The F5 DDoS Protection Reference Architecture

F5 offers guidance to security and network architects in designing, deploying, and managing architecture to protect against increasingly sophisticated, application-layer DDoS attacks.

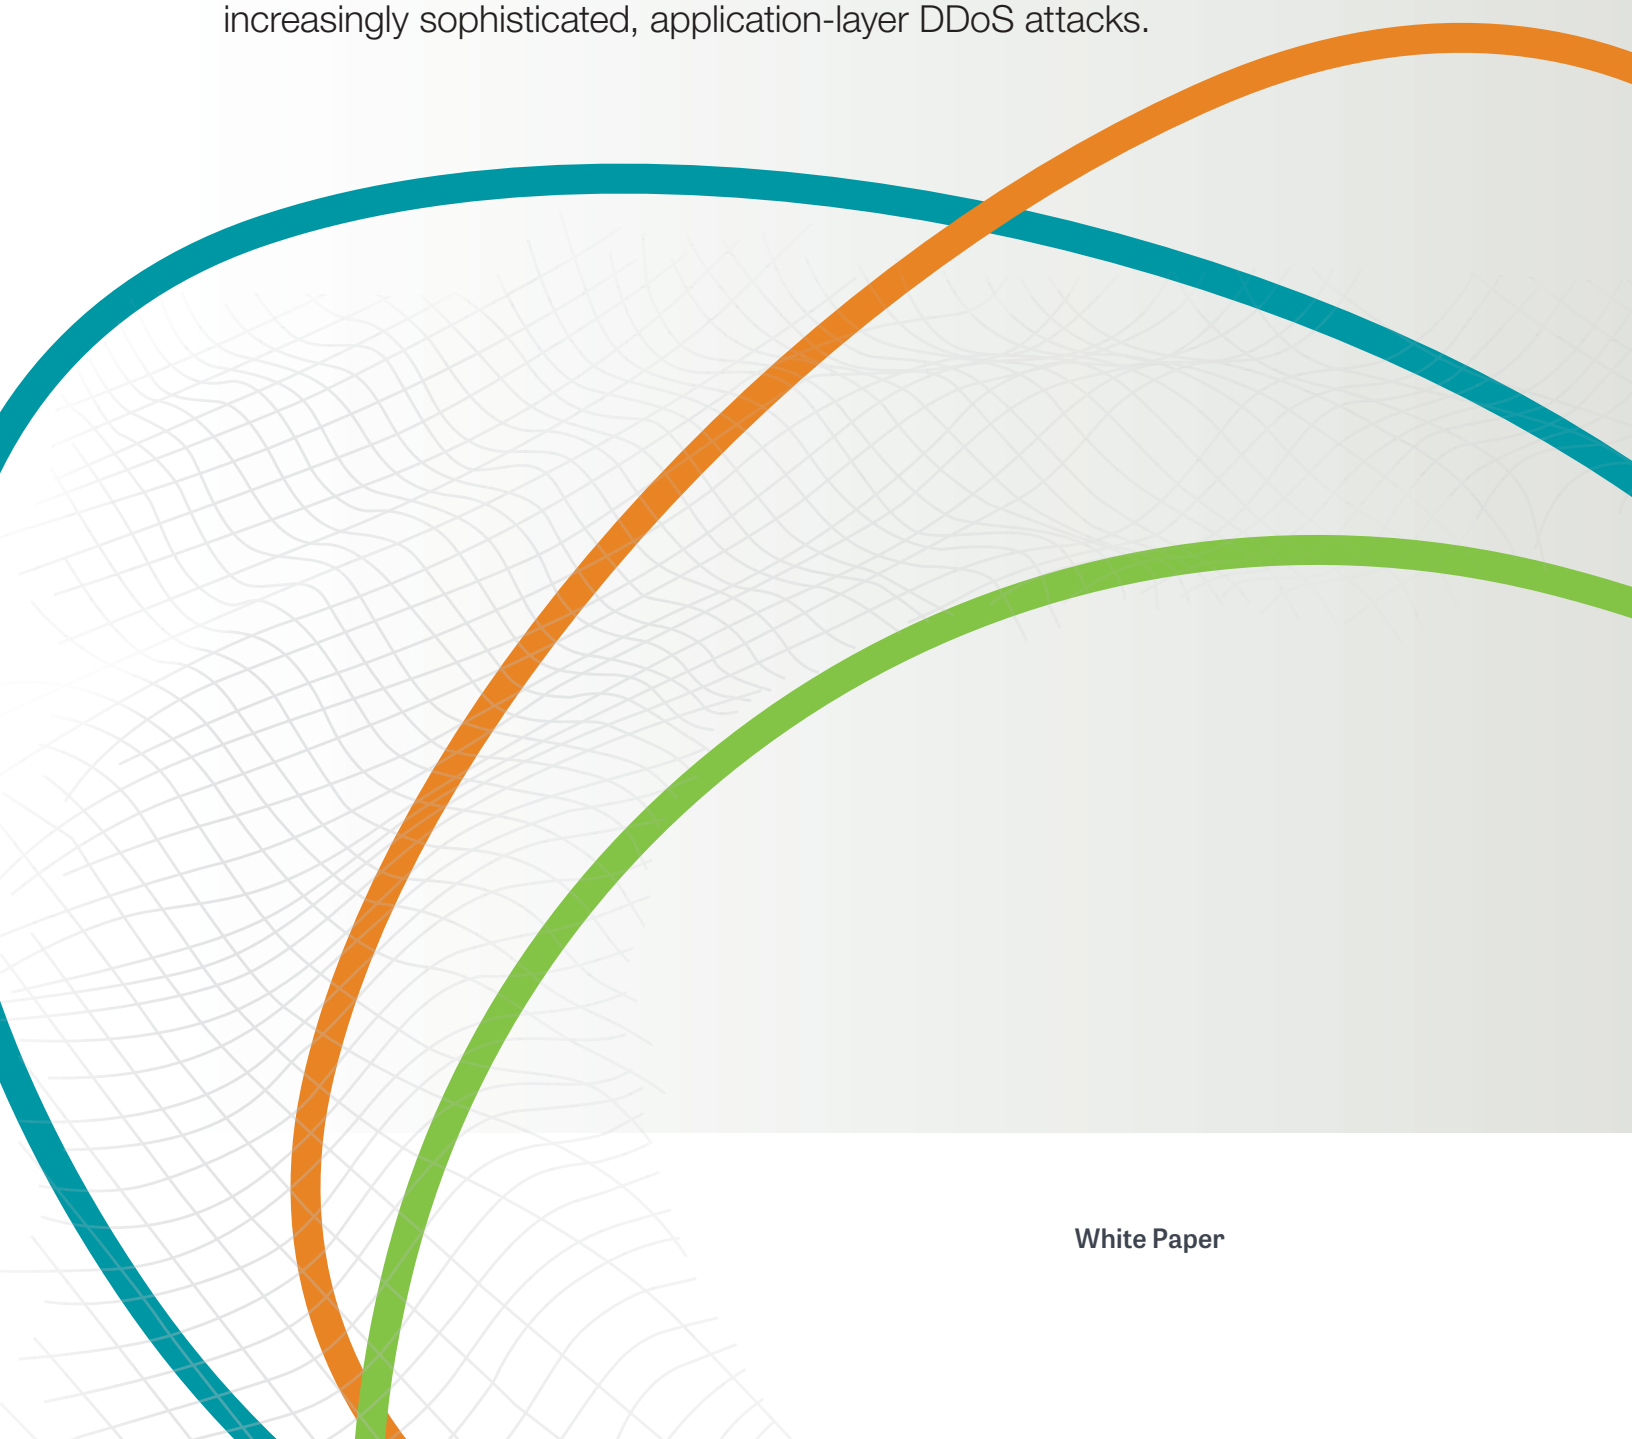

White Paper

Contents

Introduction

Since 2012, a wave of crippling DDoS attacks has pushed large financial customers and enterprises to redesign their networks to include DDoS protection. Working with these customers, F5 has developed a multi-tier DDoS protection architecture. Tier 1 handles DDoS mitigation for DNS and layers 3 and 4. Freed from the noise of the network attacks, tier 2 can use its CPU resources to protect the higher-layer application protocols. This strategy enables organizations to defend against all types of DDoS attacks and is already providing benefits at several F5 customer data centers.

The Four Categories of DDoS

While the DDoS threat landscape is constantly evolving, F5 has found that attacks continue to fall within four attack types: volumetric, asymmetric, computational, and vulnerability-based. These attack categories have the following characteristics:

- Volumetric—Flood-based attacks that can be at layer 3, 4, or 7
- Asymmetric—Attacks designed to invoke timeouts or session-state changes
- Computational—Attacks designed to consume CPU and memory
- Vulnerability-based—Attacks that exploit software vulnerabilities

Defensive mechanisms have evolved to deal with these different categories, and today's high-profile organizations have learned to deploy them in specific arrangements to maximize their security posture. By working with these companies and fine-tuning their components, F5 has developed a recommended DDoS mitigation architecture that can accommodate specific data center size and industry requirements.

Building a DDoS Protection Solution

The following architecture describes a multi-tier DDoS protection system built around well-known industry components. Some of these devices may be provided by other vendors and suppliers (e.g., the firewall and the cloud-based scrubbing service) but some are specific F5 components for which we recommend no other replacement.

Components of a DDoS Protection Architecture

Table 1 shows the mapping of DDoS architecture components to the four DDoS attack categories they mitigate.

Attack Category	Mitigation Component	Recommended Supplier
Volumetric	Cloud-Based Scrubbing Service	Prolexic, Verizon, VeriSign
	Web Application Firewall	F5
Asymmetric	Web Application Firewall	F5
Computational	Application Delivery Controller	F5
	Network Firewall	F5
Vulnerability-Based	IP Reputation Database	Webroot
	Intrusion Prevention/Detection Systems (IDS/IPS)	Sourcefire
	Application Delivery Controller	F5

Table 1: Mapping of DDoS mitigation components to attack types

Cloud-based DDoS scrubbing service

An external, cloud-based DDoS scrubbing service is a critical component of any DDoS mitigation architecture. When an attacker is sending 50 Gbps of data at an organization's 1 Gbps ingress point, no amount of on-premises equipment is going to solve that problem. The issue is akin to too many people trying to get through a doorway at once. The cloud service, hosted either from a true public cloud or within the organization's bandwidth service provider, solves this problem by crudely sorting out the obvious bad from the likely good.

Even though only a fraction of today's attacks are volumetric enough to consume all available ingress bandwidth, there are enough of these attacks to require an agreement with a cloud service supplier and make it a critical part of the overall solution. However, cloud-based scrubbing is not adequate on its own; even cloud service vendors will admit this.

As many of today's organizations have found, hackers know all about these cloud services. They can either leverage this fact by deliberately causing the target to incur the costs of those services, or they can avoid them altogether with application-layer attacks the services do not recognize. Low and slow attacks can be especially effective at evading detection by cloud-based scrubbing services. In addition, these services typically cannot handle encrypted traffic and web form POSTs.

How to choose a cloud-based DDoS scrubbing service:

- Choose BEFORE you get attacked. The premium you'll pay during attacks can be double.
- If you're in an industry that is attacked often, choose a flat-rate monthly plan.
- If you expect only occasional attacks, pay per engagement. Review this policy periodically.
- When negotiating with providers, insert a clause to allow annual testing of the service.

DDoS-aware network firewall

The network firewall has been the keystone of perimeter security for a long time. However, many network firewalls are not resistant to DDoS attacks at all. In fact, many of the best-selling firewalls can be disabled with the simplest layer 4 attacks. Sheer throughput is not the answer if the firewall does not recognize and mitigate the attack.

For a layer 3- and 4-based security control device, F5 recommends that architects choose a high-capacity network firewall that is a DDoS-aware. Specifically, architects should be looking to support millions (not thousands) of simultaneous connections and be able to repel SYN floods without affecting legitimate traffic.

Application Delivery Controller

Application Delivery Controllers provide strategic points of control in the network. When chosen, provisioned, and controlled properly, they can significantly strengthen a DDoS defense. For example, the full-proxy nature of the F5 Application Delivery Controller reduces computational and vulnerability-based threats by validating common protocols such as HTTP and DNS. For these reasons, F5 recommends a full-proxy Application Delivery Controller.

Web application firewall with integrated DDoS protection

The web application firewall is a higher-level component that understands and enforces the security policy of the application. This component can see and mitigate application-layer attacks whether they are volumetric HTTP floods or vulnerability-based attacks. There are several vendors that provide web application firewalls. For an effective DDoS architecture, F5 recommends only its own web application firewall module for the following reasons:

- The F5 web application firewall can provide additional services such as anti-hacking, web scraping protection, and PCI compliance.
- F5 customers benefit from using a combination of the Application Delivery Controller and web application firewall to apply application delivery and application security policy at the same time.
- The F5 Application Delivery Controller offloads and inspects SSL traffic. By combining it with the web application firewall, customers can consolidate SSL termination and security analysis of the encrypted payload in one device.

Intrusion detection and prevention systems

Intrusion detection and prevention systems (IDS/IPS) can play a small role in DDoS mitigation. F5 recommends that IDS/IPS functionality should not be deployed in a single location (for example, integrated into a layer 4 firewall). IDS/IPS should be deployed in certain instances in front of back-end components that may need specific, additional protection, such as a database or specific web server.

IP reputation database

An IP reputation database helps defend against asymmetric denial-of-service attacks by preventing DDoS attackers from using known scanners to probe an application for later exploitation and penetration. An IP reputation database can tie in at either tier of the two-tier architecture F5 recommends and may be generated internally or come from an external subscription service.

Multi-Tier DDoS Protection Architecture

For high-bandwidth customers, F5 recommends a two-tier DDoS solution. The first tier at the perimeter is composed of layer 3 and 4 network firewall services and simple load balancing to a second tier. The second tier consists of more sophisticated—and also more CPU-intensive—services including SSL termination and a web application firewall stack.

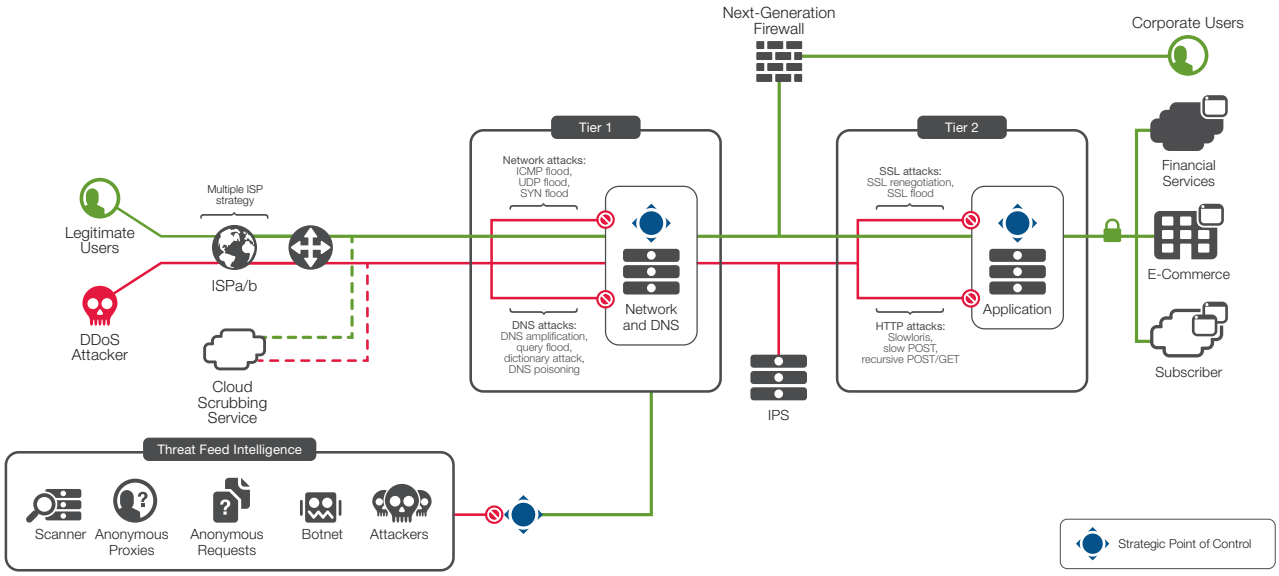

Figure 1: The two-tier F5 DDoS protection reference architecture

There are compelling benefits to the multi-tier approach:

1. The tiers can be scaled independently of one another. For example, when web application firewall usage grows, another appliance (or blade) can be added to the second tier without affecting the first tier.
2. The tiers can use different hardware platforms and even different software versions.
3. When new policies are applied at the second tier, the first tier can direct just a portion of traffic to the new policies until they are fully validated.

F5 Components and Capabilities

Table 2 shows which components are needed to provide specific capabilities. The F5 components of the DDoS protection architecture include:

- BIG-IP® Advanced Firewall Manager™ (AFM)
- BIG-IP® Local Traffic Manager™ (LTM)
- BIG-IP Global Traffic Manager™ (GTM)
- BIG-IP Application Security Manager™ (ASM)

	Tier 1	Tier 2	DNS
F5 Components	BIG-IP AFM BIG-IP LTM	BIG-IP LTM BIG-IP ASM	BIG-IP GTM with DNS Express™
OSI Model	Layers 3–4	Layer 7	DNS
Capabilities	Network firewall Tier 1 load balancing IP reputation blacklists	SSL termination Web application firewall Secondary load balancing	DNS resolution DNSSEC
Attacks Mitigated	SYN floods ICMP floods Malformed packets TCP floods Known bad actors	Slowloris Slow POST Apache Killer RUDY/Keep Dead SSL renegotiation	UDP floods DNS floods NXDOMAIN floods DNSSEC attacks

Table 2: Mapping of F5 components to DDoS mitigation capabilities

Alternative, Single-Tier Approach

While the two-tier architecture is preferred in high-bandwidth environments, F5 understands that for many customers, building multiple DDoS tiers may be overkill for their low-bandwidth environment. These customers are deploying a DDoS mitigation perimeter device that consolidates application delivery with network and web application firewall services.

The recommended practices in this document still apply to these customers. References to tier 1 and tier 2 can simply be applied to the single, consolidated tier in the alternate architecture.

Using the DDoS Protection Architecture to Maintain Availability

Tier 1 – Network Defense

The first tier is built around the network firewall. It is designed to mitigate computational attacks such as SYN floods and ICMP fragmentation floods. This tier also mitigates volumetric attacks up to the congestion of the ingress point (typically 80 to 90 percent of the rated pipe size). Many customers integrate their IP reputation databases at this tier and have controls to IP addresses by source during a DDoS attack.

Some organizations pass DNS through the first tier to a DNS server in the DMZ. In this configuration, with the right layer 4 controls they can validate the validity of DNS packets before sending them on to the server.

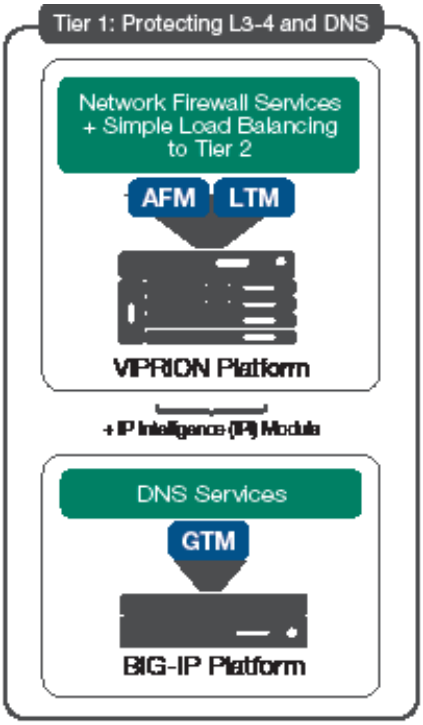

Figure 2: Tier 1 protects against network-layer DDoS attacks

Computational DDoS attack spotlight: Mitigating TCP and SSL connection floods

TCP connection floods are layer 4 attacks and can affect any stateful device on the network, especially firewalls that are not DDoS-resistant. The attack is designed to consume the memory of the flow connection tables in each stateful device. Often these connection floods are empty of actual content. Tier 1 can mitigate these by absorbing the connections into high-capacity connection tables. TCP connection floods are also mitigated by full-proxy firewalls.

SSL connection floods are designed specifically to attack the devices that terminate encrypted traffic. Due to the cryptographic context that must be maintained, each SSL connection can consume 50,000 to 100,000 bytes of memory. This makes SSL attacks especially painful. F5 recommends both capacity and the full-proxy technique for mitigating

TCP and SSL connection floods. Table 3 shows the connection capacity of F5-based network firewalls.

Platform Series	TCP Connection Table Size	SSL Connection Table Size
VIPRION Chassis	12–144 million	1–32 million
High-End Appliances	24–36 Million	2.5–7 million
Mid-Range Appliances	24 million	4 million
Low-Range Appliances	6 million	0.7–2.4 million
Virtual Edition	3 million	0.7 million

Table 3: Connection capacity of F5 hardware platforms

Tier 2 – Application Defense

The second tier is where F5 recommends deploying application-aware, CPU-intensive defense mechanisms like login walls, web application firewall policies, and dynamic security context using F5® iRules®. Often these components will share rack space with targeted IDS/IPS devices at this tier.

Tier 2 is also where SSL termination typically takes place. While some organizations terminate SSL at tier 1, it is less common due to the sensitivity of SSL keys and policies against keeping them at the security perimeter.

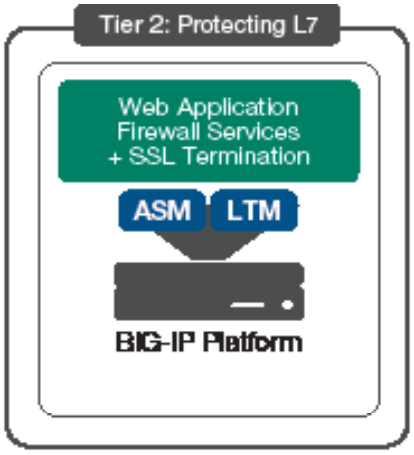

Figure 3: Tier 2 defends against application-layer DDoS attacks

Asymmetric DDoS attack spotlight: Mitigating GET floods

Recursive GETs and POSTs are among today's most pernicious attacks. They can be very hard to distinguish from legitimate traffic. GET floods can overwhelm databases and servers, and they can also cause a "reverse full pipe." F5 recorded one attacker that was sending 100 Mbps of GET queries into a target and bringing out 20 Gbps of data.

Mitigations strategies for GET floods include:

- The login-wall defense
- DDoS protection profiles
- Real browser enforcement
- CAPTCHA
- Request-throttling iRules
- Custom iRules

The configuration and setup for these strategies can be found in the F5 DDoS Recommended Practices documentation. Each of these strategies is available using a stack of web application firewall devices at tier 2.

DNS DDoS Mitigation

DNS is the second-most targeted service after HTTP. When DNS is disrupted, all external data center services (not just a single application) are affected. This single point of total failure, along with the often under-provisioned DNS infrastructure makes DNS a very tempting target for attackers.

Overprovision DNS services against query floods

DNS services have been historically under-provisioned. A significant percentage of DNS deployments are under-provisioned to the point where they are unable to withstand even small-to-medium-size DDoS attacks.

DNS caches have become popular as they can boost the perceived performance of a DNS service and they provide some resilience against standard DNS query attacks. Attackers have switched to what is called "no such domain" (or NXDOMAIN) attacks, which quickly drain the performance benefits provided by the cache.

To remedy this, F5 recommends front-ending the BIG-IP GTM DNS service with the special, high-performance DNS proxy module called DNS Express. DNS Express acts as an absolute resolver in front of the existing DNS servers. It loads the zone information from the servers then resolves every single request or returns NXDOMAIN. It is not a cache and cannot be emptied via NXDOMAIN query floods.

Consider the placement of DNS services

Often the DNS service exists as its own set of devices apart from the first security perimeter. This is done to keep DNS independent of the applications it serves. For example, if part of the security perimeter goes dark, DNS can redirect requests to a secondary data center or to the cloud. Keeping DNS separate from the security and application tiers can be an effective strategy for maintaining maximum flexibility and availability.

Some large enterprises with multiple data centers serve DNS outside the main security perimeter using a combination of BIG-IP GTM with DNS Express and the BIG-IP AFM firewall module. The main benefit of this approach is that the DNS services remain available even in the event that tier 1 firewalls go offline due to DDoS.

Regardless of whether DNS is served inside or outside the DMZ, either BIG-IP GTM or BIG-IP AFM can validate the DNS requests before they hit the DNS server.

Reference Architecture Use Cases

Following are three uses cases for the reference architecture that map to three typical customer scenarios:

1. Large financial service institution (FSI) data center
2. Enterprise data center
3. Small-to-medium business (SMB) data center

Each use case below contains a deployment scenario diagram, a short description of the specifics of the use case, and recommended sizing for the F5 components within that scenario. See table 7 for additional sizing information.

Large FSI DDoS Protection Reference Architecture

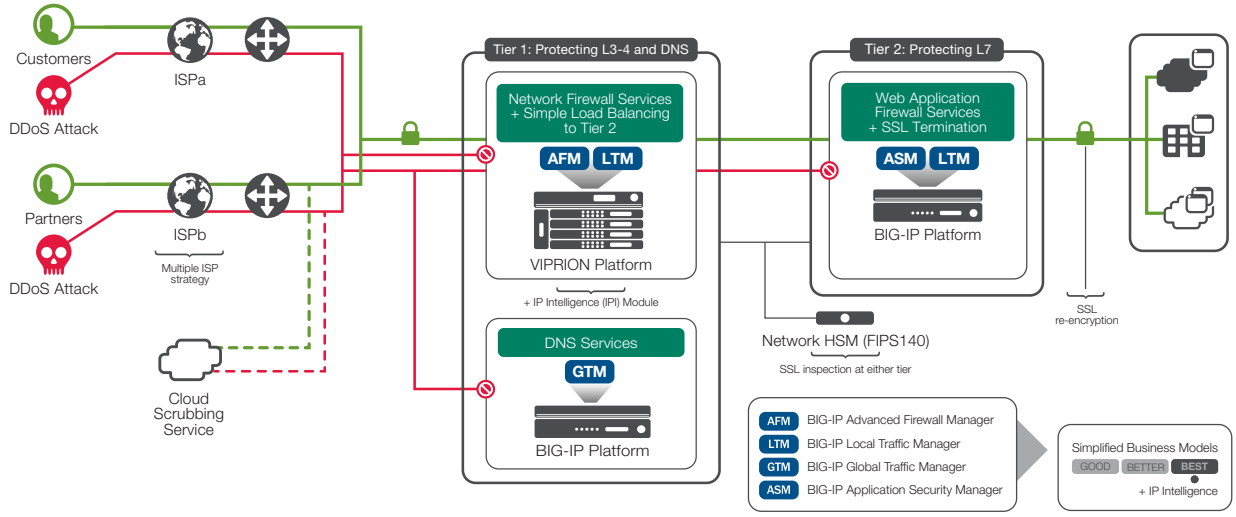

Figure 4: The F5 DDoS protection large FSI data center deployment scenario

FSI customer scenario

The large FSI data center scenario is a mature, well-recognized use case for DDoS. This is what FSIs are building right now. Typically the FSI will have multiple service providers but may forgo those service providers' volumetric DDoS offerings in favor of a cloud-based scrubbing service. The FSI data center often has few corporate staff within it so there is no need for a next-generation firewall.

FSIs have the most stringent security policy outside of the federal/military vertical. For example, nearly all FSIs must keep the payload encrypted through the entire data center. FSIs have the highest-value asset class (bank accounts) on the Internet, so they are frequent targets—not just for DDoS but also for hacking. The two-tier architecture enables FSIs to scale their CPU-intensive, comprehensive security policy at tier 2 independently of their investment in tier 1.

This use case allows FSIs to create a DDoS-resistant solution while retaining (indeed, while leveraging) the security equipment that they already have. The firewalls at tier 1 continue to do their job, and the BIG-IP ASM devices at tier 2 continue to prevent breaches.

Location	F5 Equipment
Tier 1	VIPRION Chassis (Pair)
	VIPRION Add-On: BIG-IP AFM
Tier 2	Mid-Range BIG-IP Appliance
	License Add-On: BIG-IP ASM
DNS	Mid-Range BIG-IP Appliance (Pair)

Table 4: Sizing recommendations for the FSI customer deployment scenario

Enterprise DDoS Protection Reference Architecture

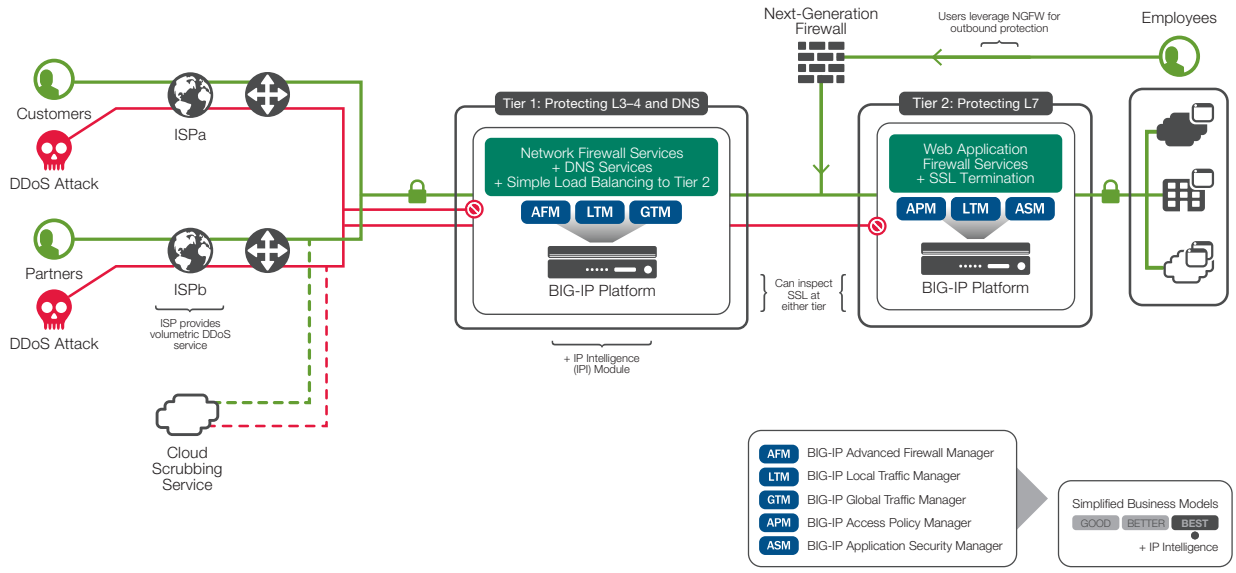

Figure 5: The F5 DDoS protection enterprise data center deployment scenario

Enterprise customer scenario

The enterprise anti-DDoS scenario is similar to the large FSI scenario. The primary difference is that enterprises do have staff inside the data center and therefore need the services of a next-generation firewall (NGFW). They are tempted to use a single NGFW for both ingress and egress, but this makes them vulnerable to DDoS attacks. Another difference is that enterprises will often use the volumetric DDoS service offered by the Internet service provider (ISP).

The recommended enterprise architecture includes a smaller NGFW on a separate path from the ingress application traffic. By using two tiers, the enterprise can take advantage of asymmetric scaling. They can add more BIG-IP ASM devices if they find that their CPU at tier 2 is at a premium.

Different verticals and companies have different requirements. By using F5 equipment at both tiers, the enterprise architecture allows the customer to decide where it makes the most sense for them to decrypt (and optionally re-encrypt) the SSL traffic. For example, an enterprise can decrypt SSL at tier 1 so that they can mirror the decrypted traffic off to a network tap that is monitoring for advanced threats.

Location	F5 Equipment
Tier 1	High-End BIG-IP Appliance (Pair)
	License Add-On: BIG-IP AFM
Tier 2	Mid-Range BIG-IP Appliance
	License Add-On: BIG-IP ASM
DNS	Mid-Range BIG-IP Appliance (Pair)

Table 5: Sizing recommendations for the enterprise customer deployment scenario

SMB DDoS Protection Reference Architecture

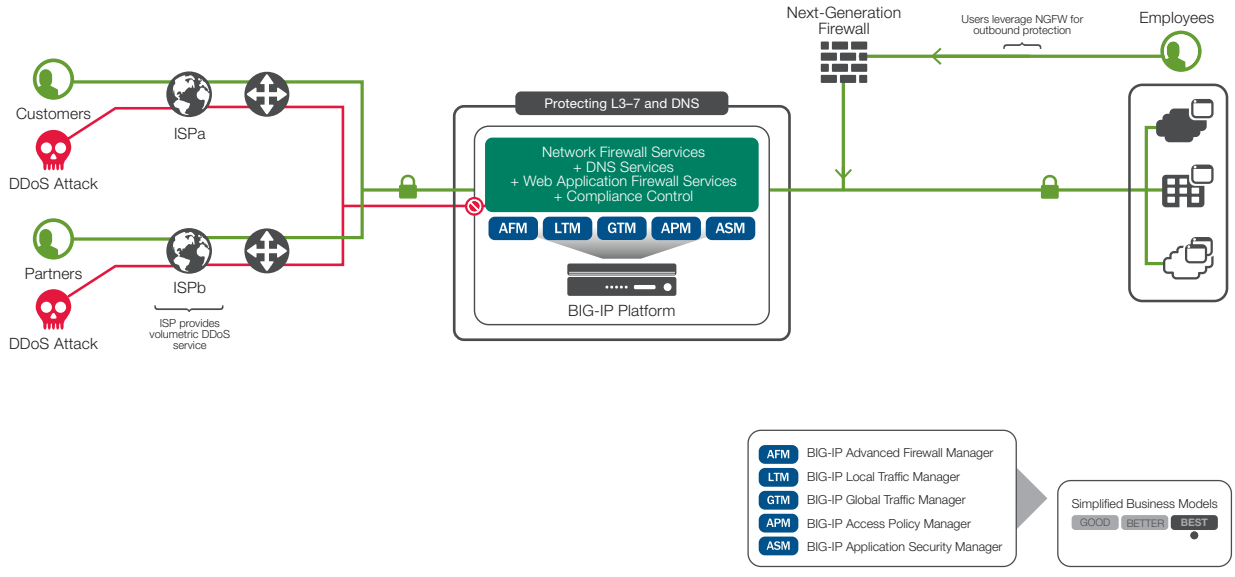

Figure 6: The F5 DDoS protection small-to-medium business data center deployment scenario

SMB customer scenario

The SMB data center use case is all about providing security while maximizing the value of consolidation. These businesses are very serious about getting the most bang for their buck. They would like to do everything from one device if they can, and they are willing to go offline during a DDoS attack.

For this use case, the customer is putting all their eggs in one basket. They will get the most cost-efficient solution but will also have the largest availability challenge. For example, if they have a policy issue with BIG-IP ASM, or if Application Visibility and Reporting cannot keep pace with the traffic, the entire box may become unavailable.

However, a one-tier architecture can actually reduce risk in smaller organizations that do not have the resources to support a larger, two-tier architecture. The organization gains efficiency by focusing specialized resources with deep knowledge on a single platform. F5 provides high availability systems, superior scale and performance, and world-class support that help further offset risk.

Certainly financial savings is the biggest benefit of the single-tier architecture. These customers get a superior DDoS solution with equipment that is already working to deliver their revenue-generating applications every day. The consolidated environment helps save on rack space, power, management, and a range of other costs.

Location	F5 Equipment
Single Tier	Mid- to High-End BIG-IP Appliance Pair
	License Add-On: BIG-IP GTM
	License Add-On: BIG-IP ASM
	License Add-On: BIG-IP AFM
	License Add-On: BIG-IP APM

Table 6: Sizing recommendations for the SMB customer deployment scenario

Sizing Specifications

Table 7 shows specifications for the range of F5 hardware devices that are available to meet customers' scaling requirements.

	Throughput	SYN Flood (per second)	ICMP Flood	HTTP Flood (JavaScript redirect)	SSL Flood (+20k attack requests)	TCP Connections	SSL Connections
VIPRION 2400 4-blade chassis	160 Gbps	196 million	100 Gbps	350,000 RPS	16,000 TPS	48 million	10 million
10200V Appliance High-end appliance	80 Gbps	80 million	56 Gbps	175,000 RPS	16,000 TPS	36 million	7 million
7200V Appliance Mid-range appliance	40 Gbps	40 million	32 Gbps	131,000 RPS	16,000 TPS	24 million	4 million
5200v Appliance Low-range appliance	30 Gbps	40 million	32 Gbps	131,000 RPS	16,000 TPS	24 million	4 million

Table 7: F5 hardware specifications for DDoS protection. See the customer use cases for specific sizing recommendations.

Conclusion

This recommended DDoS protection reference architecture leverages F5's long experience combatting DDoS attacks with its customers. Small- and medium-size businesses are finding success with a consolidated approach. Global financial services institutions are recognizing that the recommended two-tier architecture represents the ideal placement for all of their security controls. Enterprise customers are re-arranging and re-architecting their security controls around this architecture as well. For the foreseeable future, a two-tier DDoS protection architecture should continue to provide the flexibility and manageability that today's architects need to combat the modern DDoS threat.

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

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

Solutions for an application world.

WP-SEC-13307-ddos-protection