent work is that it fools the attacker into thinking the connection is merely stalled, fully negating the attack.

If an organization is a frequent target or handles traffic that is a primarily SSL traffic, the following iRule can be deployed at every virtual server that requires protection.

```
when RULE_INIT {
    set static::maxquery 5
    set static::mseconds 60000
}
when CLIENT_ACCEPTED {
    set ssl_hs_reqs 0
}
when CLIENTSSL_HANDSHAKE {
    incr ssl_hs_reqs after $static::mseconds {
        if {$ssl_hs_reqs > 0}
    } {
        incr ssl_hs_reqs -1
    }
}
if {
    $ssl_hs_reqs > $static::maxquery
} {
    after 5000 log "Handshake attack detected, dropping [IP::client_
addr]:[TCP::client_port]"
    drop
}
}
```



Mitigating Application Attacks

At the top of the OSI stack is the application layer. This is the area where it's most difficult to detect or defend against malicious behavior, and in particular, conventional firewalls provide little defensive value. Consequently, the application layer is being targeted by most of today's attackers.




Figure 6: Application attacks are the most prevalent today.

An application attack is different from a network attack in that it is specific to the application being targeted. Whereas a SYN flood can be launched against an IP address, an application attack will usually exploit properties specific to the victim, such as the repeated downloading of a single PDF file on the website. To lower-level security devices such as firewalls, the attack connections are indistinguishable from normal traffic.

BIG-IP ASM brings together a variety of anti-attack and DDoS prevention technologies specifically designed to mitigate application layer attacks, including the majority of the OWASP Top 10. BIG-IP ASM learns the expected input for every page in the site it protects and generates a security policy to protect that page. Because BIG-IP ASM is application-aware, it can foil application-layer attacks that abuse the application, the database, or the business logic.



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

BIG-IP ASM can distinguish between humans and robots as the sources of traffic and use this information during an attack to block non-human visitors. It can also inject JavaScript redirect code into the stream to foil the majority of botnet slaves while allowing access to legitimate browsers. Finally, BIG-IP ASM can also rate-limit traffic to specific application servers when it detects that an attack may be underway.

Mitigating Specific Application Attacks

Today's DDoS attack tools often use multiple attack vectors, mixing flood types. As attacks against the application layer increasingly grow multi-pronged, they've sometimes earned the name diverse distributed denial-of-service (3DoS) attacks. Whether they use high- or low-bandwidth approaches or both, these attacks can be very difficult to identify and defeat.

A solution that can provide early warning about the attack vectors and defend against multiple, simultaneous vectors is therefore the most effective. The combination of BIG-IP LTM, appropriate iRules, and BIG-IP ASM defeats a large number of application-layer attacks.

OSI Layer	Attack	BIG IP LTM + iRule	BIG-IP ASM
Application (Layers 6–7)	Slowloris (Nuclear DDoSer, Slowhttptest)	✓	✓
	Keep-Dead	✓	✓
	Slow POST (R-U-Dead-Yet, Tor Hammer, Nuclear DDoSer, Slowhttptest)	✓	✓
	HashDoS	✓	✓
	Apache Killer (Slowhttptest)	✓	✓
	HTTP GET Flood, Recursive GET Flood (Web Scraping), Dirt Jumper (HTTP Flood)	✓	✓
	#RefRef (exploits SQLi / OWASP Top 10 vulnerability as entry)		✓
	XML Bomb (DTD Attack), XML External Entity DoS		✓

Figure 7: Multiple attack vectors can be defeated by BIG-IP technologies and products working together



Simple GET floods

One of the most common application layer attacks is a GET flood that simply requests static URLs. BIG-IP LTM can mitigate these attacks with an iRule that filters on the requested URL, and BIG-IP ASM can rate-limit requests based on server performance, client requests per IP address, and increases in requests from specific URLs.

Recursive GET floods

Recursive GET floods are GET flood attacks that iterate through the website, retrieving every object that can be requested. Unlike simple GET floods, recursive floods cannot be filtered with a URL-matching iRule.

BIG-IP ASM can mitigate these attacks from a different angle, however, by monitoring the application's response time (which is by itself the most accurate detection method) and then sequentially applying three different countermeasures:

1. A smart JavaScript injection that will verify that the user is indeed using a browser. Most attacking tools are not browser-based, since browsers are not designed to send a lot of requests per second. In addition, this countermeasure can deal even with an attacker using a website behind a proxy without affecting the traffic of legitimate users connecting through the same proxy. In either case, the identified attacker's connection is dropped.
2. If the JavaScript injection doesn't solve the problem, (for example, when it doesn't effect a positive change in latency), then BIG-IP ASM will rate-limit GET requests from even the chattiest IP addresses.
3. If neither the first nor the second countermeasures solves the issue, BIG-IP ASM escalates to rate-limiting per URL.

Malicious POST floods

POST floods are gaining momentum as attackers have figured out that this technique is a good way to get around various intermediaries, such as content delivery networks (CDNs) and caching services. Typically POST floods bypass these and go straight to the origin servers. Sending a POST, which is nearly as easy for a client as sending a GET, has a much greater chance of tying up valuable resources on the origin server.

BIG-IP ASM can use its techniques for identifying human vs. robotic connections to foil POST attacks. As with recursive GET floods, it can also rate-limit based on the URI, server performance, or the number of requests per client.



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

Mitigating Low Bandwidth HTTP Attacks

Low-bandwidth attacks are a specific form of application-layer attack that are often undetectable by conventional means because they use very little incoming bandwidth.

Slowloris attacks

The Slowloris and PyLoris attack tools achieve denial of service by feeding an HTTP header to a server in an extremely slow fashion. Slowloris starts by probing the target service to determine its inactivity timeout—usually about five minutes or 300 seconds. Once the interval is known, Slowloris opens connections that emulate a simple browser and sends a bogus HTTP header just ahead of the timeout (for instance, every 299 seconds):

```
HTTP/1.1
GET /
X: a <299 second pause> X: a <299 second pause> X: a <299 second pa
use>
```

The connections will go on like this forever. When enough of them have engaged a specific web server, that server will no longer have enough connections to accept new requests, resulting in a denial of service.

BIG-IP LTM, as a standard, layer 7, full-proxy virtual server for HTTP, mitigates these attacks in its TMOS high-performance traffic management microkernel or simply dilutes the attack with the PVA. It will never pass along Slowloris and PyLoris requests because it will be waiting for the final double carriage return that marks the end of the headers. Since the attack tools never send that token, BIG-IP LTM does not consider the connections valid. Eventually they will be discarded without ever consuming resources behind the ADC.

For distributed Slowloris attacks, where millions of Slowloris connections may pile up at the BIG-IP device, a Slowloris iRule takes a more proactive approach to dealing with the attack.

Slow POST attacks

The slow POST attack is similar to the Slowloris attack but can only be mitigated with the BIG-IP ASM module. Slow POST works by starting an HTTP POST operation (like an upload) and then feeding the upload data in very slowly:



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

```
HTTP/1.1
POST /target-url
Content-Length: 1048576
Host: a a <pause> b <pause> c <pause>
```

BIG-IP ASM mitigates this and other low-bandwidth attacks by cataloging the performance of each request and then limiting the number of very slow connections per CPU core.

By establishing and enforcing a limit on these kinds of attacks, BIG-IP ASM allows access to legitimate clients with poor connections while defending the resources from malicious overloading.

HashDoS

All major web services platforms (e.g., Java, ASP.NET, and Apache) use the same fast hash algorithm for the dictionary tables. Their reliance on the same hash function made all of these platforms vulnerable to a clever attack released in late 2011 called the HashDoS attack. It worked by sending a single large POST filled with thousands of tailored form variables that overwhelmed the hashing function of any single target server. A single POST message, pre-computed and sent over a 33 K connection by a client as weak as a handset, could tie up a server for over an hour.

BIG-IP LTM mitigates this HashDoS attack through the application of a public iRule that drops any POST that contains an excessive number of form variables or an excessively large payload. By mitigating the problem at the ADC, organizations protect all back-end web server platforms at the same time. BIG-IP ASM mitigates this attack by using a signature and limiting the total number of parameters that can be sent on a single request.

WHITE PAPER

Mitigating DDoS Attacks with F5 Technology




Figure 8: F5 solutions protect all web service platforms against HashDoS attacks.

Mitigating Attacks Using Network Reconnaissance

The most sophisticated of today's attacks are asymmetric attacks designed to tie up the application tier—specifically, the database tier—by sending a flood of legitimate requests that trigger resource-expensive database queries. Before the flood, attackers collect the necessary information using network reconnaissance to crawl a website and measure the return time of each URI. This information can be collected by one party and sold to another, or simply saved for later. Attacks that target high-response-time database queries are very difficult for many vendors to mitigate.

The full proxy application awareness of F5 products prevents these attacks at the network reconnaissance stage. Many network scanners are known by BIG-IP ADCs subscribed in the F5 IP Intelligence service. The scanners cannot switch to anonymous networks, because those are known, too. Requests from identified scanners can therefore be blocked with the application of iRules or by BIG-IP ASM or BIG-IP AFM.

BIG-IP ASM can also reject one-off network scanners, which may be used only once and therefore cannot be globally tracked, by allowing access to the application only by humans with browsers. BIG-IP ASM can complete the defense by enforcing that visitors go through a set of pages sequentially before gaining access to valuable, resource-intensive services.



WHITE PAPER

Mitigating DDoS Attacks with F5 Technology

Mitigating Business Logic Attacks

Although it's not on the OSI model, business logic represents a layer higher than the application layer that is defined by work flow and processes. In a 2011 paper, "[How to Shop for Free Online](#)¹," security researchers demonstrated attacks against e-commerce payment systems by manipulating input to take advantage of security gaps in business logic. Traditionally such attacks are not considered DDoS attacks per se, but they can be similarly automated.

They can also be prevented. By excluding access by bots, BIG-IP ASM prevents web-scraping, brute-force password cracking, and forceful browsing.

Conclusion

Organizations today have to face facts: DDoS attacks are increasing in volume, frequency, and sophistication, and they are targeting every level in the data center. Smart organizations are moving to defend not only their network, session, and application layers, but also their business logic and database tiers as well.

Fortunately, the same set of F5 technologies can mitigate attacks at every layer. At the network level, F5 technology for high-performance, scalable network firewalls protects against layer 3 and 4 DDoS attacks.

As the DDoS threat expands into the session layer, the trend demonstrated by the SSL renegotiation and HashDoS attacks of 2011 will continue and probably accelerate. Those examples made the adaptability of F5 technology clear when support staff provided solutions to these previously unseen vulnerabilities—without requiring new software.

Even higher in the OSI model—in the application layer and in business logic—the full proxy architecture of BIG-IP LTM and BIG-IP GTM and the control of BIG-IP ASM and BIG-IP AFM provide security that no other set of solutions can match. That is why these F5 products are already a critical part of many networks today. With an ideal mix of technologies, application understanding and awareness, a strategic point of control in the Application Delivery Network, and a world class price-to-performance ratio, the F5 platform mitigates the evolving DDoS threat.

¹ Rui Wang, Shuo Chen, XiaoFeng Wang, Shaz Qadeer. Indiana University and Microsoft Research. 2011.

WHITE PAPER

Mitigating DDoS Attacks with F5 Technology



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

©2015 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS01-1943 0113