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

by F5

Introduction

For over 15 years, F5 has worked with customers to defend their applications against distributed denial of service (DDoS) attacks. Over time, many core features of the F5® TMOS® system have been made resilient against DDoS attacks. The high-profile attacks since 2012 have large financial customers and enterprises redesigning their networks to include DDoS protection. Working with these customers, F5 has developed a DDoS Protection reference architecture that includes both cloud and on-premises components.

The cloud component of the DDoS Protection reference architecture works as an insurance policy for volumetric attack mitigation. On premises, the reference architecture includes multiple tiers of defense to protect layers 3 through 7. The network defense tier protects DNS and layers 3 and 4. Freed from the noise of the network attacks, the application defense tier can use its CPU resources to protect the high-layer applications. 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 DDoS Protection architecture is built around well-known industry components. Some of these devices may be provided by other vendors and suppliers, but some are specific F5 components.

Components of a DDoS Protection Architecture

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

Attack Category	Mitigation Component
Volumetric	Cloud-Based Scrubbing Service Web Application Firewall
Asymmetric	Web Application Firewall
Computational	Application Delivery Controller Network Firewall
Vulnerability-Based	IP Reputation Database Intrusion Prevention/Detection Systems (IDS/IPS) Application Delivery Controller

Figure 1: Mapping of DDoS mitigation components to attack types.

Cloud-based DDoS scrubbing service

A 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 cloud service, hosted either from a true public cloud or within the organization's bandwidth service provider, solves the problem by sorting out the obvious bad from the likely good.

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, DDoS-aware network firewall. 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 (ADCs) 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 ADC reduces computational and vulnerability-based threats by validating common protocols such as HTTP and DNS. For these reasons, F5 recommends a full-proxy ADC.

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. Several vendors provide web application firewalls. For an effective DDoS architecture, however, 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 ADC and web application firewall to apply application delivery and application security policy at the same time.
- The F5 ADC 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 may be generated internally or come from an external subscription service.

Multi-Tier DDoS Protection Architecture

F5 recommends a hybrid cloud/on-premises DDoS solution. Volumetric attacks will be mitigated by F5 Silverline™ DDoS Protection—a service delivered via the F5 Silverline cloud-based platform. Silverline DDoS Protection will analyze and remove the bulk of the attack traffic. Sometimes, a DDoS campaign may include application layer attacks that must be addressed on premises. These asymmetric and computational attacks can be mitigated using the network defense and application defense tiers. The network defense tier is composed of layer 3 and 4 network firewall services and simple load balancing to the application defense tier. The application defense tier consists of more sophisticated (and also more CPU-intensive) services including SSL termination and a web application firewall stack.

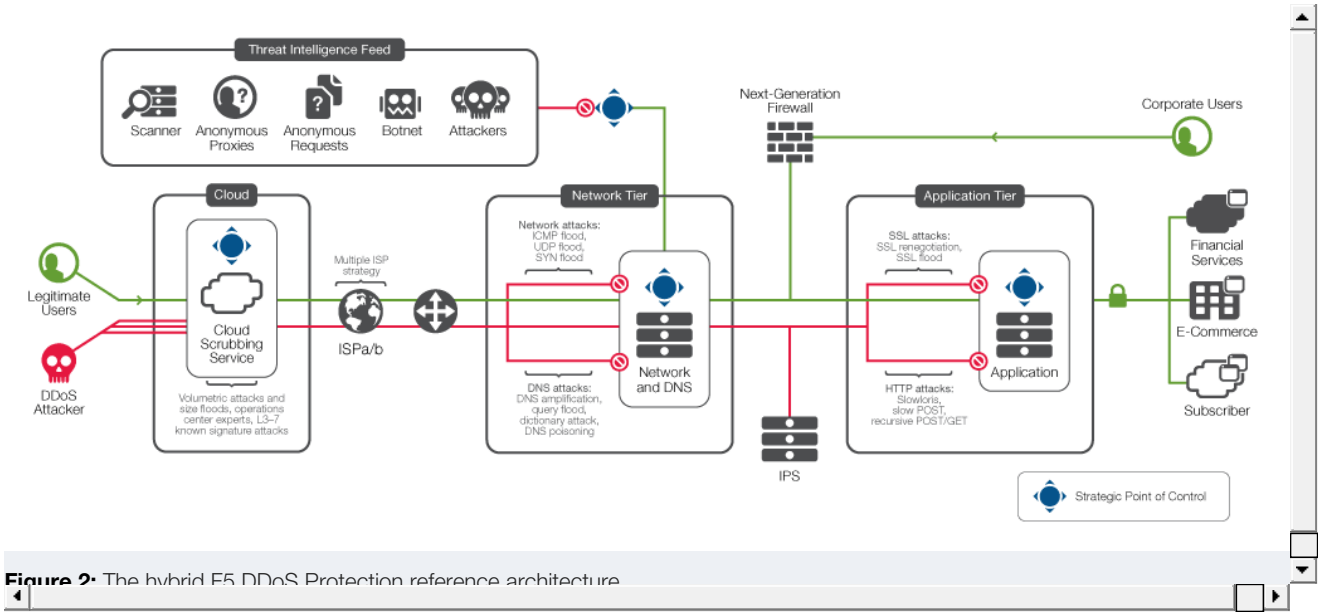

Figure 2: The hybrid F5 DDoS Protection reference architecture

There are compelling benefits to separating network defense and application defense for the on-premises portion of the DDoS Protection architecture.

1. The network and application defense 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 application tier without affecting the network tier.
2. The network and application defense tiers can use different hardware platforms and even different software versions.
3. When new policies are applied at the application defense tier, the network defense tier can direct just a portion of traffic to the new policies until they are fully validated.

F5 Components and Capabilities

Figure 3 shows components needed to provide specific capabilities. The F5 components of the DDoS Protection reference architecture include:

- Silverline DDoS Protection
- BIG-IP® Advanced Firewall Manager™ (AFM)
- BIG-IP® Local Traffic Manager™ (LTM)
- BIG-IP® Global Traffic Manager™ (GTM) with DNS Express™
- BIG-IP® Application Security Manager™ (ASM)

	Cloud	Network Defense	Application Defense	DNS
F5 Components	SilverLine DDoS Protection	BIG-IP AFM BIG-IP LTM	BIG-IP LTM BIG-IP ASM	BIG-IP GTM with DNS Express™
OSI Model	Layers 3 and 4	Layers 3 and 4	Layer 7	DNS
Capabilities	Volumetric scrubbing Traffic dashboarding	Network firewall Layer 4 load balancing IP blacklists	SSL termination Web application firewall Secondary load balancing	DNS resolution DNSSEC
Attacks Mitigated	Volumetric floods Amplification Protocol whitelisting	SYN floods ICMP floods Malformed packets TCP floods Known bad actors	Slowloris Slow POST Apache Killer RUDY/Keep Dead SSL attacks	UDP floods DNS floods NXDOMAIN floods DNSSEC attacks

Figure 3: Mapping of F5 components to DDoS mitigation capabilities

Alternative, Consolidated Approach for On-Premises Protection

While the multi-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 network and application defense tiers can simply be applied to the single, consolidated tier in the alternate architecture.

Using the DDoS Protection Architecture to Maintain Availability

Cloud for Volumetric Defense

There is always a risk of a volumetric attack sufficiently large enough to overflow an organization's ingress capacity. The defense against these attacks is to re-route the incoming attack through a set of high-bandwidth data centers that can scrub the traffic clean before returning it to the origin data center.

The factors that influence the choice of a cloud provider include capacity, latency, and value. As figure 4 shows, modern DDoS attacks are in the hundreds of gigabits per second. A modern cloud scrubber has the capacity to absorb attacks of those volumes.

Latency is added when the cloud scrubber does not have a scrubbing center sufficiently close to the customer's own data centers. Small-to-medium business (SMB) and regional companies can find cloud scrubbers within their region, but multinationals have requirements for scrubbing centers in each of the global regions.

Capacity and capability

- Global coverage—Data centers in North America, Europe, and Asia.
- Terabits of global capacity or hundreds of gigabits per center.

Organizations will say that the true value of the cloud scrubber is found only after the campaign. Questions that determine their satisfaction include:

- Was it expensive?
- What was the level of false positives?
- Did we have visibility and control into the delivery of the legitimate traffic?

Ready Defense subscription as a backup cloud-scrubbing service

Many customers already have an agreement with an external DDoS scrubbing service. These organizations can also benefit from having a backup scrubbing service. Silverline DDoS Protection can be used in this manner with its Ready Defense™ subscription. As the organization's primary DDoS scrubber, Ready Defense can take over to either assist or completely mitigate the attack.

Always Available subscription as the primary service

Organizations can use the Silverline DDoS Protection Always Available™ subscription as their primary service to respond to DDoS attacks. They can replace their existing primary service or delegate their existing service to be the secondary service.

Deployment models

Silverline DDoS Protection has two main deployment models: routed configuration and F5 IP Reflection™.

Routed configuration is for enterprises that need to protect their entire network infrastructure. Silverline DDoS Protection leverages Border Gateway Protocol (BGP) to route all the traffic to its scrubbing and protection center, and utilizes a Generic Routing Encapsulation (GRE) tunnel to send the clean traffic back to the origin network. Routed configuration is a scalable design for enterprises with large network deployments. It does not require any application-specific configuration and provides an easy option to turn on or off Silverline DDoS Protection.

IP Reflection is an alternative asymmetric technique to provide network infrastructure protection without the need for GRE tunnels. Organizations with devices that support destination NAT can leverage IP Reflection. With IP Reflection, there is no need to change any IP address and the IP address space is not affected as it is with GRE.

Return traffic methods used by Silverline DDoS Protection include:

- (AWS) Direct Connect
- IP Reflection
- GRE tunnels
- Proxy
- Customer bundles (fiber)

Volumetric attack spotlight: amplification attacks

Figure 4 shows that in 2014 the record for the world's largest DDoS attack was broken several times. Each of these attacks used a technique called "amplification," where the attackers leveraged weaknesses in NTP, DNS, and SNMP protocols to direct responses from thousands of unwitting public Internet hosts at an intended victim.

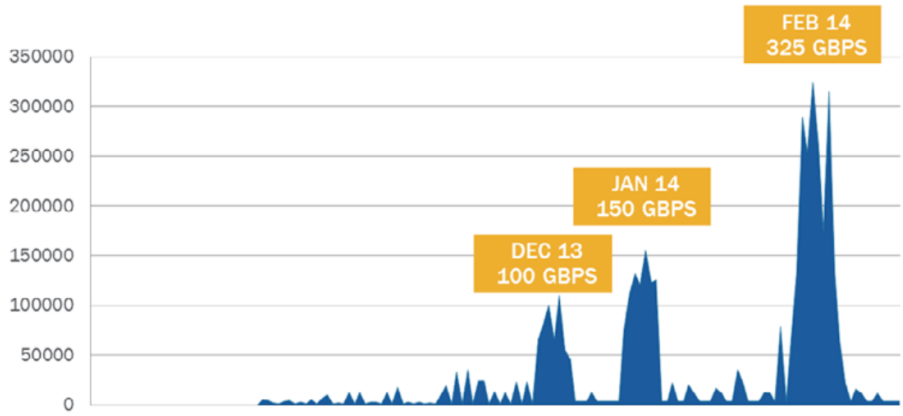

Figure 4: Increasingly large volumetric attacks in 2014.

On-Premises Network Defense

The network defense 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.

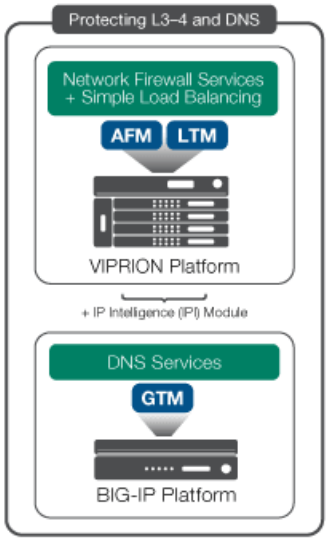

Figure 5: Network defense tier protects against network-layer DDoS attacks.

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. They can be absorbed into high-capacity connection tables in the network tier or 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. Figure 6 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

Figure 6: Connection capacity of F5 hardware platforms

On-Premises Application Defense

The application defense 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.

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

Figure 7: A web application firewall 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.

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 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 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 F5 DNS Express™. DNS Express acts as an absolute resolver in front of the existing DNS servers. It loads the zone information from the servers and 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 if the network defense tier goes 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. SMB data center

Each use case below contains a deployment scenario diagram, a short description of the specifics of the use case, and recommended F5 components within that scenario. See figure 14 for additional sizing information.

Large FSI DDoS Protection Reference Architecture

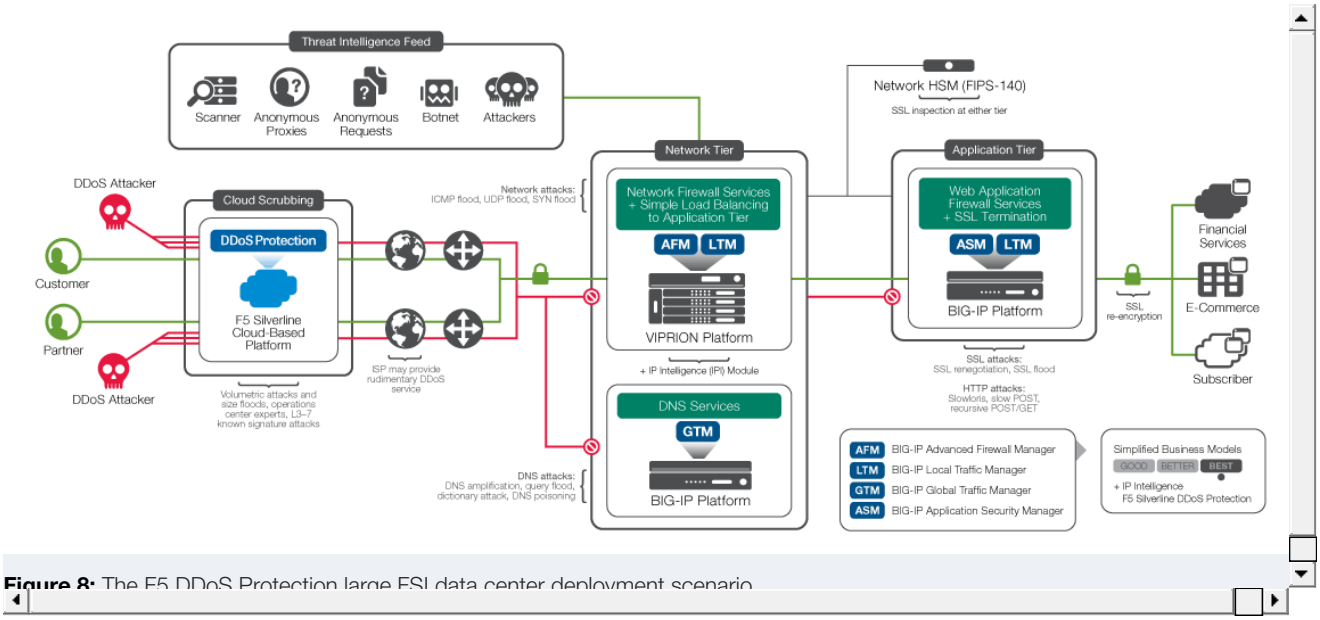

Figure 8: The F5 DDoS Protection large FSI data center deployment scenario

Large FSI customer scenario

The large FSI data center scenario is a mature, well-recognized use case for DDoS. Typically the FSI will have multiple service providers but may forgo those service providers' volumetric DDoS offerings in favor of another scrubbing service. Many of these may also have a backup volumetric DDoS service as an insurance policy against the failure of their primary cloud scrubber.

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 on-premises architecture enables FSI organizations to scale their CPU-intensive, comprehensive security policy at the application tier independently of their investment in the network tier.

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 the network defense tier continue to do their job, and the BIG-IP ASM devices at the application defense tier continue to prevent breaches.

Location	F5 Equipment
Cloud	Silverline DDos Protection: Ready Defense Subscription Always Available Subscription
Network Tier	VIPRION Chassis (Pair) VIPRION Add-On: BIG-IP AFM
Application Tier	Mid-Range BIG-IP Appliance License Add-On: BIG-IP ASM
DNS	Mid-Range BIG-IP Appliance (Pair)

Figure 9: Sizing recommendations for the FSI customer deployment scenario.

Enterprise DDoS Protection Reference Architecture

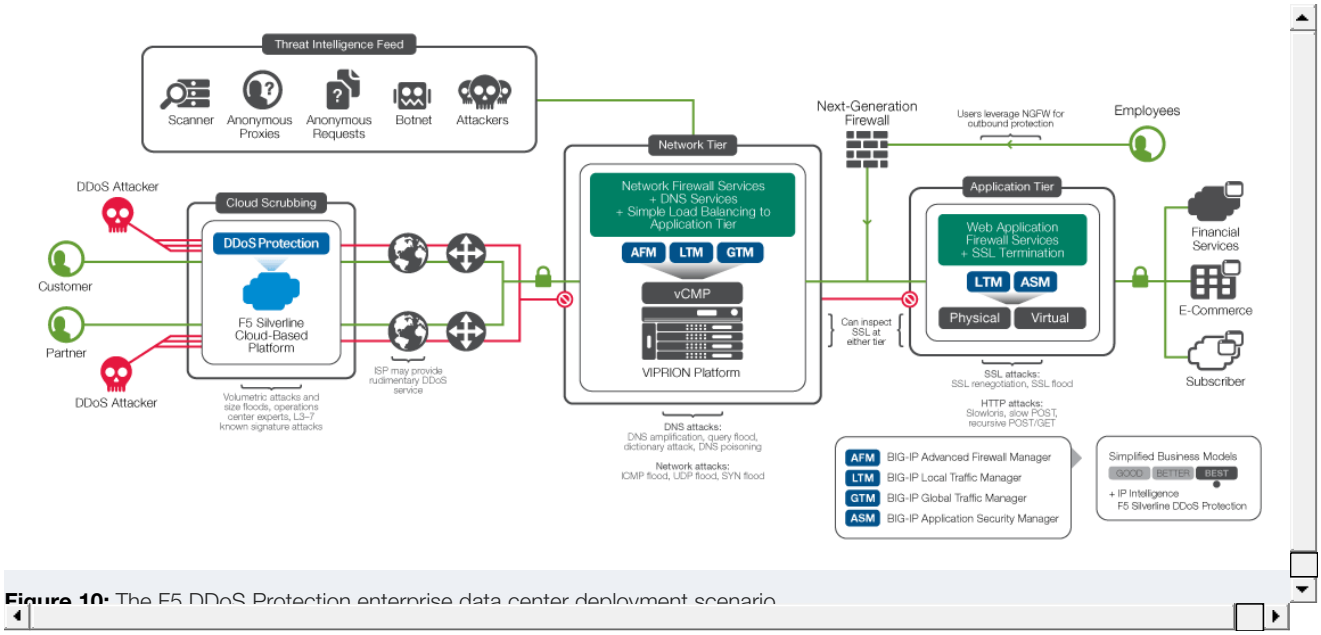

Figure 10: The F5 DDoS Protection enterprise data center deployment scenario

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).

F5 recommends that enterprises have a backup volumetric DDoS service as an insurance policy against the failure of the ISP cloud scrubber. These customers can use the Ready Defense subscription as that secondary service for volumetric protection.

On premises, the recommended enterprise architecture includes a smaller NGFW on a separate path from the ingress application traffic. By using a network defense tier and an application defense tier, enterprises can take advantage of asymmetric scaling—adding more BIG-IP ASM devices if they find that CPU is at a premium.

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

Location	F5 Equipment
Cloud	Silverline DDoS Protection: Ready Defense Subscription Always Available Subscription
Network Tier	High-End BIG-IP Appliance (Pair) License Add-On: BIG-IP AFM
Application Tier	Mid-Range BIG-IP Appliance License Add-On: BIG-IP ASM
DNS	Mid-Range BIG-IP Appliance (Pair)

Figure 11: Sizing recommendations for the enterprise customer deployment scenario.

SMB DDoS Protection Reference Architecture

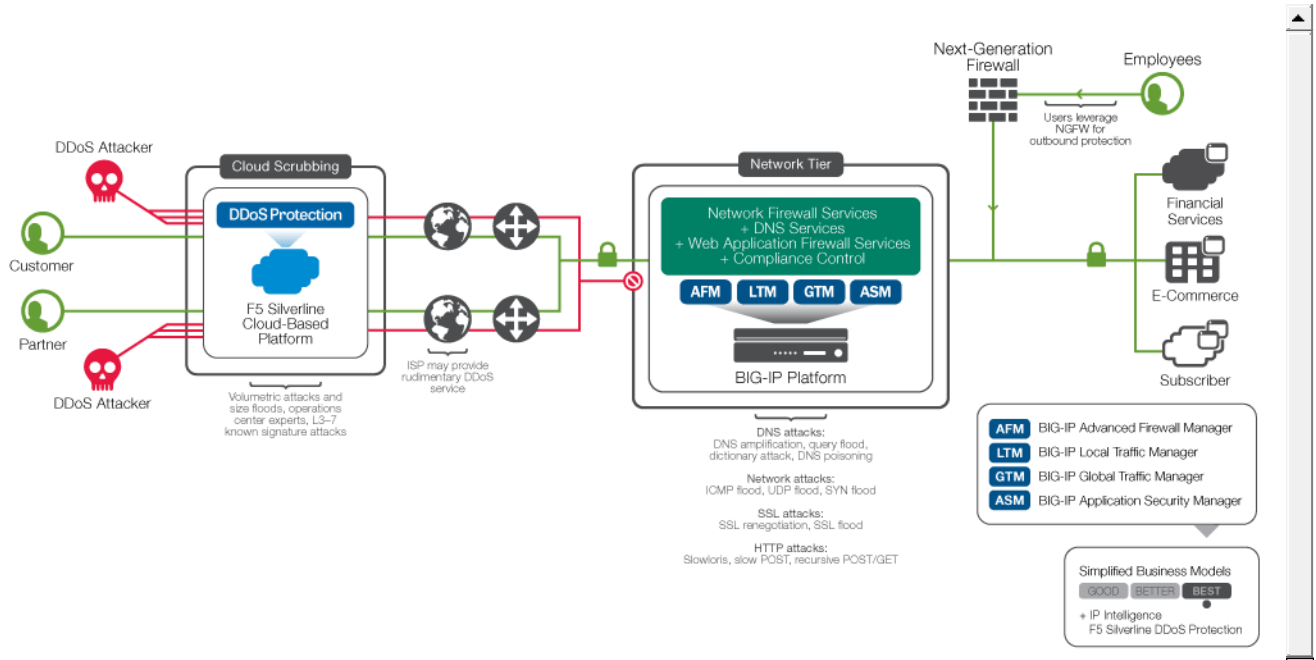

Figure 12: The F5 DDoS Protection small-to-medium business data center deployment scenario

SMB customer scenario

The SMB data center use case is all about providing security while maximizing the value of consolidation. These businesses are 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 of its eggs in one basket. It will get the most cost-efficient solution but will also have the largest availability challenge.

On the other hand, 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 this consolidated 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, and management.

Location	F5 Equipment
Cloud	Silverline DDoS Protection: Ready Defense Subscription Always Available Subscription
Consolidated On-Premises 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

Figure 13: Sizing recommendations for the SMB customer deployment scenario.

Sizing Specifications

Figure 14 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)	TCP Connections	SSL Connections
VIPRION 2400 4-blade chassis	160 Gbps	196 million	100 Gbps	350,000 RPS	48 million	10 million
10200V Appliance High-end appliance	80 Gbps	80 million	56 Gbps	175,000 RPS	36 million	7 million
7200V Appliance Mid-range appliance	40 Gbps	40 million	32 Gbps	131,000 RPS	24 million	4 million
5200v Appliance Low-range appliance	30 Gbps	40 million	32 Gbps	131,000 RPS	24 million	4 million

Figure 14: 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 combating 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 hybrid architecture represents the ideal placement for all of their security controls. Enterprise customers are rearranging and rearchitecting their security controls around this architecture as well. For the foreseeable future, a hybrid 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

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

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

Japan
f5j-info@f5.com