Digitally signed by Vicol Irina
Date: 2022.06.17 16:42:23 EEST
Reason: MoldSign Signature
Location: Moldova

Maximize the Value of Your Enterprise Application Delivery

To make it easier and more affordable to get the capabilities your organization needs, F5 provides three offerings: Good, Better, Best.

With the F5 Good, Better, Best options, you receive:

- Flexibility—Choosing from available modules makes it easier to adopt advanced F5 functionality.
- Simplicity—Having a standardized image and fewer configurations can help simplify operations.
- Best Value—You can save up to 65% vs. buying as components.

Select your right fit

Get advanced traffic management, optimization, and security services on one F5® BIG-IP® platform with the right offering for your organization.

GOOD

Provides intelligent local traffic management for increased operational efficiency and peak network performance of applications.

BETTER

All the benefits of “Good” plus advanced application delivery optimization.

BEST

“Better” plus advanced access management and application security. Delivers optimal security, performance, and availability for your applications and network.

Choose the right platform for your organization

VE **Virtual**

Get flexible deployment options for virtual environments and the cloud.

Hardware

Achieve high performance with specialized and dedicated hardware.

VE **Hybrid**

Combine virtual and physical editions for flexibility and performance.

Features and Capabilities	Good	Better	Best
BIG-IP® Local Traffic Manager™			
Load balancing and monitoring	●	●	●
Application visibility and monitoring	●	●	●
L7 intelligent traffic management	●	●	●
Core protocol optimization (HTTP, TCP, HTTP/2, SSL)	●	●	●
SSL proxy and services	●	●	●
IPv6 support	●	●	●
Programmability (iRules®, iCall™, iControl®, iApps®)	●	●	●
ScaleN™ (on-demand scaling of performance and capacity)	●	●	●
BIG-IP® APM® Lite (user authentication, SSL VPN for 10 concurrent users)	●	●	●
SYN flood DDoS protection	●	●	●
Software Services			
Advanced routing (BGP, RIP, OSPF, ISIS, BFD)	Optional	●	●
BIG-IP® DNS			
Global server load balancing		●	●
DNS services		●	●
Real-time DNSSEC solution		●	●
Global application high availability		●	●
Geolocation		●	●
DNS DDoS attack prevention		●	●
BIG-IP® Advanced Firewall Manager™			
High-performance ICSA firewall		●	●
Network DDoS protection		●	●
Application-centric firewall policies		●	●
Protocol anomaly detection		●	●
BIG-IP® Application Security Manager™			
PCI-compliant web application firewall			●
Web scraping prevention			●
Integrated XML firewall			●
Violation correlation and incident grouping			●
Application DDoS protection			●
BIG-IP® Access Policy Manager®			
500 concurrent user sessions; scalable up to 200,000			●
BYOD enablement			●
Full proxy for VDI (Citrix, VMware)			●
Single sign-on enhancements (identity federation with SAML 2.0)			●

SDN services are automatically included with all iSeries appliances. BIG-IP Application Acceleration Manager (AAM) is end-of-sale.

To learn more about F5 market-leading technologies, visit f5.com.

DATA SHEET

BIG-IP Virtual Edition

Software-based application delivery services are critical to maintaining the adaptable and secure application infrastructure demanded by enterprises undergoing digital transformation. F5 accelerates your transition to the cloud and software-defined architectures with virtual application delivery platforms that provide an agile, flexible, and efficient way to deploy advanced application and security services.

Many enterprises have or are planning to deploy applications across multiple cloud environments—both public and private—making it more difficult to implement advanced, consistent and compliant application services for every app in their portfolio. Furthermore, they are expanding beyond traditional monolithic applications and deploying more modern, dynamic application architectures including containers and microservices that have unique requirements.

Standardizing on F5 app services accelerates migration to, and between clouds, while providing consistent and advanced services for both monolithic and modern applications running in those environments—helping you more easily support and manage your growing multi-cloud application portfolio.

F5® BIG-IP® Virtual Editions (VEs) are the industry's most scalable virtual application delivery controllers (vADCs)—facilitating high-performance application traffic processing across all leading hypervisors and cloud platforms—easing your transition from hardware to software. VE's deliver all the same market-leading application delivery services—including advanced traffic management, application security, application acceleration, DNS, network firewalling and secure access management—that run on F5 purpose-built hardware. This similarity enables service configurations and policies from existing F5 appliances to be reused and replicated on VE's, simplifying cloud migrations. VE's can easily be provisioned and configured automatically by network operators and developers alike, allowing them to be integrated within existing CI/CD pipelines and ensuring all applications are deployed with the necessary security, compliance and traffic management capabilities. When used in conjunction with F5 BIG-IQ® Centralized Management, you can rapidly create, provision, and manage application services anywhere while gaining visibility into the health and performance of your multi-cloud apps, all from a centralized point of control.

AVAILABLE BIG-IP MODULES:

- BIG-IP Local Traffic Manager (LTM)
- BIG-IP DNS
- BIG-IP Advanced Firewall Manager (AFM)
- BIG-IP Access Policy Manager (APM)
- Advanced WAF
- SSL Orchestrator
- BIG-IP Carrier Grade NAT (CGNAT)
- BIG-IP Policy Enforcement Manager (PEM)

Key Benefits

Increase multi-cloud agility
Quickly and easily spin up, spin down, or migrate application delivery services across the data center and public cloud, using instant deployment options as needed.

Accelerate deployments with automation
Automate app services insertion with F5's Automation Toolchain. It enables declarative provisioning and configuration of BIG-IP VE across cloud environments and integration with automation and CI/CD tools including Ansible, Jenkins, and Terraform.

Optimize application and security services
Implement robust security and traffic management services to keep your apps available, protected and compliant—regardless of deployment location.

Use modern application architectures
Native integration with container orchestration environments lets you implement advanced app services that are as dynamic as your containers.

Support high-performance requirements in the cloud
Make the transition from hardware to software without the typical performance degradation issues.

Gain ultimate deployment and consumption flexibility
Deploy BIG-IP VE across the broadest array of supported hypervisor and cloud platforms with the freedom to consume through perpetual, utility, subscription, or enterprise licensing agreement (ELA).

Primary Cloud Scenarios

BIG-IP VEs can be used to deliver a consistent set of advanced application services in the four primary cloud scenarios described below: private cloud/software-defined data center (SDDC), public cloud, multi/hybrid cloud, and colocation with cloud interconnect.

PRIVATE CLOUD USING SOFTWARE-DEFINED ARCHITECTURES

Enterprises are migrating to private cloud/SDDCs to achieve agility, application time to market, and to provide control to application owners and developers via a self-service portal or catalog. A private cloud or SDDC using F5 application services is ideal for speeding application deployments, enabling dynamic changes in the data center, and matching infrastructure services to workloads using a per-app model. F5 products and solutions integrate with the leading private cloud technology platforms, including OpenStack, VMware, Cisco, and Microsoft Azure Stack. F5 provides cloud solution templates and supports open source tools like Heat, Ansible, and open-vm-tools to orchestrate and automate the deployment of app delivery and security services.

Flexibility and high performance in a two-tier hybrid architecture

Some enterprises are moving to a two-tier architecture as part of their SDDC transformation. At the edge of the network is the application tier that provides front-door services including L4 traffic management, DDoS firewall, or SSL offload—for all traffic entering the network, based on overall business and security policies. Services that deal with high-volume traffic require the highest performance and scalability, a case where dedicated, purpose-built hardware can be more cost-efficient than commodity servers. The per-app tier manages the application stack inside the data center, which leverages highly scalable, flexible software to deliver advanced application and security services on a per-application basis. This two-tier hybrid data center model (see Figure 1) offers the best of both worlds: hardware where it's needed and software agility close to the app.

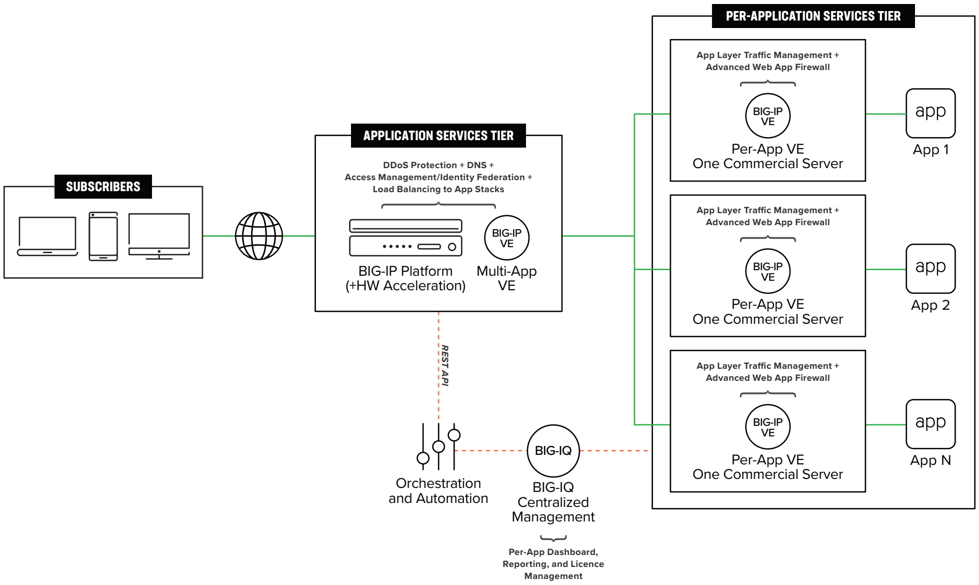
SUBSCRIBERS
APPLICATION SERVICES TIER
DDoS Protection + DNS +
Access Management/Identity Federation +
Load Balancing to App Stacks
BIG-IP Platform
(+HW Acceleration)
BIG-IP VE
Multi-App VE
REST API
Orchestration and Automation
BIG-IQ
BIG-IQ Centralized Management
Per-App Dashboard, Reporting, and Licence Management
PER-APPLICATION SERVICES TIER
App Layer Traffic Management + Advanced Web App Firewall
BIG-IP VE
Per-App VE
One Commercial Server
App Layer Traffic Management + Advanced Web App Firewall
BIG-IP VE
Per-App VE
One Commercial Server
App Layer Traffic Management + Advanced Web App Firewall
BIG-IP VE
Per-App VE
One Commercial Server
app
App 1
app
App 2
app
App N

Figure 1: Two-tier architecture with F5 hardware or shared multi-app VE at the edge and per-app VEs.

DEPLOY APPLICATIONS IN AND ACROSS PUBLIC CLOUD ENVIRONMENTS

Deploying applications in the leading public clouds gives you the flexibility and scalability you want, without the investment and capital costs associated with building out additional private data centers. Using F5 application and security services delivered by BIG-IP VEs provides the following benefits:

- **Repeatable architectures across cloud environments**—as you expand and adopt new clouds, reuse the same secure, validated, and compliant architecture to accelerate multi-cloud adoption and simplify operations.
- **Reduced tool sprawl and operational complexity**—standardizing on familiar services that are cloud-agnostic makes deploying and maintaining apps across cloud environments quicker and easier.

- **Consistent levels of availability, performance, and security**—provide your customers with an excellent user experience while protecting both your revenue and reputation.
- **Faster time to market**—rapidly provision advanced application services when launching new applications or migrating existing applications to the public cloud.
- **Deep integration with public cloud providers**—dynamically scale out app services through integration with AWS Auto Scaling, or easily apply advanced application security with an out-of-the-box, pre-configured web application firewall (WAF) solution in the Azure Security Center.
- **Flexible licensing models**— consume with a licensing model supportive of your business requirements, whether that's as a subscription, enterprise licensing agreement (ELA), pay-as-you-go, or on a perpetual-basis.

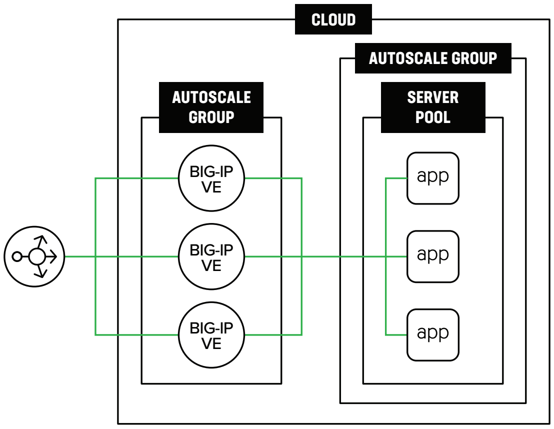

Figure 2: BIG-IP VE's deployed within an autoscaling architecture–either within or across availability zones–to ensure that your apps are available and secure while optimizing costs as your apps scale to match demand.

APPLICATION PORTABILITY ACROSS HYBRID AND MULTI-CLOUD ENVIRONMENTS

Despite the many benefits of public cloud deployments, enterprises often avoid moving all applications or data to the public cloud due to perceived loss of control, risk, regulatory compliance, and lack of support for legacy application design. As a result, many elect to operate within a hybrid cloud or hybrid multi-cloud model whereby part of their operations run in the public cloud(s) while components unable to move to the cloud or that require advanced security and compliance monitoring remain on-premises. In some scenarios, applications operate across environments to increase redundancy or to allow greater scale-out capacity when needed. F5 increases the portability of these apps while reducing management overhead by providing a set of standardized application services that can be reused wherever an app is currently running, or wherever it's redeployed to. In Figure 3, internet-facing front-end applications are deployed in the public cloud while mission-critical workloads with greater security and compliance requirements run on-premises. A direct connection links the two environments to reduce latency.

COLOCATION DEPLOYMENTS WITH DIRECT CONNECT TO PUBLIC CLOUD

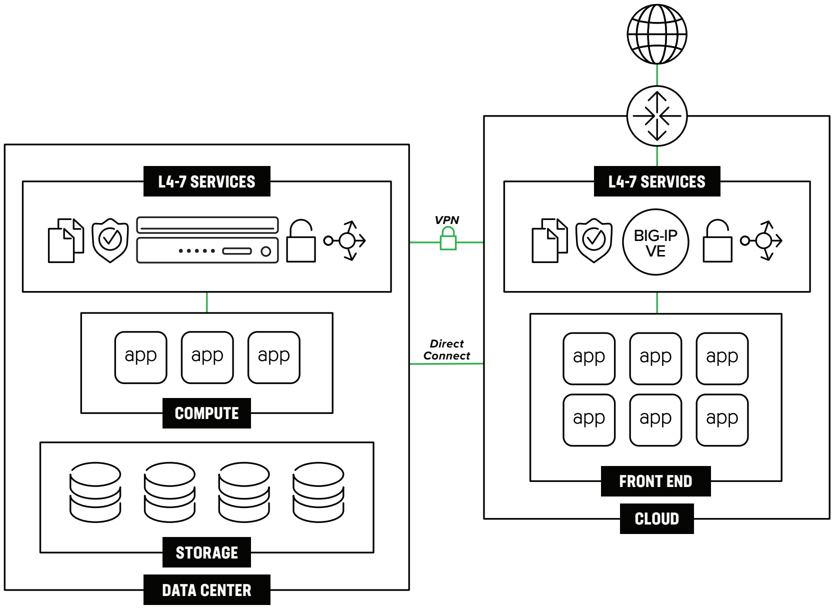

Figure 3: Hybrid cloud deployment with BIG-IP Virtual Editions supporting apps across public cloud and data center.

Many enterprises operate their application portfolio in a hybrid cloud model similar to that shown in Figure 3. But, for some, there may be an associated latency increase caused by large distances between their data center and cloud edge locations. For these organizations, the best option is to deploy on-premises apps within a colocation facility and use direct connections to connect both ends of their hybrid architecture. F5 BIG-IP VE can also be deployed in these colocation facilities and used to provide application service insertion for both apps deployed in the colocation and those running in the public cloud. As a result, consistent app services can be implemented for apps running in different cloud environments.

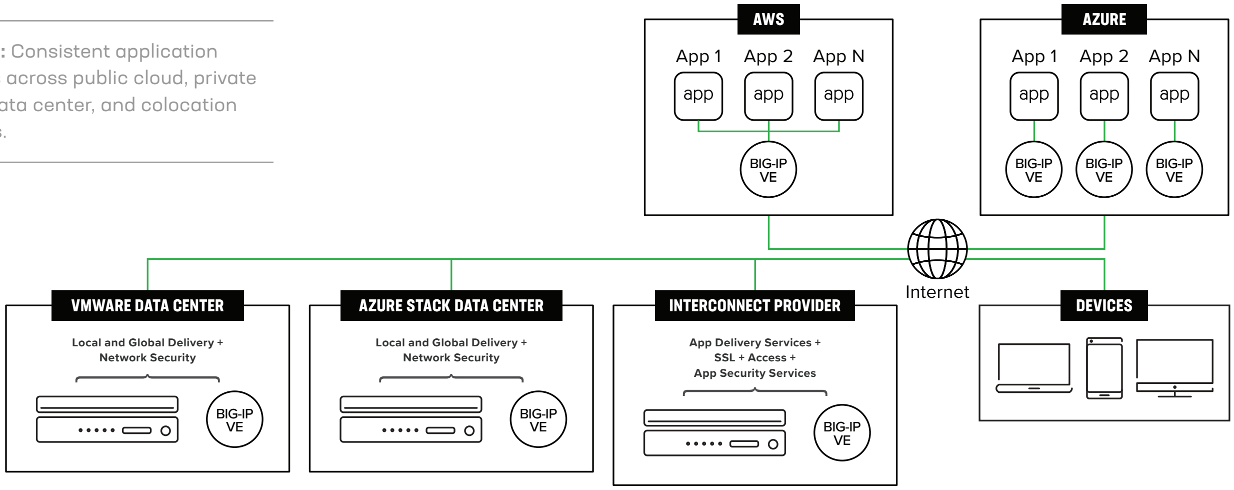

Figure 4: Consistent application services across public cloud, private cloud, data center, and colocation facilities.

Integration with SDN Frameworks

Software-defined networking (SDN) achieves agility, flexibility, and cost-efficiency in terms of overcoming the complexity of networking infrastructure in data centers today. SDN seeks to operationalize the network through virtualization and abstraction, similar to what has occurred for servers and storage. However, while SDN has focused on stateless L2–3 connectivity, there remains the need for stateful and flow-aware L4–7 services. Through its Technology Alliance partnerships, F5 is completing the SDN vision by integrating its intelligent app delivery services with leading SDN architectures (VMware NSX, Cisco ACI) via BIG-IP plug-ins and REST APIs. In addition, BIG-IP platforms can serve as SDN gateways, bridging virtualized networks and traditional network architectures to provide a smooth transition and investment protection.

Achieve Hardware-Comparable Performance with Software

A significant inhibitor of cloud adoption among large enterprises and especially service providers, is the reduction in performance typically associated with a transition from hardware to software. This means that, for many, the promise of increased deployment agility and scalability the cloud offers may not be worth sacrificing low latency, highly responsive user experiences their data center delivers.

The BIG-IP Virtual Edition is the most scalable, high performing virtual ADC available, capable of supporting 100Gbps NICs within a single instance, meaning you don’t have to choose between agility and high performance—you can have both. Below are a few examples of how BIG-IP Virtual Edition has been augmented to provide even greater performance.

- **High Performance VEs**—these VE instances aren’t limited by a throughput cap, but are instead licensed by the number of vCPU cores that can be allocated. That lets you optimize the underlying host hardware and achieve 85Gbps+ of L4 throughput.
- **SR-IOV and Advanced Network Interface Card (NIC) support**—BIG-IP VE’s driver is optimized to interact directly with underlying NIC’s using Single Root I/O Virtualization (SR-IOV), significantly improving throughput performance and reducing latency. SR-IOV can be enabled in AWS using AWS ENA, in Azure with Azure Accelerated Networking, and in private cloud environments with select Intel, Mellanox, Broadcom, and Emulex NIC’s.
- **Accelerated cryptographic and compression processing**—the BIG-IP VE can offload compute-intensive cryptographic functions and compression using Intel’s Quick Assist Technology, freeing up CPU cycles to focus on other important application tasks.

- **Offload to FPGA-enabled SmartNIC**—Offload various compute-intensive tasks to a high-performance Intel SmartNIC, including DDoS mitigation, Carrier Grade NAT (CGNAT) and layer 4 traffic transmission. Doing so significantly improves performance by over 30%, while reducing strain on BIG-IP VE compute resources by up to 80%.

Dynamic App Services for Container Environments

Organizations are rapidly adopting containerized environments to develop more agile and portable applications, typically using management and orchestration frameworks to coordinate the provisioning and automation of these workloads. But these apps still need services like SSL offload, routing, and web application protection

F5 Container Ingress Services (CIS) is a container integration solution that helps developers and system's teams manage front-door ingress control and advanced application delivery and security services for container and Platform as a Service (PaaS) deployments. CIS integrate BIG-IP VE with native container environments and orchestration systems, including Kubernetes and RedHat OpenShift. That integration enablers dynamic Ingress HTTP routing, load balancing, and security for containers as they're spun up.

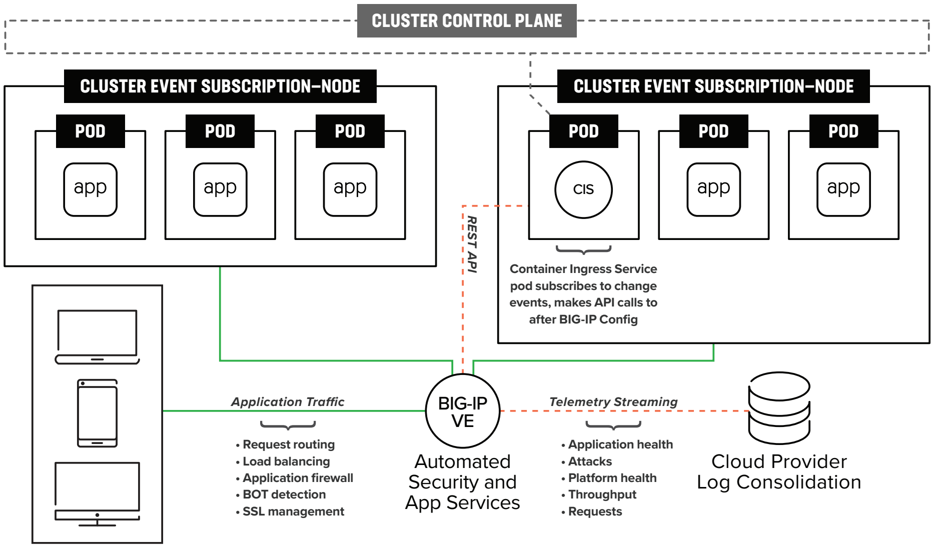

Figure 5: BIG-IP VE providing front-door app services to containers using F5 Container Ingress Services.

Automation, Orchestration, and Programmability

F5 offers many ways to program the application services fabric and network, enabling organizations to react in real time to operational and business events, automate deployment and configuration, and easily integrate into home-grown or third-party orchestration systems.

- **F5 Automation Toolchain** Provides a set of open-source automation tools that make it faster and easier to deploy and configure BIG-IP VE via simple, yet powerful declarative interfaces—all of which can be consumed as part of a complete CI/CD pipeline. It includes:
    - **Declarative onboarding** for L1–3 provisioning
    - **Application services extension 3 (AS3)** for L4–7 configuration
    - **Telemetry streaming** for aggregating, normalizing, and forwarding app stats and events to third-party analytics tools
- **F5 Cloud Solution Templates** Enable automatic deployment and bootstrapping of BIG-IP VEs across all leading public and private cloud environments and across a diverse range of architectural topologies, including HA and autoscaling.
- **F5 Cloud Failover Extension (CFE)** An iControl LX extension that provides L3 failover functionality in cloud environments, effectively replacing Gratuitous ARP (GARP).
- **F5 iRules** Scripting that provides granular traffic control and visibility, enabling customization, rapid response to errors in application code and security vulnerabilities, and support for new protocols.

Visit F5's GitHub repository for additional information on the F5 Automation Toolchain, Cloud Solution Templates, and other open-source extensions and integrations.

Centralized Management of BIG-IP VE

F5® BIG-IQ® Centralized Management provides a unified point of control for your entire F5 portfolio, ensuring your finger remains on the pulse of devices, modules, and licenses—helping you deliver optimal application availability, performance, and security. It provides a single pane of glass to manage and deploy F5 devices, including key BIG-IP modules like BIG-IP Local Traffic Manager (LTM), BIG-IP Application Security Manager (ASM), BIG-IP Advanced Firewall Manager (AFM), BIG-IP Access Policy Manager (APM), and BIG-IP DNS as well as other F5 solutions including SSL Orchestrator, Secure Web Gateway, DDoS Hybrid Defender, WebSafe, and MobileSafe.

Use BIG-IQ Centralized Management to:

- Automatically back-up images and configurations.
- Monitor dashboards, reporting, and alerting.
- Provide role-based access control (RBAC).
- Obtain detailed analytics on a per-app basis.
- Manage BIG-IP VE licenses.
- Ensure consistent security and traffic management policies across your infrastructure.
- Create, provision, and deploy new BIG-IP VE devices and app services.
- Align to modern development practices and CI/CD workflows through Automation Toolchain.
- Assign and manage machine identities and certificates via Venafi integrations.

BIG-IQ's VE license management lets you automate large-scale virtual ADC deployments, including per-app VEs, in supported clouds with an F5 subscription or ELA licensing. With BIG-IQ Centralized Management, you can spin up and provision individual VE licenses from a single license pool on demand. When resource requirements decrease, you can spin down the VE and return it to the license pool for future use.

Specifications

Available in a range of performance options, F5 virtual editions can be sized and configured to suit the application services required. Maximum performance is based on applicable VE licensed performance ranges and resources (number of CPU cores/memory) allocated.

Minimum resource requirements: 1vCPU, 2 GB RAM, and 10 GB disk.

THROUGHPUT LICENSED VE

Figure 6: BIG-IP performance with Dell PowerEdge R620 with Intel Xeon CPU E5-2670 0 @ 2.6GHz and Intel 82599EB 10-Gigabit SFP+ NIC–configured for PCI pass-through with support for SR-IOV.

Performance	Starting	Maximum*
L7 requests per second	3,000	450,000
L4 connections per second	2,000	135,000
L4 throughput	25 Mbps	10 Gbps**
Maximum L4 concurrent connections	1 million	10 million

SSL	Starting	Maximum*
SSL RSA TPS (2K keys)	900	3,800
SSL throughput (RSA)	23 Mbps	4 Gbps
SSL ECC TPS	1,200	20,000***
SSL throughput (ECC)	23 Mbps	5.4 Gbps
Software Compression	**Starting**	**Maximum***
Compression throughput	20 Mbps	4 Gbps
DNS	**Starting**	**Maximum***
Query response per second	1,000	250,000

Note: BIG-IP APM specifications are maintained within this support.f5.com article.

* Maximum performance specs are based on ideal lab testing conditions with maximum supported vCPUs and may vary due to customer or cloud provider environmental conditions, type of hypervisor used, and capacity of host server hardware. Please refer to SOL14810 on askf5.com for specific license and performance details that may impact your performance.

** 10 Gbps throughput requires use of NICs that support SR-IOV.

*** Based on ECDHE_ECDSA_AES256_GCM_SHA384 cipher string, running BIG-IP TMOS v12.1.

HIGH PERFORMANCE VE

Figure 7: BIG-IP LTM VE performance on SuperMicro 2U server with dual Intel® Xeon® Scalable Processors @ 28cores (2.7GHz) and Intel XL710 40G NIC—configured for SR-IOV using VMware ESXi 6.5 hypervisor. High performance VE licensed for 24 vCPUs, running BIG-IP TMOS v15.x and later required.

Performance	Maximum*
L7 requests per second	4.6 million
L4 connections per second	1.4 million
L4 throughput	85 Gbps**
SSL	**Maximum***
SSL RSA TPS (2K keys)	30,000
SSL throughput (RSA)	32 Gbps
SSL ECC TPS	100K
SSL Throughput (ECC)	37 Gbps

Figure 8: BIG-IP LTM VE performance on Neon City Platform with 2x Intel Xeon® Gold E5-6230N Processor, Intel® QuickAssist Adapter 8970 with 3x QAT Physical Functions (End-Point) and Intel XL710 40G NIC – configured for SR-IOV using KVM CentOS 7.5. High Performance VE licenses for 16vCPUs (for ECC) and 20vCPUs (for RSA), running BIG-IP TMOS v14.1.0.3 and later.

SSL with Intel QAT	Maximum*
SSL RSA TPS (2K keys)	95K
SSL throughput (RSA)	60 Gbps
SSL ECC TPS	59K
SSL Throughput (ECC)	46 Gbps
BIG-IP DNS	**Maximum***
Query responses per second	1.8 million

BIG-IP VE FOR SMARTNICS

Figure 9: High Performance VE (8vCPU/16GB RAM) with enabled & Intel FPGA PAC N3000 SmartNIC.

	DDoS Attack Size Mitigation	VE CPU Utilization
DDoS Protection without SmartNIC	2.5 Gbps	100%
DDoS Protection with SmartNIC	40 Gbps	27%
	L4 Throughput	**VE CPU Utilization**
CGNAT (NAT44 NAPT) without SmartNIC	37 Gpbs	87%
CGNAT (NAT44 NAPT) with SmartNIC	48 Gbps	4%
	L4 Throughput	**VE CPU Utilization**
L4 Acceleration without SmartNIC	36 Gbps	81%
L4 Acceleration with SmartNIC	48 Gbps	4%

Note: BIG-IP APM specifications are maintained within this support.f5.com article.

SUPPORTED HYPERVISORS AND LINUX DISTRIBUTIONS

F5 offers the most flexible deployment options in the industry, with support across all major virtualization platforms.

Figure 10: F5 BIG-IP VE support for the leading hypervisors. (For the full list of supported versions, please go to the VE Supported Hypervisors Matrix on ask.f5.com.)

	Lab	25 Mbps	200 Mbps	1 Gbps	3 Gbps	5 Gbps	10 Gbps
VMware vSphere	●	●	●	●	●	●	●
KVM and Community Xen	●	●	●	●	●	●	●
Microsoft Hyper-V	●	●	●	●	●		

* Maximum performance specs are based on ideal lab testing conditions, optimized host and guest settings, maximum supported vCPUs, SR-IOV capable NICs, and may vary due to customer or cloud provider environmental conditions, type of hypervisor used, and capacity of host server hardware and NICs. Please refer to SOL14810 on askf5.com for specific license and performance details that may impact your performance.

** 85 Gbps throughput achieved using Mellanox CX-5 100G NIC configured for SR-IOV using KVM CentOS 7.5.

Figure 11: High-performance, VE-supported hypervisors. Note: The high performance paravirtualized driver is used as the default driver for throughput licensed and high performance VE.

	High-Perf. SR-IOV	High-Perf. Paravirtualized Driver
KVM	●	● (Virtio)
VMware vSphere	●	● (Vmxnet3)

SUPPORTED PUBLIC CLOUD IAAS PROVIDERS

F5 offers support for leading public cloud providers including Amazon Web Services, Microsoft Azure, Google Cloud Platform, and IBM Cloud.

Figure 12: F5 BIG-IP VE support for the leading public cloud IaaS providers. For details and a list of validated cloud providers, visit F5.com.

	Lab	25 Mbps	200 Mbps	1 Gbps	3 Gbps	5 Gbps	10 Gbps*	HPVE*
Amazon Web Services** and GovCloud	●	●	●	●	●†	●	●	● (20G)***
Amazon IC Marketplace		●	●	●				
Microsoft Azure and Government	●	●	●	●	●†	●	● (10G)****	● (10G)****
Google Cloud Platform	●	●	●	●		●	●	●
VMware on IBM Cloud††		●	●	●		●		
Alibaba Cloud International		●	●	●		●	●	
Oracle Cloud Infrastructure†		●	●	●		●	●	

* 10Gbps & HPVE throughput limit applies to non-Internet facing IP traffic only—due to cloud platform ingress throughput limitations.

** Includes VMware on AWS.

*** Achievable using AWS ENA NIC with Gen5 EC2 instances (multi—NIC interfaces and v14.1.x and higher)

**** Achievable using Azure Accelerated Networking (multi-NIC interfaces and v15.0 and higher)

† BYOL only

†† Utility (PAYG) billing only

Please refer to this support matrix on askf5.com to learn more about support for BIG-IP VE in the cloud. You can also leverage the BIG-IP Image Generator Tool to create custom VE images for specific TMOS releases or hot-fixes that may not be available in cloud marketplaces.

F5 BIG-IP Virtual Editions: Simplified Licensing and Choices

F5 virtual editions are available for all BIG-IP modules and can be purchased based on throughput tier from the 10M non-production lab license to the 25 Mbps, 200 Mbps, 1 Gbps, 3 Gbps, 5 Gbps, and 10 Gbps production licenses. As performance requirements increase, F5 offers pay-as-you-grow upgrade licenses. In addition, F5 offers High-Performance VE licenses with no throughput limits and allows you to increase the number of vCPUs to increase performance—up to a maximum of 24 vCPUs.

BIG-IP Virtual Editions are available in a range of licensing models to suit your individual business and budget and budgeting requirements, including:

- **Perpetual (Bring-your-own-license)**—One-time CapEx purchase, supporting 3 major software releases.
- **Subscription**—1- to a 3-year subscription with unlimited version upgrades and premium support included
- **Utility (Pay-as-you-go)**—Hourly or monthly billing for maximum flexibility and no long-term commitment
- **Hybrid-Enterprise Licensing Agreement**—3-year subscription with maximum architectural flexibility across hybrid environments, annual budget protection and premium support included.

The Good, Better, Best bundle offerings from F5 provide you with the best value through flexibility to provision additional advanced application traffic management and security modules as needed.

BIG-IP Cloud Edition

In the age of CI/CD, increased automation, and DevOps methodologies, flexibility, agility, and speed-to-market are the names of the game. Achieving these goals can be difficult, however, especially when different teams are aligned to different priorities.

For teams with app delivery and security concerns, F5 BIG-IP Cloud Edition is a great option. It was designed to enable app teams with self-serve app services in public and private cloud environments. BIG-IP Cloud Edition delivers dedicated, right-sized, and secure F5 application services with per-app manageability and analytics at every stage of the application lifecycle—from initial development to production deployment.

Using BIG-IP Per-App Virtual Editions (VEs) in tandem with BIG-IQ Centralized Management provides app owners, SecOps, and NetOps teams with industry-leading F5 services including traffic management, advanced WAF, role-based control, health monitoring, actionable analytics, and autoscaling. These services are dedicated to individual apps, which means more agility, lower TCO, shorter maintenance windows, and reduced ticket times. It’s the first solution that aligns the priorities of app owners, NetOps, and SecOps.

And it does it through the same scalable, secure, and customizable services provided by traditional F5 VEs—at a price and with a license model (limited to one Virtual IP and three Virtual Servers) appropriate for supporting individual applications—a more flexible and agile way to ensure your apps are always available and performing optimally.

Get Started Today

See for yourself how BIG-IP Virtual Editions can provide an agile, flexible, and efficient way to deploy and optimize application services.

Download the free BIG-IP VE trial

Start testing how you can make your application fast, secure, and available with a full-featured BIG-IP VE—including BIG-IQ Centralized Management—in the environment of your choice. Download a 30-day trial of a BIG-IP VE now. Please review the "Getting Started" documentation.

Get a full evaluation license

Request a free evaluation license to gain access to the latest versions of F5 virtual editions.

Buy BIG-IP for your development lab

Build, test, configure, and stage BIG-IP modules in your development lab.

Try BIG-IP VEs in the public cloud

Try BIG-IP VEs through public cloud providers with free trials and pay-as-you-go hourly billing. See how to get started in AWS, Azure, and GCP by watching the videos.

F5 Global Services

Demands on you and your teams are high. You have to balance implementing business solutions rapidly while maintaining a very high level of solution availability. Accordingly, F5 Global Services and its partners offer world-class consulting, support, and training to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services and its partners can help ensure that your applications scale and are always secure, fast, and available. For more information about F5 Global Services, contact consulting@f5.com or visit f5.com/support.

DevCentral

The F5 DevCentral™ user community of more than 200,000 members is your source for additional technical documentation, discussion forums, blogs, media, and more related to BIG IP virtual editions, application services in virtualized data centers, and cloud deployments.

More Information

To learn more about the BIG-IP family of products, visit f5.com to find these and other resources:

Data sheets

BIG-IP Local Traffic Manager
BIG-IP DNS
BIG-IP Advanced Firewall Manager
BIG-IP Application Security Manager
BIG-IP Access Policy Manager
BIG-IP Carrier-Grade NAT
BIG-IP Policy Enforcement Manager
BIG-IQ Centralized Management
Container Ingress Services

Web pages

Virtual Editions
Cloud Computing
Cloud Solution Templates
F5 on AWS
F5 on Azure
F5 on GCP
F5 on IBM Cloud
F5 on Alibaba
F5 on VMware
F5 on OpenStack
F5 Automation Toolchain

Case studies

American Systems Launches Secure EMNS for Service Members with F5 and Microsoft Azure
Maximus Streamlines Operations with F5 in AWS
Ricacorp Properties Strengthens Website Security with F5 on Microsoft Azure

White papers

Migrating Tier 1 Application Workloads to AWS with F5
How to Add F5 Application Delivery Services to OpenStack
The BIG-IP Platform and Microsoft Azure: Application Services in the Cloud

Overview

VE FIPS Solution Overview

DATA SHEET

BIG-IP Local Traffic Manager

WHAT'S INSIDE

Application Delivery with Scale, Automation, and Customization

Applications drive innovation and profitability, allowing your business to leverage cloud computing, mobility, and software-defined networking (SDN). Your organization from AppDev and DevOps teams to Infrastructure and IT Ops depends on your app services and network infrastructure running at peak performance with app-centric security to meet the challenges of today—and tomorrow.

F5® BIG-IP® Local Traffic Manager™ (LTM) delivers your applications to users in a reliable, secure, and optimized way. You get the extensibility and flexibility of application services with the programmability you need to manage your cloud, virtual, and physical infrastructure. With BIG-IP LTM, you have the power to scale, automate, and customize application services faster and more predictably.

KEY BENEFITS

Scale applications rapidly and reliably
Optimize for today's web applications with HTTP/2 to ensure that your customers and users have access to the applications they need—whenever they need them.

Automate and customize with programmable infrastructure
Control your applications—from connection and traffic to configuration and management—with F5® iRules® LX for network programmability, with Node.js language support in BIG-IP. Use the F5 Automation Toolchain for a declarative approach to efficiently provision, configure, and manage appliances.

Migrate to virtual and cloud environments
Realize operational consistency and comply with business needs across physical, virtual, and cloud environments with deployment flexibility and scalability.

Simplify deployment and management of apps
User-defined F5 iApps® and FAST templates make it easy to deploy, manage, and gain complete visibility into your applications.

Secure your critical applications
Protect the apps that run your business with industry-leading SSL performance and visibility.

RELEVANT RESOURCES

Load Balancing Your Applications

SSL Visibility, Control, and Performance

Deploy Consistent Policies across Any Cloud

Troubleshoot App Performance Issues

APPLICATION INTELLIGENCE

Application Traffic Management

BIG-IP LTM includes static and dynamic load balancing to eliminate single points of failure. Application proxies give you protocol awareness to control traffic for your most important applications. BIG-IP LTM also tracks the dynamic performance levels of servers in a group, ensuring that your applications are not just always on, but also are easier to scale and manage.

Secure Application Delivery

BIG-IP LTM delivers industry-leading SSL performance and visibility for inbound and outbound traffic, so you can cost-effectively protect your entire user experience by encrypting everything from the client to the server. It also defends against potentially crippling DDoS attacks and provides ICAP services for integration with data loss protection and virus protection.

Application Delivery Optimization

BIG-IP LTM scales dramatically, improving page load times and the user experience with HTTP/2, intelligent caching, extensive connection optimization and management, compression, RAMCache performance, F5 TCP Express™, and F5 OneConnect™. It also makes real-time protocol and traffic-management decisions based on application and server conditions, enables rules customization and programmability, and TCP and content offloading.

Application Visibility and Monitoring

Monitor exactly how your application is performing for real users based on application response times, network conditions, and user context. F5 Analytics captures application-specific statistics, such as URL, throughput, and server latency, reported at different levels of the service. BIG-IP LTM makes it simple to integrate with your existing tools using industry standards such as sFlow, SNMP, and syslog.

IoT Protocol Visibility

BIG-IP LTM enables support of IoT clients, publishing useful information to the MQTT Brokers (servers) through the MQTT protocol. The MQTT Brokers then send information to all the subscribers of this information and MQTT support allows BIG-IP LTM to leverage MQTT traffic load balancing for customers' IoT clients.

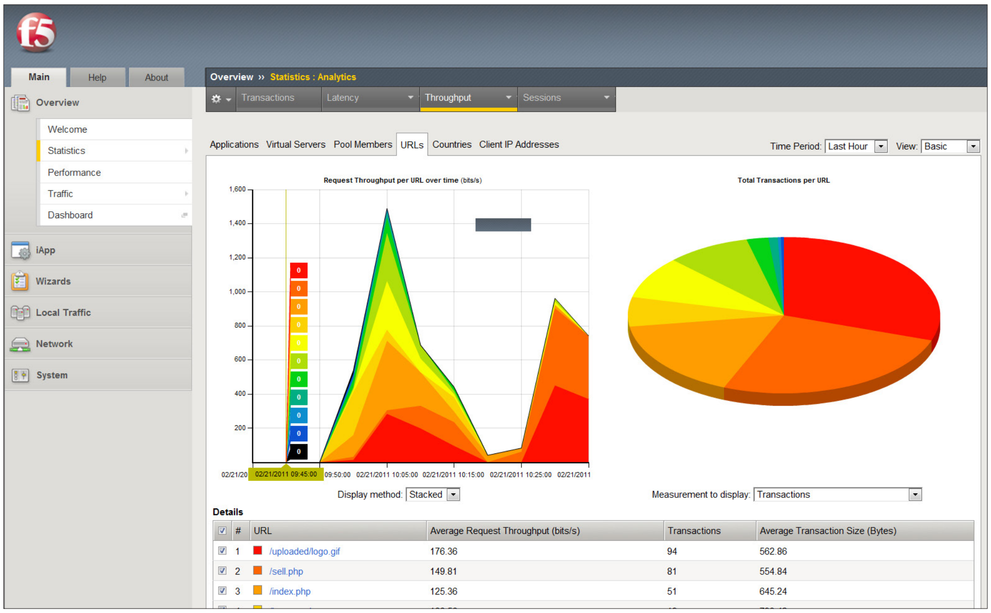

Details

☑	#	URL	Average Request Throughput (bits/s)	Transactions	Average Transaction Size (Bytes)
☑	1	🟥 /uploaded/logo.gif	176.36	94	562.86
☑	2	🟧 /sell.php	149.81	81	554.84
☑	3	🟧 /index.php	125.36	51	645.24

Figure 1: F5 Analytics provides real-time, application-level statistics.

AUTOMATION AND CONTAINER INGRESS

RELEVANT RESOURCE

Integrate into Container Environments

F5 Automation Toolchain allows network and application services such as traffic management and application security to be managed programmatically, through simple, declarative APIs versus traditional manual imperative configurations.

At the core of the F5 Automation Toolchain is the Application Services 3 Extension (AS3) which enables administrators and developers to automate layer 4–7 application services. AS3 also provides a sustainable foundation to enable F5's Infrastructure as Code (IaC) strategy and future integration with third-party orchestration, SDN, and NFV solutions.

F5 Declarative Onboarding enables initial provisioning of F5 solutions, as well as configuration of layer 2–3 objects such as route domains, routes, self IPs, and VLANs. The Declarative Onboarding Extension, like the Application Services 3 Extension, accepts a JSON declaration that defines the desired onboarding end-state via a single REST API.

The F5 Telemetry Streaming Extension is an iControl LX extension that aggregates, normalizes, and forwards statistics and events to consumer applications such as Splunk, Azure Log Analytics, AWS CloudWatch, AWS S3, Graphite, and more. This tool uses a declarative model, meaning you provide a JSON declaration rather than a set of imperative commands.

The F5 API Services Gateway is a TMOS-independent Docker container which runs F5's iControl LX framework and provides a lightweight, fast, portable, TMOS-independent vehicle for customers to leverage iControl LX.

The F5 Automation Toolchain delivers a process-driven approach to automation. Use the components of the Automation Toolchain to efficiently provision, configure, and manage the services that support your apps. The Automation Toolchain is available, free of charge, on GitHub and Docker Hub.

F5 ecosystem integrations with Ansible, Terraform, Puppet, Chef, and Cisco ACI help you simplify orchestration and configuration management across public and private clouds and on-premises, delivering software-defined networking with policy-driven automation and increasing the speed of app deployment through automated provisioning.

F5 Container Ingress Services (CIS) makes it easy to deliver advanced application services to your container deployments, enabling Ingress control HTTP routing, load balancing, and application delivery performance as well as robust security services.

Container Ingress Services easily integrates BIG-IP solutions with native container environments, such as Kubernetes, and PaaS container orchestration and management systems, such as RedHat OpenShift.

PROGRAMMABLE INFRASTRUCTURE

Local Traffic Policies

BIG-IP® local traffic policies are a structured, data-driven collection of rules created by populating tables in a web UI. The policy tables are filled using readable English; no programming skills are required. These policies allow you to inspect, analyze, modify, route, re-direct, discard, or manipulate traffic, and solve common use cases previously covered by simple iRules. For example, you might create a policy that determines whether a client is using a mobile device, and then redirect requests from mobile devices to the applicable mobile web site URL.

iRules

The F5 iRules® scripting language—F5's traffic scripting interface—enables programmatic analysis, manipulation, and detection of all aspects of the traffic in your networks. Customers routinely implement security mitigation rules, support new protocols, and fix application-related errors in real time. With robust and flexible iRules, you can easily and rapidly develop solutions that you can confidently deploy across multiple applications.

iRules LX

iRules LX is the next stage of evolution for network programmability that brings Node.js language support to the BIG-IP platform. Node.js allows JavaScript developers access to over 250,000 npm packages that make code easier to write and maintain. Development teams can access and work on code with the new iRules LX Workspace environment and the new plug-in available for the Eclipse IDE, which can be used for continuous integration builds.

iApps and FAST

F5 iApps and FAST templates are powerful tools that enable you to deploy, manage, and analyze enterprise application services as a whole rather than individually managing configuration and objects. iApps and FAST give you greater visibility into and control over application delivery—and help you deploy in hours rather than weeks. This application-centric approach aligns the network with your applications and adapts application delivery to business needs.

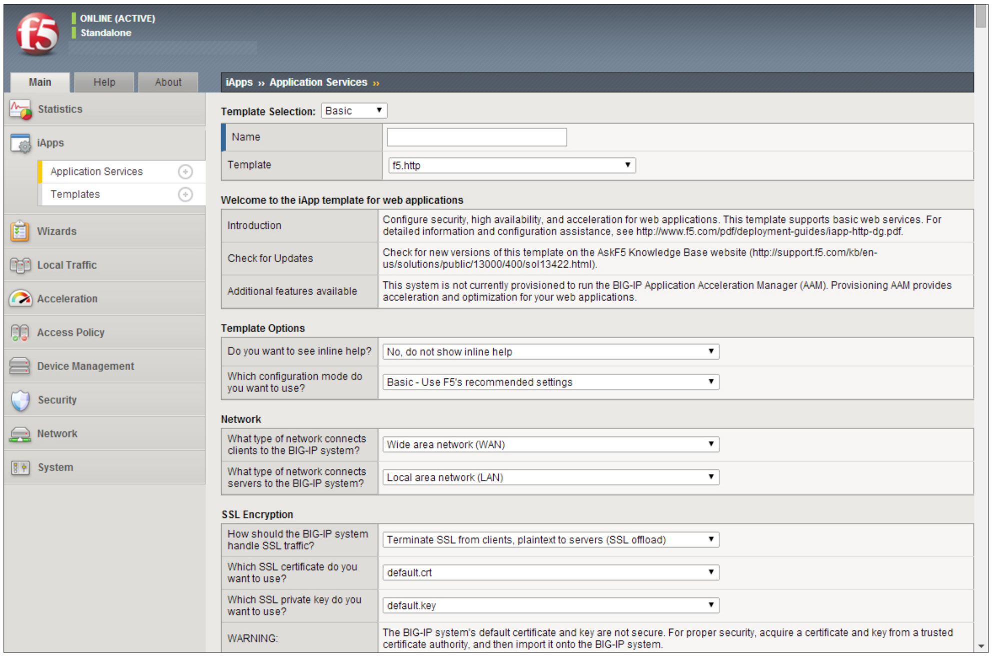

Template Selection: Basic

Name	
Template	f5.http

Welcome to the iApp template for web applications

Introduction	Configure security, high availability, and acceleration for web applications. This template supports basic web services. For detailed information and configuration assistance, see http://www.f5.com/pdf/deployment-guides/iapp-http-dg.pdf.
Check for Updates	Check for new versions of this template on the AskF5 Knowledge Base website (http://support.f5.com/kb/en-us/solutions/public/13000/400/sol13422.html).
Additional features available	This system is not currently provisioned to run the BIG-IP Application Acceleration Manager (AAM). Provisioning AAM provides acceleration and optimization for your web applications.

Template Options

Do you want to see inline help?	No, do not show inline help
Which configuration mode do you want to use?	Basic - Use F5's recommended settings

Network

What type of network connects clients to the BIG-IP system?	Wide area network (WAN)
What type of network connects servers to the BIG-IP system?	Local area network (LAN)

SSL Encryption

How should the BIG-IP system handle SSL traffic?	Terminate SSL from clients, plaintext to servers (SSL offload)
Which SSL certificate do you want to use?	default.crt
Which SSL private key do you want to use?	default.key
WARNING:	The BIG-IP system's default certificate and key are not secure. For proper security, acquire a certificate and key from a trusted certificate authority, and then import it onto the BIG-IP system.

Figure 2: iApps Templates simplify application deployments.

iControl

The F5 iControl® APIs and SDK allow automation and integration of custom applications into all aspects of BIG-IP LTM and other BIG-IP modules. iControl is delivered as both REST and SOAP APIs to fit the model best suited for your organization. With iControl, every aspect of BIG-IP LTM configuration, including most aspects of all BIG-IP modules—from device and application provisioning to application tuning and health and support initiation—can be programmatically automated to achieve dynamic infrastructures.

iCall

F5 iCall® is a powerful scripting framework, based on TMSH (the F5 TMOS® Shell command-line interface) and Tcl, that helps customers maintain their environment and reduce downtime by automating tasks. It monitors for events and executes scripts to resolve issues quickly and predictably. iCall enables administrators to react to specified events by executing services on the management plane, such as generating a TCP stack dump on a failure, executing a specific iApp to reconfigure application network service settings, or adjusting load balancing weights on application services based on a change in health monitoring data.

RELEVANT RESOURCES

Optimize ADC Performance

Consistent App Services in Any Cloud

SCALABLE INFRASTRUCTURE

Cloud-Ready

BIG-IP LTM makes it easy to realize operational consistency and comply with business needs across physical, virtual, and cloud environments, removing the friction of migrating applications between traditional physical and cloud architectures. Available in public clouds and for migration across multi-cloud. Learn more in the BIG-IP Virtual Edition data sheet.

ScaleN

F5 ScaleN® technology uses the F5 VIPRION® chassis, Device Service Clusters, and the scaling capabilities of F5 Virtual Clustered Multiprocessing™ (vCMP) to enable more efficient, elastic, and multi-tenant solutions for data centers, clouds, and hybrid deployments. ScaleN moves beyond traditional infrastructure limitations and offers multiple scalability and consolidation models to help you meet your specific business needs.

Virtual Networking

The BIG-IP® SDN Services module natively supports VXLAN and NVGRE to offer gateway capabilities with BIG-IP LTM bridging virtual and traditional networks. This lets you keep things simple, applying application delivery network services across both virtual and traditional networks.

Advanced Routing

The BIG-IP® Advanced Routing™ Module allows BIG-IP LTM to provide network routing capabilities such as BGP, RIP, OSPF, ISIS, and BFD for enhanced interoperability within the network, increasing the resilience and capacity of your network.

ATTACK MITIGATION

Deploy F5 Distributed Cloud Bot Defense directly from your BIG-IP

Bots cause significant financial pain through scraping that slows performance, scalping and inventory hoarding that frustrate loyal customers, enumerating gift card codes to steal balances, creating fake accounts to commit fraud, and credential stuffing—the testing of stolen credentials—that leads to account takeovers.

Today’s advanced, persistent bots are more sophisticated than ever. To stay ahead of attackers, F5 Distributed Cloud Bot Defense uses rich client-side signal collection, industry-leading code obfuscation, aggregate telemetry collection, and AI for unparalleled long-term efficacy and near-zero false positives while maintaining access for good bots. And because F5 defends the most targeted sites on the web—including those of the world’s largest banks, retailers, and airlines, F5 is ready when these attacks target your organization.

Deploy F5 Distributed Cloud Bot Defense directly from your BIG-IP or through a connector that’s right for your application, with support services tailored to your needs, from self-service to managed service.

BIG-IP PLATFORMS

Only F5's next-generation, cloud-ready ADC platform provides DevOps-like agility with the scale, depth of security, and investment protection needed for both established and emerging apps. The BIG-IP® iSeries appliances deliver quick and easy programmability, ecosystem-friendly orchestration, and record-breaking, software-defined hardware performance. As a result, customers can accelerate private clouds and secure critical data at scale while lowering TCO and future-proofing their application infrastructures. F5 solutions can be rapidly deployed via integrations with open source configuration management tools and orchestration systems.

In addition to the iSeries, F5 offers VIPRION modular chassis and blade systems designed specifically for performance and for true, on-demand, linear scalability without business disruption. VIPRION systems leverage F5's ScaleN clustering technology so you can add blades without reconfiguring or rebooting.

The F5 VELOS® platform is the next generation of F5's industry-leading chassis-based systems, which delivers unprecedented performance and scalability in a single Application Delivery Controller (ADC). You can seamlessly scale capacity by adding modular blades in a chassis, without disruption, and VELOS allows a mix of traditional BIG-IP tenants as well as next-generation BIG-IP tenants in the future.

The next-generation ADC solution, F5 rSeries, bridges the gap between traditional and modern infrastructures with a rearchitected, API-first platform designed to meet the needs of your traditional and emerging applications. The new F5 rSeries delivers unprecedented levels of performance, a fully automatable architecture, and the highest reliability, security and access control for your critical applications.

Virtual Editions (VEs) of BIG-IP software run on commodity servers and support the range of hypervisors and performance requirements. VEs provide agility, mobility, and fast deployment of app services in software-defined data centers and cloud environments.

See the BIG-IP System Hardware, VIPRION, VELOS, and Virtual Edition data sheets for more details. For information about specific module support for each platform, see the latest release notes on AskF5. For the full list of supported hypervisors, refer to the VE Supported Hypervisors Matrix.

With BIG-IQ® Centralized Management, you can manage F5 platforms with a single-pane-of-glass view, including:

- BIG-IP iSeries Appliances
- rSeries Appliances
- BIG-IP Virtual Editions
- VIPRION Chassis
- VELOS Chassis

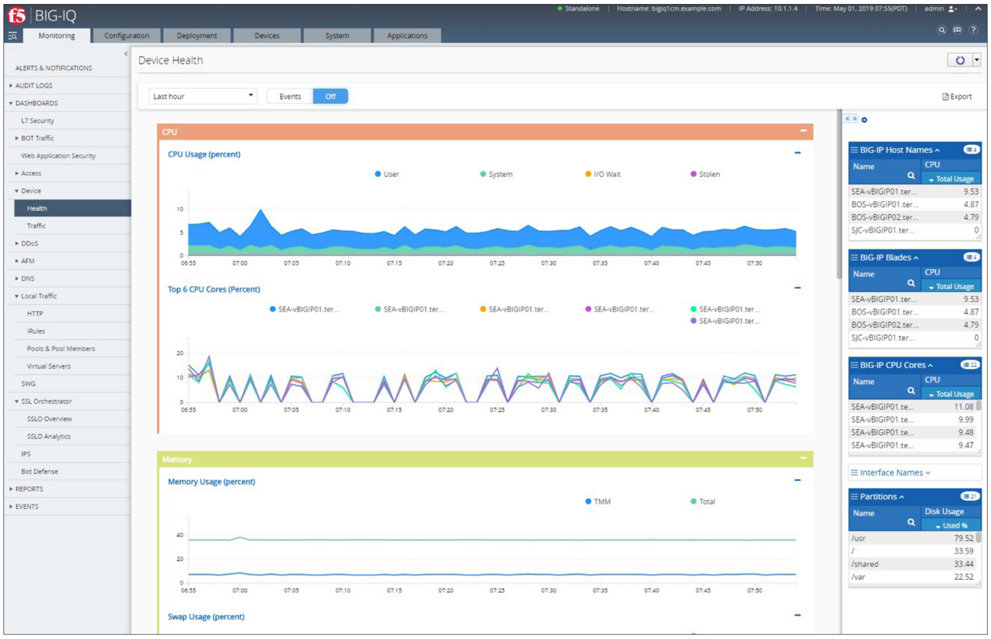

BIG-IP Host Names	CPU Total Usage
SEA-vBIGIP01.ter...	9.53
BOS-vBIGIP01.ter...	4.87
BOS-vBIGIP02.ter...	4.79
SJC-vBIGIP01.ter...	0

BIG-IP Blades	CPU Total Usage
SEA-vBIGIP01.ter...	9.53
BOS-vBIGIP01.ter...	4.87
BOS-vBIGIP02.ter...	4.79
SJC-vBIGIP01.ter...	0

BIG-IP CPU Cores	CPU Total Usage
SEA-vBIGIP01.te...	11.08
SEA-vBIGIP01.te...	9.99
SEA-vBIGIP01.te...	9.48
SEA-vBIGIP01.te...	9.47

Partitions	Disk Usage Used %
/usr	79.52
/	33.59
/shared	33.44
/var	22.52

Figure 3: Manage your BIG-IP appliance health and track CPU and memory usage of physical and virtual platforms, hardware blades, and cores with BIG-IQ. Use logging and reporting to understand overall trends and spot areas needing correction. Easily manage policies, certificates, and licensing management to push out to all BIG-IP ADCs for centralized control of app services infrastructure.

FLEXIBLE LICENSING TO MEET YOUR NEEDS

To align with differing purchasing directives, BIG-IP Next may also be licensed via an assortment of consumption models. Choose the licensing model that best suits your needs, including subscription, perpetual, or the utility program:

- **Subscription**—Renewable one- to three-year subscriptions afford initial upfront savings and include access to F5 premium support.
- **Perpetual**—A one-time CapEx investment provides complete solution ownership.
- **Utility**—Pay-as-you-go model includes access to F5 premium support without the need for a long-term commitment.

F5 GLOBAL SERVICES

F5 Global Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services can help ensure your applications are always secure, fast, and reliable. For more information about F5 Global Services, contact consulting@f5.com or visit f5.com/support.

DEVCENTRAL

The F5 DevCentral® technical community is an active and engaged source for the best technical, how-to articles, discussion forums, shared code, media, and more related to programmability and Application Delivery Networking.

BIG-IP LTM FEATURES

Application Traffic Management

- Intelligent load balancing
- Application protocol support (HTTP/2, SSL/TLS, SIP, etc.)
- Application health monitoring
- Application connection state management
- F5 OneConnect
- Advanced routing (BGP, RIP, OSPF, ISIS, BFD)
- SDN services (VXLAN, NVGRE)

Application Delivery Optimization

- Symmetric adaptive compression
- RAM cache and compression
- TCP Express
- HTTP/2 gateway

Secure Application Delivery

- SSL connection and session mirroring
- Hybrid crypto services (Hardware SSL offload for BIG-IP VE)
- SSL/TLS encryption offload (hardware accelerated)
- Algorithm agility (GCM, ECC, Camellia, DSA, RSA)
- Suite B support including forward secrecy
- Internal/Network/Cloud HSM (FIPS 140-2)
- SSL visibility

Application Visibility and Monitoring

- F5 Analytics
- Performance dashboard
- High-speed logging
- sFlow

Programmable Infrastructure

- iRules and iRules LX for data plane programmability
- iCall for event-based control-plane scripting
- iApps for app-level config management and deployment
- iControl for Management API (SOAP, REST)

Automation and Container Ingress

- Automation Toolchain for declarative apps services configurations
- Application Services 3 Extension (AS3) automates Layer 4-7 services
- Declarative Onboarding for initial provisioning and configurations
- Telemetry Streaming for data stream export to 3rd party analytics
- FAST templates for declarative app services configurations
- Container Ingress Services for automation of container app services

Ecosystem Integrations

- Ansible templates for app services automation
- Terraform modules for deployment automation
- Cisco ACI and F5 BIG-IP for integrated network fabric and control
- Puppet for automation of configurations and app services
- Chef for configuration management integrations
- F5 Distributed Cloud Bot Defense for attack mitigation

ScaleN

- On-demand scaling
- All-active application clustering

MORE INFORMATION

To learn more about BIG-IP LTM, visit f5.com to find these and other resources.

Web

BIG-IP Local Traffic Manager
DevCentral

Data sheets

BIG-IP System Hardware
BIG-IP Virtual Editions
rSeries
VELOS
VIPRION

Articles and guides

State of Application Strategy 2021: Unpacking the Current and Future State of Application Security and Delivery
3 Tips for Maintaining a High-Performing App Portfolio
Choose advanced cloud solutions that will scale into the future
Load Balancing Your Applications

Case studies

Varolii: SaaS Provider Ensures High Uptime and Resiliency for Critical Customer Apps with F5
Motorists Insurance Group Gives Customers A Seamless Experience With F5 + Okta Solution
Pandora Scales to Serve Tens of Millions of Internet Radio Users with F5 Solution
MarketAxess Increases Productivity with F5 and Ansible Automation

Use cases

Deploy Consistent Policies across Any Cloud
Integrate into Container Environments
Troubleshoot App Performance Issues
Integrate into CI/CD Pipelines

DATA SHEET

BIG-IP DNS

WHAT'S INSIDE

Hyperscale and Protect Your DNS While Optimizing Global App Delivery

Scaling and securing every environment helps protect your business from site outages and improves DNS and application performance. Securing DNS infrastructures from the latest distributed denial-of-service (DDoS) attacks and protecting DNS query responses from cache-poisoning redirects will help keep your business online and viable. To fully achieve these goals, you need efficient ways to monitor DNS infrastructure and application health, and to scale on-demand.

F5® BIG-IP® DNS distributes DNS and user application requests based on business policies, data center and cloud service conditions, user location, and application performance. The BIG-IP platform delivers F5's high-performance DNS services with visibility, reporting, and analysis; hyperscales and secures DNS responses geographically to survive DDoS attacks; delivers a real-time DNSSEC solution; and ensures high availability of global applications in all cloud environments.

KEY BENEFITS

Hyperscale DNS up to 100 million RPS with a fully loaded chassis
BIG-IP DNS hyperscales authoritative DNS up to 100 million query responses per second (RPS) and controls DNS traffic. It ensures that users are connected to the best site and delivers on-demand scaling for DNS and global apps.

Protect against DNS attacks and ensure availability
Ensure DNS and application availability and protection during DNS DDoS attacks or volume spikes. Mitigate DNS threats by blocking access to malicious IP domains.

Improve global application performance
Send app users to the cloud or on-premises site with the best performance based on application, geolocation, business, and network conditions.

Deploy flexibly, scale as you grow, and manage efficiently
BIG-IP DNS delivers flexible global application management in virtual and multi-cloud environments. The web-based UI provides easy DNS configuration with centralized menus; advanced logging, statistics, and reporting along with export to 3rd party analytics.

UNMATCHED DNS PERFORMANCE

BIG-IP DNS delivers hyperscale performance that can handle even the busiest apps and websites. When apps have a volume spike in DNS queries due to legitimate requests or DDoS attacks, BIG-IP DNS manages requests with multicore processing and F5 DNS Express™, dramatically increasing authoritative DNS performance up to 50 million RPS to quickly respond to all queries.

This scalability helps your organization provide the best quality of service (QoS) for your users while eliminating poor application performance. DNS Express improves standard DNS server functions by offloading DNS responses as an authoritative DNS server. BIG-IP DNS accepts zone transfers of DNS records from the primary DNS server and answers DNS queries authoritatively.

Benefits and features of multicore processing and DNS Express include:

- High-speed response and DDoS attack protection with in-memory DNS
- Authoritative DNS replication in multiple BIG-IP or DNS service deployments for faster responses
- Authoritative DNS and DNSSEC in multi-clouds for disaster recovery and fast, secure responses
- Scalable DNS performance for quality of app and service experience
- The ability to consolidate DNS servers and increase ROI

In cases of very high volumes for apps and services or a DNS DDoS attack, BIG-IP DNS with DNS Express enabled plus in Rapid Response Mode (RRM) hyperscales up to 100 million RPS. It extends availability with unmatched performance and security—absorbing and responding to queries up to 200 percent of the normal limits. See page 17 for performance metrics and details.

DNS CACHING AND RESOLVING

DNS latency can be reduced by enabling a DNS cache on BIG-IP DNS and having it respond immediately to client requests. BIG-IP DNS can consolidate the cache and increase the cache hit rate. This reduces DNS latency up to 80 percent, with F5 DNS Caching reducing the number of DNS queries for the same site. When used in hardware on the F5 VIPRION® platform, DNS caching hyperscales for ultimate query response performance and delivers linear scalability across multi-bladed chassis. In addition to caching, BIG-IP DNS allows the device to do its own DNS resolving without requiring the use of an upstream DNS resolver.

Caching profiles available to select for multiple caches include:

- Transparent cache
- BIG-IP DNS site between client and DNS internal/external
- Hot cache
- Caching resolver
- No cache response - BIG-IP DNS sends out requests with responses returned for resolving and caching
- Validating caching resolver

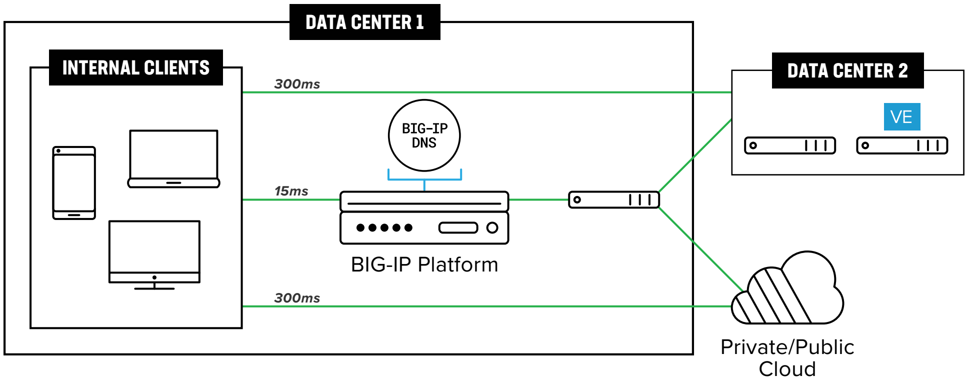

Figure 1: BIG-IP DNS supports all common DNS deployments that are either authoritative or locally resolved DNS. Specific zone requests not cached are forwarded to name servers for faster DNS resolving, allowing users to receive expedient responses.

BIG-IP DNS reduces the average DNS response time and latency for mobile and desktop devices from an average of 300 milliseconds (ms) and 100 ms respectively to as little as 15 ms, depending on workloads.

SECURE APPLICATIONS

DNS denial-of-service attacks, cache poisoning, and DNS hijacking threaten the availability and security of your applications. BIG-IP DNS protects against DNS attacks and enables you to create polices that provide an added layer of protection for your applications and data.

DNS attack protection features include:

- Hardened device—BIG-IP DNS is ICSA Labs Certified as a network firewall, and resists common teardrop, ICMP, and daemon attacks.
- DNS attack protection—BIG-IP DNS offers built-in protocol validation in software to automatically drop high-volume UDP, DNS query, NXDOMAIN floods, and malformed packets. You can use BIG-IP DNS in hardware to mitigate these high-volume attacks.

- DNS load balancing—The BIG-IP platform can be used to front-end static DNS servers. If the DNS request is for a name controlled by the BIG-IP platform, F5 DNS services will answer the request.
- Security control—F5 iRules® for DNS can help you create policies that block requests from rogue sites.
- Packet filtering—BIG-IP DNS uses packet filtering to limit or deny websites' access based on source, destination, or port.

DNS firewall

DNS DDoS, cache poisoning of LDNS, and other unwanted DNS attacks and volume spikes can cause DNS outage and lost productivity. These attacks and traffic spikes increase volume dramatically and can take down DNS servers.

BIG-IP DNS, with security, scale, performance, and control functionality, provides DNS firewall benefits. It shields DNS from attacks such as reflection or amplification DDoS attacks and other undesired DNS queries and responses that reduce DNS performance.

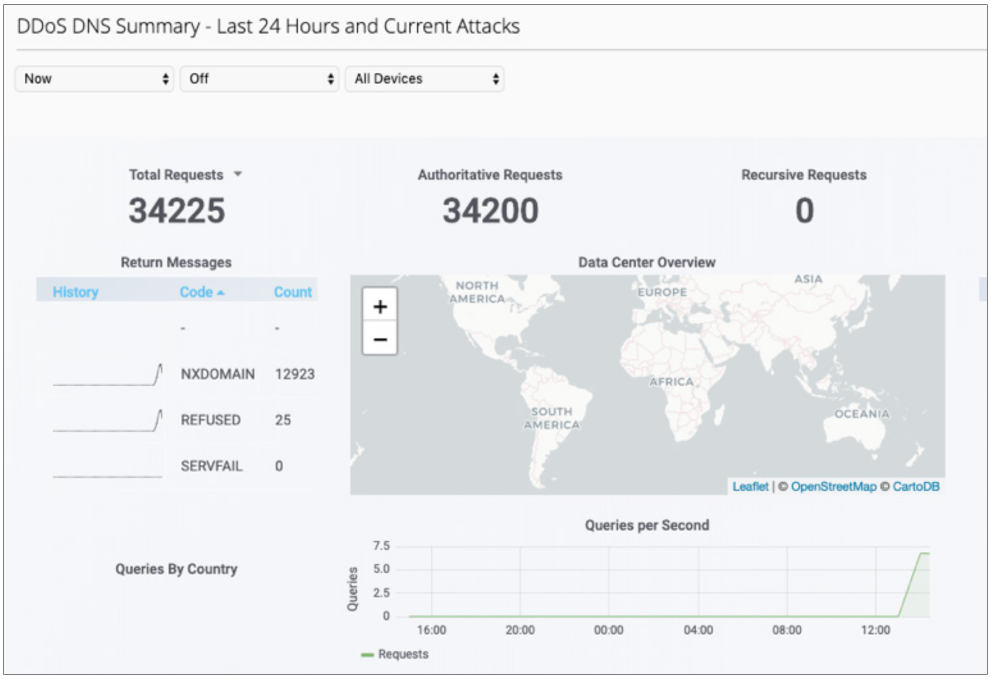

DDoS DNS Summary - Last 24 Hours and Current Attacks

Now | Off | All Devices

Total Requests 34225

Authoritative Requests 34200

Recursive Requests 0

Return Messages

History	Code	Count
	-	-
	NXDOMAIN	12923
	REFUSED	25
	SERVFAIL	0

Figure 2: Visualize DNS DDoS attacks, the top 25 attack URLs, Queries Per Second (QPS) and by country as available, and other DNS traffic such as Reponses Per Second (RPS) by record types for a full picture of your DNS performance and unwanted attacks.

In addition, you can mitigate complex DNS security threats by blocking access to malicious IP domains with Response Policy Zones. With BIG-IP DNS, you can install a third-party domain filtering service such as SURBL or Spamhaus and prevent client infection or intercept infected responses to known sources of malware and viruses. F5 DNS firewall services reduce the costs of infection resolution and increase user productivity.

F5 DNS SERVICES INCLUDE:

- Protocol inspection and validation
- DNS record type ACL*
- High-performance authoritative DNS, which scales responses exponentially
- Authoritative DNS hyperscaling up to 200% to absorb DDoS attacks
- Reducing latency and hyperscaling DNS caching
- DNS load balancing
- Stateful inspection (never accepts unsolicited responses)
- ICSA Labs certification (can be deployed in the DMZ)
- The ability to scale across devices using IP Anycast
- Secure responses (DNSSEC)
- DNSSEC response rate limits
- DNS over HTTPS support resolves queries and mitigates attacks
- Complete DNS control using DNS iRules
- DDoS threshold alerting*
- Threat mitigation by blocking access to malicious IP domains
- DNS logging and reporting
- Hardened F5 DNS code (not BIND protocol)

*Requires provisioning BIG-IP® Advanced Firewall Manager™ to access functionality.

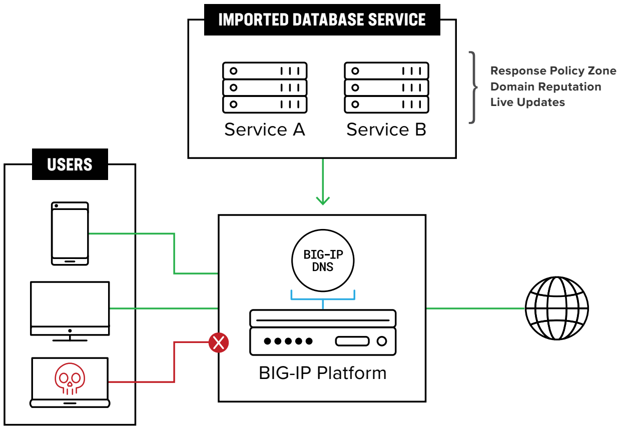

Figure 3: Lower your risk of malware and virus communication and mitigate DNS threats by blocking access to malicious IP domains with a domain reputation service such as SURBL or Spamhaus.

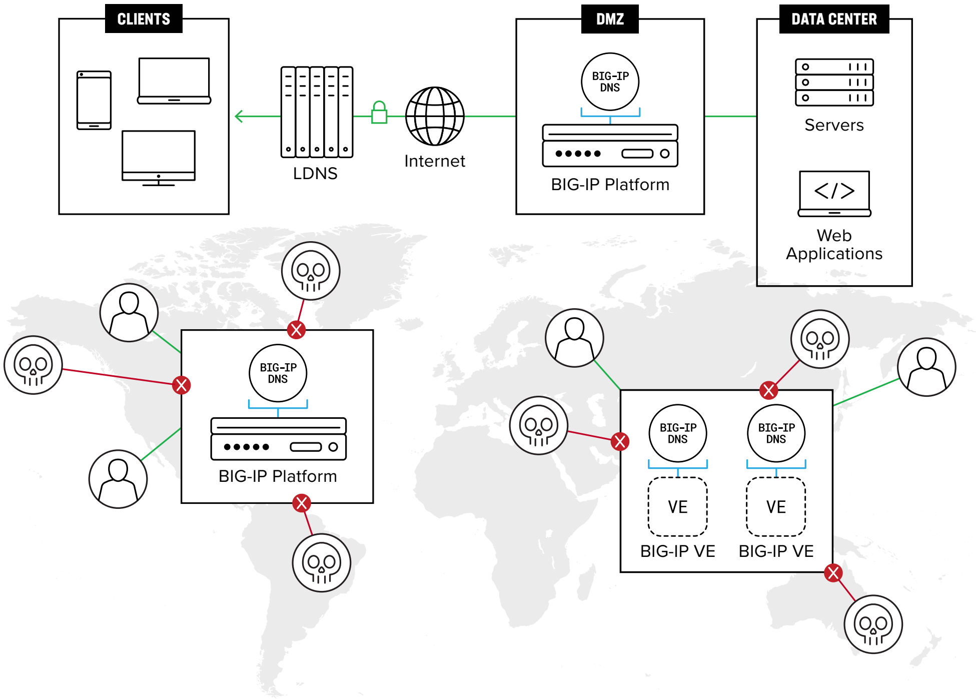

Figure 4: BIG-IP DNS keeps apps available with firewall services protecting DNS infrastructure from high-volume attacks and malformed packets.

Complete DNSSEC signing

With BIG-IP DNSSEC support, you can digitally sign and encrypt your DNS query responses. This enables the resolver to determine the authenticity of the response, preventing DNS hijacking and cache poisoning. In addition, receive all the benefits of global server load balancing while also securing your DNS query responses. Alternatively, if a zone has already been signed, BIG-IP DNS manages static DNSSEC responses for higher performance.

Centralized DNSSEC key management

Many IT organizations have or want to standardize on FIPS-compliant devices and secure DNSSEC keys. You can use BIG-IP DNS with FIPS cards that provide 140-2 support for securing your keys. In addition, BIG-IP DNS integrates and uses hardware security modules (HSMs) from Thales for implementation, centralized management, and secure handling of DNSSEC keys, reducing OpEx and delivering consolidation and FIPS compliance.

Top-level domain support for DNSSEC

For DNS administrators who want to delegate to other secure sub-domains, BIG-IP DNS allows easy management of DNSSEC as a top-level domain, becoming a parent zone.

DNSSEC validation

In most networks, DNS resolvers offload DNSSEC record requests and crypto calculations to validate that the DNS response being received is correctly signed. DNSSEC responses coming into the network require high CPU loads on DNS resolving servers.

DNS over HTTPS

DNS over HTTPS (DOH) is encrypted DNS using SSL for full protection. It is fully enabled by popular web browsers and can create delays and security issues for service providers and enterprises who are not able to terminate and respond to these DNS inquires. F5 BIG-IP DNS allows your network to unencrypt and resolve DNS queries over HTTPS without impacting responses-per-second (RPS). In addition, DoH support removes HTTPS as a DNS Spoofing vector for malicious amplification attacks and protects the last mile with DNS message encapsulation.

DNS over TLS

DNS over TLS (DoT) ensures that DNS requests and responses are not tampered with or forged via on-path attacks. DoT adds TLS encryption on top of the transmission control protocol (TCP), which is used for DNS queries. DoT is a protocol that authenticates communication between a DNS client and a DNS server. It uses cryptographic signatures for secured transmission. DNS over TLS or DoT, is a standard for encrypting DNS queries to keep them secure and private. DoT uses the same security protocol, TLS to encrypt and authenticate communications.

DEPLOY F5 DISTRIBUTED CLOUD BOT DEFENSE DIRECTLY FROM YOUR BIG-IP

Bots cause significant financial pain through scraping that slows performance, scalping and inventory hoarding that frustrate loyal customers, enumerating gift card codes to steal balances, creating fake accounts to commit fraud, and credential stuffing—the testing of stolen credentials—that leads to account takeovers.

Today's advanced persistent bots are more sophisticated than ever. To stay ahead of attackers, F5 Distributed Cloud Bot Defense uses rich client-side signal collection, industry-leading code obfuscation, aggregate telemetry collection, and AI for unparalleled long-term efficacy and near-zero false positives while maintaining access for good bots. And because F5 defends the most targeted sites on the web—including those of the world's largest banks, retailers, and airlines—F5 is ready when these attacks target your organization.

Deploy Distributed Cloud Bot Defense directly from BIG-IP or through a connector that's right for your application, with support services tailored to your needs, from self-service to managed service.

Advanced global load balancing

BIG-IP DNS includes the industry's most advanced traffic distribution capabilities to match the needs of any organization or globally deployed application.

- Round robin
- Global availability
- LDNS persistence
- Application availability
- Geography
- Virtual server capacity
- Least connections
- Packets per second
- Round trip time
- Hops
- Packet completion rate
- User-defined QoS
- Dynamic ratio
- LDNS
- Ratio
- Kilobytes per second

With BIG-IP DNSSEC validation, administrators can easily offload and validate DNSSEC on the client side using BIG-IP DNS for resolving. This results in superior DNS performance and a dramatic increase in the site response to users.

GLOBALLY AVAILABLE APPLICATIONS

BIG-IP DNS offers global application availability and sophisticated health monitoring that support a wide variety of application types, giving organizations the flexibility to adapt quickly and stay competitive.

These global availability and health monitoring features include:

- Global load balancing—BIG-IP DNS provides comprehensive, high-performance application management for hybrid environments.
- Dynamic ratio load balancing—BIG-IP DNS routes users to the best resource based on site and network metrics (for example, based on the number of hops between the client and the local DNS).
- Wide area persistence—To ensure user connections persist across apps and data centers, BIG-IP DNS synchronizes data, propagates local DNS, and maintains session integrity.
- Geographic load balancing—BIG-IP DNS includes an IP database identifying location at the continent, country, and state/province level to connect users to the closest app or service for the best performance.
- Custom topology mapping—With BIG-IP DNS, organizations can set up custom topology maps. By defining and saving custom region groupings, you can configure topology based on intranet app traffic policies that match your internal infrastructure.
- Infrastructure monitoring—BIG-IP DNS checks entire infrastructure health, eliminating single points of failure and routing app traffic away from poorly performing sites.

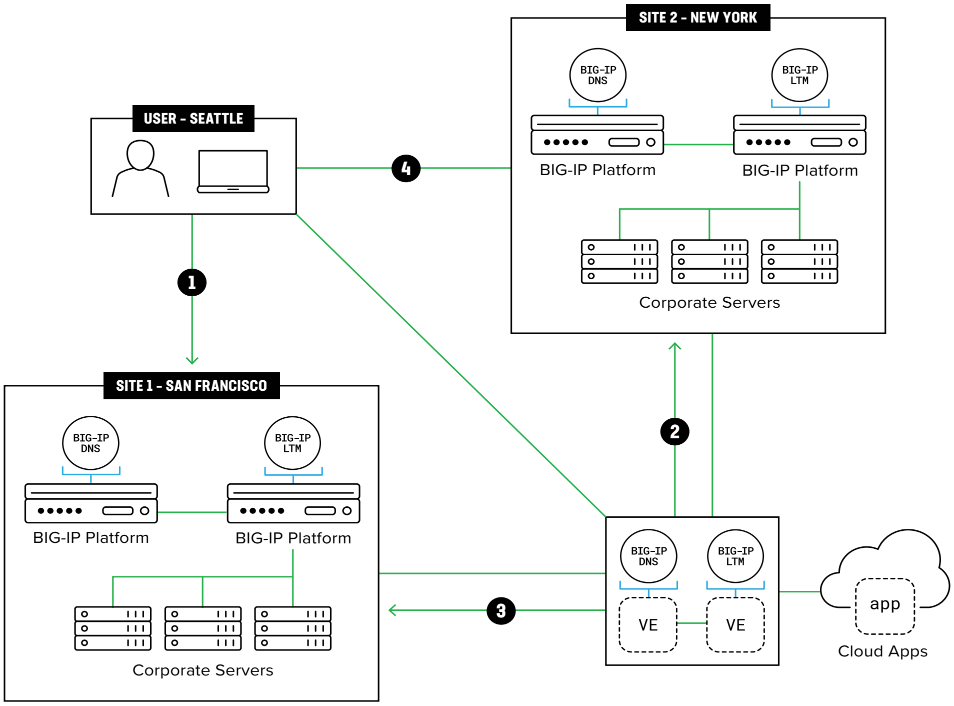

Figure 5: BIG-IP DNS ensures users are always connected to the best site. 1 User queries local DNS to resolve domain, and local DNS queries BIG-IP DNS. 2 BIG-IP DNS uses metrics collected for each site and identifies the best server. 3 BIG-IP DNS responds to local DNS with IP address. 4 User is connected to best site on premises or in multi-cloud.

Application health monitoring

BIG-IP DNS improves the application experience by intelligently monitoring the availability of resources. It expands application resilience by flexibly selecting and using the best available BIG-IP solutions for health monitoring. BIG-IP DNS reduces application downtime and enables easy availability with multiple settings in application monitoring.

Today’s sophisticated applications require intelligent health checks to determine availability. Instead of relying on a single health check, BIG-IP DNS aggregates multiple monitors so that you can check the application state at multiple levels. This results in the highest availability, improves reliability, and eliminates false positives to reduce management overhead.

BIG-IP DNS provides pre-defined, out-of-the-box health monitoring support for more than 18 different applications, including SAP, Oracle, LDAP, and mySQL. BIG-IP DNS performs targeted monitoring of these applications to accurately determine their health, reduce downtime, and improve the user experience.

Disaster recovery/business continuity planning

In addition to performing comprehensive site availability checks, you can define the conditions for shifting all traffic to a backup data center, failing over an entire site, or controlling only the affected applications.

SIMPLE MANAGEMENT

Managing a distributed, multiple-site network from a single point is an enormous challenge. BIG-IP DNS provides tools that give you a global view of your infrastructure with the means to manage the network and add polices to ensure the highest availability for your business-critical applications. Manage global infrastructure from a centralized UI with features including:

- Web-based user interface.
- Streamlined and centralized DNS and GSLB menus for fast configuration.
- Efficient list and object management for complete visibility of global resources.
- Unique naming of objects to reduce administration and build business policies.
- Enhanced management of distributed applications as part of one collective group.
- Context-sensitive help for information on objects, commands, and configuration examples.

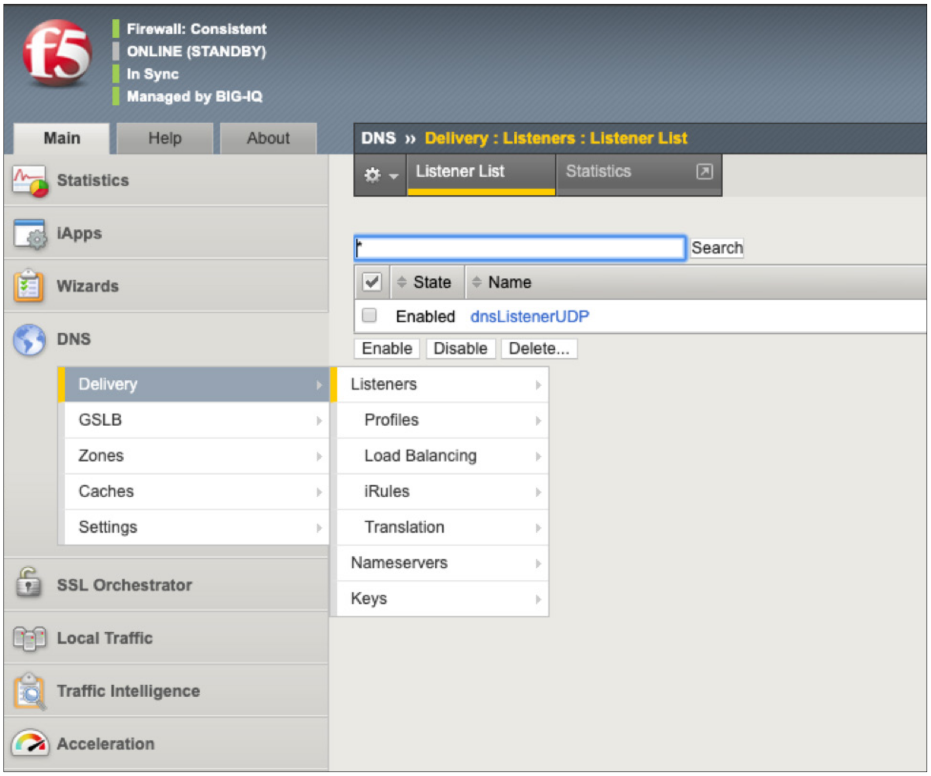

☑	State	Name
□	Enabled	dnsListenerUDP

Figure 6: Reduce DNS delivery deployment time with centralized and easy-to-find configuration and management sequences.

- **Powerful command-line interface**—The TMSH command-line interface delivers integrated search, context-sensitive help, and batch-mode transactions.
- **Automated setup and synchronization**—Autosync automates and secures multiple BIG-IP DNS devices, eliminating difficult hierarchical management common to DNS.
- **Improves scale and analysis with unlimited N+1 devices**—In a failover situation, when BIG-IP DNS services are part of a Device Service Cluster (DSC) group, the BIG-IP solution performs at its peak capacity—across all appliances or virtual editions synced with DNS and GSLB services. BIG-IP DNS provides highly scalable apps and services, performing smart analysis on all incoming traffic to better understand patterns and anomalies.
- **Scalable and optimized GSLB configurations**—Incremental Sync delivers high performance for large deployments. With more devices synced, configuration changes transpire rapidly. For large deployments with GSLB configurations and rapid user changes, you can protect changes by manually saving when most convenient.
- **Configuration retrieval**—AutoDiscovery enables retrieved configurations from distributed BIG-IP instances, removing repeat configurations across devices.
- **Data center and sync groups**—Create logical groups of network equipment to ensure efficient use of monitoring and metrics collection for intelligently sharing with members in the logical group.

- **Distributed application management**—You can define dependencies between application services and manage them as a group, building scalable traffic distribution policies and improving efficiency with granular control of objects.
- **iRules**—Use the F5 iRules scripting language to customize the distribution of global traffic. BIG-IP DNS looks deep inside DNS traffic to customize app traffic to the desired data center, pool, or virtual server. This reduces latency, increases attack protection, and improves performance.
- **Customize traffic with QoS**—Design traffic decisions and easily develop custom load balancing algorithms using quality of service metrics in iRules, such as round-trip time, hops, hit ratio, packet rate, topology, and more.
- **DNS iRules**—Manage DNS queries, responses, and actions for a fast, customized DNS infrastructure. For instance, configure DNS iRules with filtering for protection and reporting.
- **F5 ZoneRunner**™—ZoneRunner is an integrated DNS zone file management tool that simplifies and reduces the risk of misconfiguration. Built on the latest version of BIND, ZoneRunner provides:
    - Auto population of commonly used protocols.
    - Validation/error checking for zone file entries.
    - Zone importation from an external server or a file.
    - Automatic reverse lookups.
    - Easy creation, editing, and searching of all records.
    - Easy management of NAPTR records for LTE and 4G requirements.

Load balancing across container environments

When migrating to public cloud and containers, engineers and architects need a robust solution for load balancing between clusters (global load balancing) across multi-cloud deployments. F5 BIG-IP DNS targets applications in container clusters for scale, routing, and security services by enabling Server Name Indication (SNI) support when using HTTPS monitor.

DNS health monitor

The DNS health monitor available in BIG-IP DNS and BIG-IP® Local Traffic Manager™ (LTM) monitors DNS server health and helps configure DNS based on reporting. The DNS health monitor detects whether the servers are operating at peak performance and helps in reconfiguring for optimal responses.

High-speed logging

You can easily manage DNS and global app logging for fast network visibility and planning. High-speed logging of DNS queries and responses, syslog, and global server load balancing decision logs improve information on data to enable fast network recognition with quick, deep search and display.

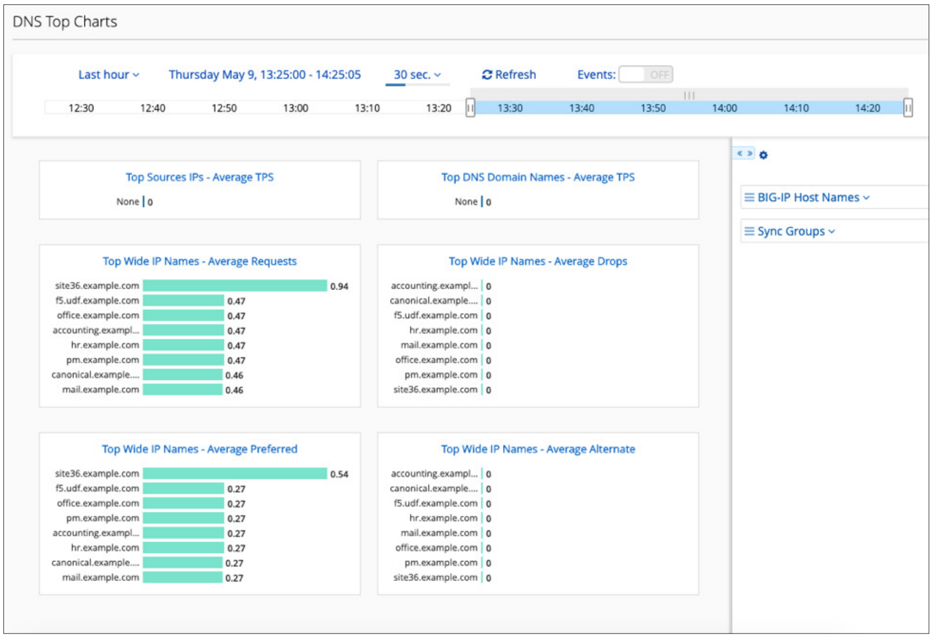

Figure 7: Understand your DNS health from seconds to years comparing the top source IPs, Domain and wide IP names and the TPS for reporting.

Enhanced DNS detailed statistics

BIG-IP DNS delivers advanced DNS statistics for administrators, with enhanced detailed data for profiles such as query type counts (A, CNAME, NS, RRSIG, AAAA, SRV, and “other” types) with requests, responses, and percentage counts. Stats are per profile and per device global count for fast visibility and capacity planning of DNS delivery infrastructure. DNS detail stats are viewable in DNS profile or in analytics reporting.

Advanced DNS reporting and analytics

F5 Analytics provides advanced DNS reporting and analysis of applications, virtual servers, query names, query types, client IPs, top requested names, and more for business intelligence, capacity planning, ROI reporting, troubleshooting, performance metrics, and tuning, enabling maximum optimization of the DNS and global app infrastructure.

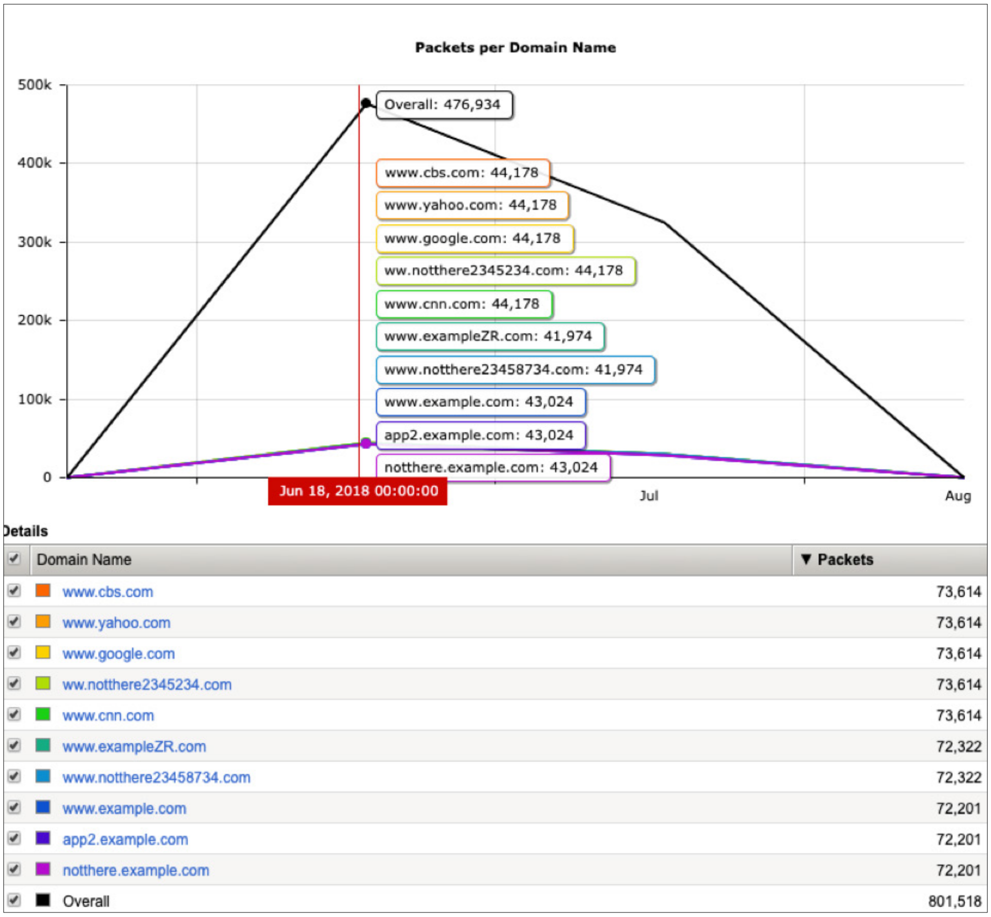

☑		Domain Name	▼ Packets
☑	🟧	www.cbs.com	73,614
☑	🟧	www.yahoo.com	73,614
☑	🟨	www.google.com	73,614
☑	🟩	ww.notthere2345234.com	73,614
☑	🟩	www.cnn.com	73,614
☑	🟩	www.exampleZR.com	72,322
☑	🟦	www.notthere23458734.com	72,322
☑	🟦	www.example.com	72,201
☑	🟪	app2.example.com	72,201
☑	🟪	notthere.example.com	72,201
☑	⬛	Overall	801,518

Figure 8: Administrators can easily manage DNS using analytics with advanced reporting and analysis of actions for fast visibility of DNS delivery and infrastructure.

BIG-IQ Centralized Management

F5® BIG-IQ® Centralized Management provides a central point of control for F5 physical and virtual devices and for the solutions that run on them. It simplifies management, helps ensure compliance, and gives you the tools you need to deliver your applications securely and effectively. BIG-IQ manages BIG-IP DNS licenses, policies, SSL certificates, images, and configurations.

BIG-IQ offers BIG-IP DNS centralized management including the ability to create, retrieve, update, and delete all global server load balancing (GSLB) objects; tools to deploy and rollback GSLB policies; and the ability to manage DNS listener and profile configurations.

ADVANTAGES OF DNS VISIBILITY

- View and manage configuration and policies on DNS devices.
- Add BIG-IP DNS and BIG-IP LTM devices to existing sync groups.
- Analyze F5 iQuery® connection information to help identify DNS sync-group issues.
- View high-level statistics across your DNS infrastructure showing status of DNS sync groups and devices.
- View both real-time and historical DNS statistics.

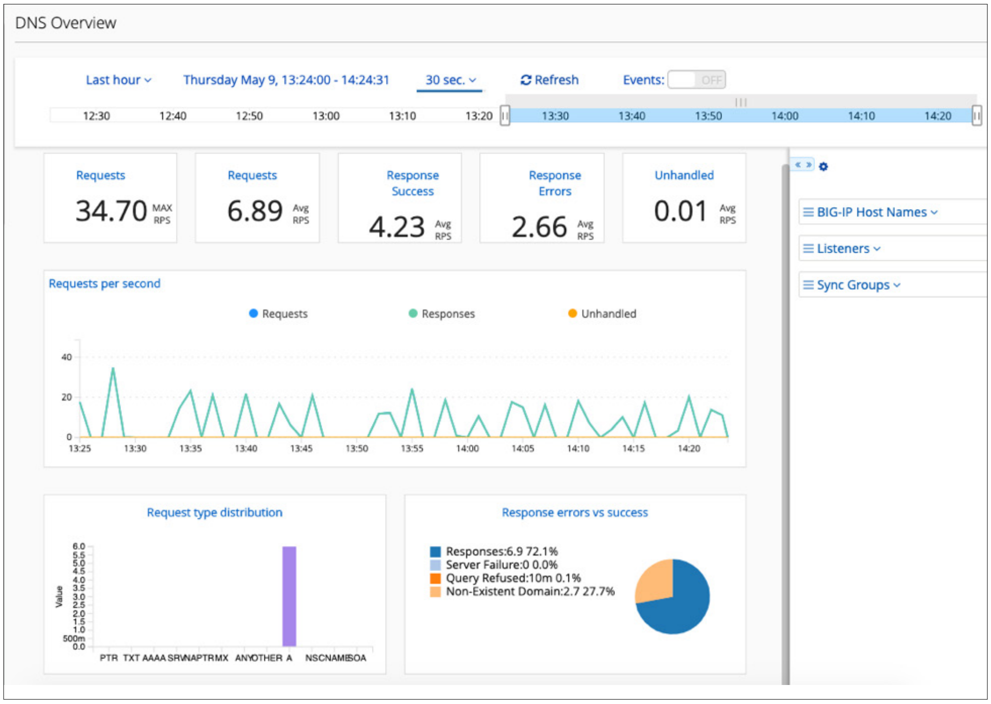

Figure 9: BIG-IQ integration showcases advanced DNS visibility with requests, responses, and unhandled query insights for in-depth analytics.

NETWORK INTEGRATION

BIG-IP DNS is designed to fit into your current network and into your plans for the future. Integration features include:

- **SNMP management application support**—BIG-IP DNS integrates its MIBs and an SNMP agent with DNS. This enables SNMP management applications to read statistical data about BIG-IP DNS performance.
- **Third-party integration**—BIG-IP DNS communicates and integrates with a broad array of network devices. This includes support for various types of remote hosts, such as SNMP agents, third-party caches, servers, routers, and load balancers to diagnose the health of network endpoints.
- **IPv6/IPv4 support**—Ease the transition to IPv6 by providing DNS gateway and translation services for hybrid IPv6 and IPv4 solutions and managing IPv6 and IPv4 DNS servers. BIG-IP LTM configured with NAT64 transforms IPv6 to IPv4 for those IPv4-only environments.

- **IP Anycast integration**—DNS query volumes directed to one IP address, whether legitimate or during a DoS attack, are easily managed by distributing the load among multiple geographic BIG-IP DNS devices. Network managers realize these benefits:
  - Improved user performance and reliability
  - Reduced network latency for DNS transactions
  - Ability to scale DNS infrastructure to manage DNS request load to one IP address
  - Lower rates of dropped query packets, reducing DNS timeouts/retries
  - Increased revenue because more users are serviced and brand equity is protected

Figure 10: BIG-IP DNS and IP Anycast integration distributes the DNS request load by directing single IP requests to multiple local devices.

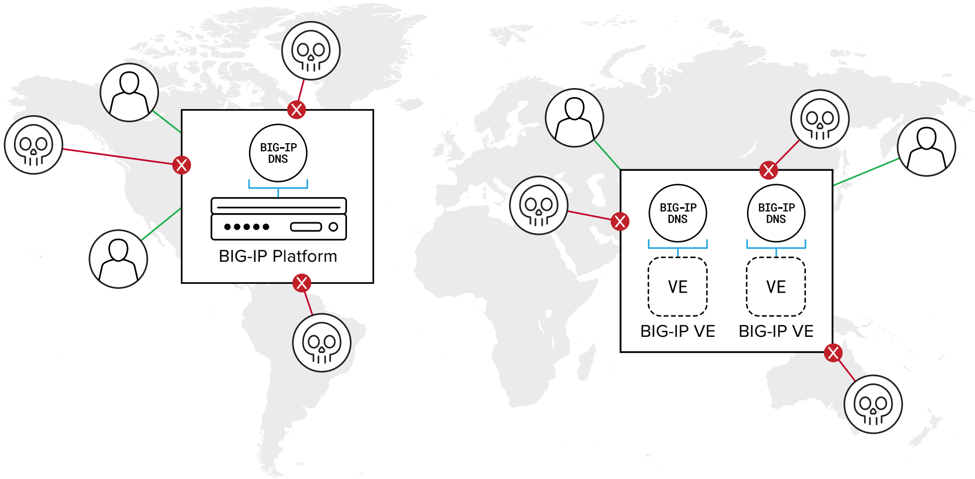

- **Global server load balancing in virtual and cloud environments**—Easily spin up virtual instances of BIG-IP DNS. Provide flexible DNS delivery and global application availability by routing users to applications in physical, virtual, and cloud environments.

ARCHITECTURE

The advanced architecture of BIG-IP DNS gives you total flexibility to control application delivery without creating traffic bottlenecks.

The BIG-IP DNS architecture includes:

- **TMOS**®—The F5 operating system, TMOS, provides a unified system for optimal DNS and application delivery, giving you total visibility, flexibility, and control across all BIG-IP services.
- **Query and response performance and scalability**—Linearly scale on larger platforms and multi-bladed chassis for increased performance by integrating functions in TMOS. BIG-IP DNS is can be provisioned for platforms that support F5 Virtual Clustered Multiprocessing™ (vCMP®).

BIG-IP PLATFORMS

Only F5's next-generation, cloud-ready ADC platform provides DevOps-like agility with the scale, security depth, and investment protection needed for both established and emerging apps. The new BIG-IP® iSeries appliances deliver quick and easy programmability, ecosystem-friendly orchestration, and record-breaking, software-defined hardware performance. Customers can now accelerate private clouds and secure critical data at scale while lowering TCO and future-proofing their application infrastructures. F5 solutions can be rapidly deployed via integrations with open source configuration management tools and orchestration systems.

In addition to the iSeries, F5 offers VIPRION modular chassis and blade systems designed specifically for performance and for true on-demand, linear scalability without business disruption. The new VELOS platform is the next generation of F5's industry-leading chassis-based systems, which delivers unprecedented performance, scalability, and customization in a single ADC system. BIG-IP® virtual edition (VE) software runs on commodity servers and supports the broadest range of hypervisors and performance requirements. VEs provide agility, mobility, and fast deployment of app services in software-defined data centers and cloud environments.

The next-generation Application Delivery Controller (ADC) solution, F5 rSeries, bridges the gap between traditional and modern infrastructures with a rearchitected, API-first platform designed to meet the needs of your traditional and emerging applications. The new F5 rSeries delivers unprecedented levels of performance, a fully automatable architecture, and the highest reliability, security, and access control for your critical applications.

See the BIG-IP System Hardware, VIPRION, VELOS, and Virtual Edition data sheets for details. For information about specific module support for each platform, see the latest release notes on AskF5. For the full list of supported hypervisors, refer to the VE Supported Hypervisors Matrix.

In addition, F5 offers BIG-IQ® Centralized Management for single-pane-of-glass management of all F5 devices for enabling orchestration of F5 application delivery policies.

BIG-IP iSeries Appliances

VIPRION Chassis

rSeries Appliances

VELOS Chassis

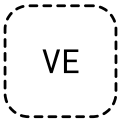

BIG-IP Virtual Editions

F5 CLOUD SERVICES

Now you can easily provision and configure the services your apps need in minutes. Built in a pay-as-you-go model, F5 Cloud Services offer predictable pricing, flexibility and the ability to auto-scale to meet changing demand. Have your on-premises DNS services and zone transferred to F5 Cloud Services:

- **DNS Cloud Service**—Easily provision and configure DNS services with a few clicks; begin responding to queries within minutes of activation.
- **DNS Load Balancer Cloud Service**—Ensure high availability and application performance with the simple, intelligent F5 DNS Load Balancer Cloud Service.

DNS ON-DEMAND SCALING

Administrators have the option to add DNS and GSLB on-demand scaling with rate limit and object limit capacity as desired to BIG-IP DNS or BIG-IP LTM appliances. This option supports requirements for exact traffic performance, resulting in lower CapEx and OpEx. On-demand scaling includes the following services: DNS, GSLB, and DNSSEC. User interface statistics show rated capacity of instances, such as query RPS and object limits, which deliver fast traffic detail for easy capacity planning. Contact your regional F5 sales representative or reseller for more information.

DNS QUERY RPS MAXIMUM PERFORMANCE

BIG-IP DNS services deliver query response per second (RPS) with high performance scalability. The table below lists many BIG-IP platforms with DNS Express enabled for authoritative DNS query response with the maximum capabilities per platform.

Platform	Max Query RPS
Virtual Edition	250,000*
i2600	240,000
i2800	500,000
i4600	480,000
i4800	880,000
i5600	1,000,000
i5800, i5820 FIPS	1,500,000
i7600	1,500,000
i7800, i7820 FIPS	2,300,000
10150s NEBS	990,000
10350v FIPS, NEBS	1,800,000
i10600	2,000,000
i10800	2,900,000
i11400-DS	1,200,000
i11600 (DS)	1,800,000
i11800 (DS)	4,500,000
i15600	4,100,000
i15800	8,100,000
VELOS CX410 Full Chassis (8 blades)	
VELOS BX110 Blade	2,500,000
VIPRION 2200 Full Chassis (2 blades) VIPRION 2400 Full Chassis (4 blades)	
VIPRION B2250 Blade	2,100,000
VIPRION 4480 Full Chassis (4 blades) VIPRION 4800 Full Chassis (8 blades)	
VIPRION B4450 Blade	6,400,000

*BIG-IP DNS Virtual Edition is available in increments of 250,000 RPS. For 5050s and above, Rapid Response Mode (RRM—see page 2) delivers up to 200 percent of normal max query RPS when turned on. See F5 Sales or reseller for details.

SIMPLIFIED LICENSING

Meeting your applications' needs in a dynamic environment has never been easier. F5's provides you with the flexibility to provision advanced modules on-demand, at the best value.

- Decide what solutions are right for your application's environment with F5's solutions.
- Specify the subscriptions you need across hybrid-cloud environments.
- Provision the modules needed to run your applications with F5's Good, Better, Best offerings.
- Spin up or down any solution needed with Enterprise Licensing Agreements.
- Implement complete application flexibility with the ability to deploy your modules on a virtual or physical platform.

F5 GLOBAL SERVICES

F5 Global Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services can help ensure your applications are always secure, fast, and available. For more information about F5 Global Services, contact consulting@f5.com or visit F5 Professional Services.

DEVCENTRAL

The F5 DevCentral™ technical community is an active and engaged source for the best technical, how-to articles, discussion forums, shared code, media, and more related to DNS delivery and global app networking.

MORE INFORMATION

To learn more about BIG-IP DNS, search on f5.com to find these and other resources.

Web pages

DNS Delivery
Global Server Load Balancing
DevCentral

Data sheet

BIG-IP System Hardware Data sheet
BIG-IP Virtual Editions
VIPRION
VELOS
rSeries Hardware Data sheet

Articles and videos

Intro to DNS services and global server load balancing (video)
DNS is key to connected customers
Addressing Cloud-based DNS—It’s time to move
Encrypted DNS - Mitigating the Impact of DoT and DoH for Service Providers
What Is a DNS Amplification Attack?

Case studies

Bank improves user experience with always available, fast, and secure access to banking services using F5
SaaS provider ensures high uptime and resiliency for critical customer apps with F5
Everbridge manages traffic and security across global cloud providers and local data center
Shawbrook Bank enlists F5 to accelerate and scale digital transformation

F5 Advanced WAF

DATA SHEET

What's Inside

Proactive Application Protection

Applications are critical to your business. Without the right protection, however, they can become an attack vector that may ultimately lead to a data breach. Consider this alarming statistic: Organizations have an average of 765 web applications and these applications are the initial target of data breaches 53% of the time.[1]

Protect your organization and its reputation by maintaining the confidentiality, availability, and performance of the applications that are critical to your business with F5® Web Application Firewall (WAF) solutions.

F5 WAF solutions are deployed in more data centers than any enterprise WAF on the market. The comprehensive suite of F5 WAF solutions includes managed rulesets for Amazon Web Services (AWS); cloud-based, self-service, and managed service in the F5 Silverline® cloud-based service delivery platform; application delivery controller (ADC) integration with F5 BIG-IP® Application Security Manager™ (ASM)[2]; and F5 Advanced Web Application Firewall™ (Advanced WAF).

Advanced WAF redefines application security to address the most prevalent threats organizations face today:

- Automated attacks and bots that overwhelm existing security solutions.
- Web attacks that steal credentials and gain unauthorized access across user accounts.
- Application layer attacks that evade static security based on reputation and manual signatures.
- New attack surfaces and threats due to the rapid adoption of APIs.

Advanced WAF is built on proven F5 technology and goes beyond reactive security such as static signatures and reputation to proactively detect and mitigate bots, secure credentials and sensitive data, and defend against application denial-of-service (DoS).

Advanced WAF delivers flexible and comprehensive protections wherever apps reside and without compromising performance. Advanced WAF is offered as an appliance, virtual edition, and as a managed service—providing automated WAF services that meet complex deployment and management requirements while protecting your apps with great precision. It is the most effective solution for guarding modern applications and data from existing and emerging threats while maintaining compliance with key regulatory mandates.

[1] 2018 Application Protection Report

[2] BIG-IP ASM continues to be offered through F5 Good/Better/Best licensing.

Key benefits

Protect web and mobile applications from malicious bots

F5 secures an organization's most valued assets, applications, and sensitive data from bots, automated attacks, web scrapers, and exploits. Advanced WAF extends bot protection to mobile applications through the F5 Anti-Bot Mobile SDK, providing rapid deployment of mobile bot protection through an easy-to-use web portal without requiring any changes to the application or mobile device. Applications fused with mobile bot protection are supported in vendor and third-party application stores.

Safeguard credentials and sensitive data from theft and abuse

Advanced WAF secures credentials and sensitive data from theft and abuse, preventing data breaches and mitigating automated attacks that leverage previously stolen credentials. F5 BIG-IP DataSafe™ application layer encryption in Advanced WAF masks sensitive fields directly within the user's web browser, rendering data stolen by bad actors through client-side attacks useless. Using BIG-IP DataSafe, customers can encrypt data at the field level transparently, without requiring any changes on clients or Web servers. Comprehensive brute force mitigation including credential stuffing protection defends against automated attacks that leverage previously stolen credentials.

Defend against sophisticated application denial-of-service (DoS)

Advanced WAF discovers and fingerprints new and unusual traffic patterns without human intervention, distinguishing and isolating potential malicious traffic from legitimate traffic. This automated mitigation capability is based on a continuous feedback loop of client behavior and server stress. If anomalous behavior is detected, Advanced WAF automatically builds a dynamic signature and begins mitigating the attack. The effectiveness of the mitigation is then monitored through the continuous feedback loop. False positives are reduced while accuracy and performance are improved through continuous mitigation tuning as the attack starts, evolves, or stops.

Mitigate sophisticated threat campaigns

Threat Campaigns provide targeted signatures to protect organizations from pervasive attacks that are often coordinated by organized crime and nation states. Based on F5 Labs research, Threat Campaigns provide critical intelligence to fingerprint and mitigate sophisticated attacks with nearly real-time updates. Metadata is used to determine both malicious requests and malicious intent, and the high accuracy of Threat Campaign signatures immediately blocks active threats with low false positives and no learning cycle.

Protect APIs

As web applications expand from connected to collaborative via the extensive use of Application Programming Interfaces (APIs), Advanced WAF ensures that API methods are enforced on URLs. It also secures applications against API attacks that commonly go undetected by traditional firewalls. With a unique defense mechanism that guards XML, JSON, and GTW APIs through rate limiting, behavioral analysis, and anti-automation, Advanced WAF automatically detects application program interface threats, enforces strict policy rules for each use case, and blocks attacks and special content types—closing the back door on application threats. With F5 Access Manager™, API protection is improved through comprehensive authentication and token enforcement.

Ensure application security and compliance

Gain comprehensive security against sophisticated layer 7 attacks, blocking threats that evade traditional WAFs and enabling compliance with key regulatory mandates.

Turn on protection immediately

Simplify security with pre-built policies, thousands of out-of-the-box signatures, and a streamlined approach to policy management that decreases operational expenses.

Patch vulnerabilities fast

Identify and resolve app vulnerabilities in minutes with leading dynamic application security testing (DAST) integration and automatic virtual patching.

Deploy flexibly

Deploy as an appliance, in virtual or cloud environments, and as a managed service supporting multi-tenant services while incorporating external intelligence that secures against known IP threats.

Defend with proven advanced protections

Defend with highly programmable technology that dynamically adapts policies, proactively stops bots and DoS attacks, and demonstrates 99.89% overall security effectiveness.

Magnify threat knowledge

Easily understand your security status with detailed forensic analysis, full visibility into HTTP and WebSocket traffic, and rich insight into all events and user types.

Ensure Comprehensive Threat Protection

The volume and sophistication of attacks makes keeping up-to-date on security threat types and protection measures a challenge for application administrators and security teams. With industry-leading capabilities and superior flexibility, F5 Advanced WAF delivers advanced, cost-effective security for the latest web and mobile applications.

Advanced WAF protects credentials from theft and abuse, and secures any parameter from client-side manipulation by validating login parameters and application flow to prevent forceful browsing and logical flaws. It also allows organizations to effectively guard against existing and emerging layer 7 application attacks—preventing costly data breaches, thwarting DoS attacks, and maintaining compliance. Advanced WAF is the first leading WAF that supports the transition from AJAX/HTTP to WebSockets for greater efficiencies and less overhead with bi-directional streaming data. Advanced WAF also provides visibility into WebSocket traffic—enabling companies to transition to protecting chat sessions and streaming information feeds (such as stock tickers) from data exposure, tampering, and theft. Users benefit from an extensive database of signatures, dynamic signature updates, DAST integration, and the flexibility of F5 iRules® scripting for customization and extensibility.

Organizations rely on Advanced WAF to protect the world's most visited web applications wherever they reside, with the highest level of security and without compromising performance. Advanced WAF enables organizations to detect and mitigate layer 7 threats including web scraping, web injection, brute force, CSRF, JSON web threats, DoS-heavy URLs, and zero-day attacks—providing early warnings, while mitigating threats per policy.

It automatically defends against multiple, simultaneous application-layer threats including stealthy, low-bandwidth DoS attacks. Advanced WAF also stops in-browser session hijacking and reports regular and repeated attacks from IPs.

Using automatic learning capabilities, dynamic profiling, unique anomaly detection methods, and risk-based policies, Advanced WAF can impose needed protections to prevent even the most sophisticated attacks from ever reaching servers. When combined with F5 BIG-IP Local Traffic Manager™ (LTM), Advanced WAF filters attacks and accelerates applications for an improved user experience.

Continuous expert security research

F5's security research team helps ensure continuous development of Advanced WAF signatures, policies, and capabilities. Researchers explore forums and third-party resources, investigate attacks, reverse engineer malware, and analyze vulnerabilities to determine effective detection and mitigation methods that guard against zero-day threats, DoS attacks, and other evasive or evolving threats. Advanced WAF offers enhanced protection from advancements in technology, regular signature updates, threat intelligence, and tightening of existing capabilities.

Defend with proactive bot protections

An always-on defense is required to successfully identify and protect against automated DoS attacks, web scraping, and brute force attacks before they occur. F5 delivers proactive bot defense capabilities that effectively provide controls to help prevent these attacks from ever taking place. Using advanced defense methods and reputation matching to identify non-human users (such as JavaScript and CAPTCHA challenges, geolocation enforcement, and other techniques), Advanced WAF slows requests to distinguish bots and then drops those requests before they reach a server. Advanced WAF thoroughly inspects user interaction, analyzes the health of the server, and discerns transaction anomalies to help detect bots that may bypass client/application challenges, established rate limits, and other standard detection methods. It also automatically mitigates layer 7 attacks that show an unusual change in request patterns. Unique from other solutions, Advanced WAF provides security experts with greater control of bot defense enforcements, allowing them to force additional action (such as high-speed logging on block or challenge actions, JavaScript challenges, URI overrides, customized HTML redirects, and more) before mitigations are applied. The Advanced WAF bot defense capabilities provide the most effective prevention methods, allowing you to identify suspicious automated activity, categorize bots detected, and mitigate attacks with the highest level of precision. The F5 Anti-Bot Mobile SDK, in conjunction with Advanced WAF, extends F5's comprehensive bot protection to mobile applications without any changes to application code.

Track malicious user attempts

Distinguishing permitted users from bad actors whenever a website is visited helps minimize security risk and prevent malicious activity. With Advanced WAF, application security teams can employ device identification tracking techniques to identify specific end-users, application sessions, and attackers. This unique capability allows IT to easily distinguish human traffic from bot traffic, spot repeat visitors, prevent malicious attempts, and help WAFs more accurately mitigate brute force, session hijacking, web scraping, and DoS attacks.

Device identification tracking enables Advanced WAF to identify the same browser, even when users switch sessions or source IPs. When activated, Advanced WAF captures and saves unique device characteristics and attributes determines which clients are suspicious, and mitigates threats based on predefined settings. Whether an automated threat, DoS attack, headless browser, or human user, Advanced WAF can distinguish between repeat attackers and customer visitors for every WAF use case.

Block malicious IP addresses

Delivering today's rich and complex Internet content to users can expose an organization to a variety of potentially malicious attacks from rapidly changing IP addresses. Inbound and outbound botnet traffic, such as DoS and malware activity, can penetrate the organization's security layers. F5 IP Intelligence Services enhances automated security decisions with IP reputation intelligence. By identifying IP addresses and security categories associated with malicious activity, IP Intelligence Services can incorporate dynamic lists of threatening IP addresses from third parties into the F5 platform, adding context and automation to Advanced WAF blocking decisions. This adds granularity to Advanced WAF rules—allowing administrators to set an alarm, stop traffic, or fully block IPs based upon a specific IP reputation category while allowlisting approved IP addresses.

Additionally, Advanced WAF alleviates computational heavy mitigation of threats from known malicious IP addresses with a unique IP shun capability (accelerated denylisting). Instead of wasting cycles on traffic from badly behaving IPs, Advanced WAF immediately denylists IPs that repeatedly fail challenges or undergo high block ratios. This temporarily blocks malicious IPs in hardware at the network layer until IP intelligence feeds are up to date.

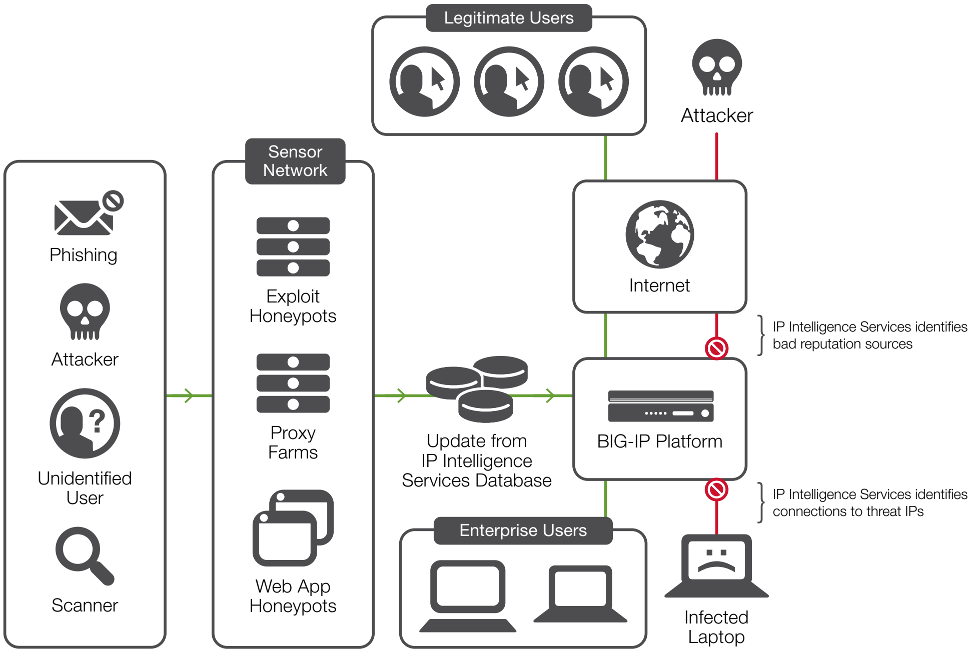

IP Intelligence Services gathers reputation data for use by F5 solutions.

Enabling secure encryption

As the increasing demand for data protection drives growth in encrypted traffic, it is important to transition to Perfect Forward Secrecy (PFS) while guarding against SSL/TLS attacks that threaten the security of applications and information in transit. Advanced WAF protects against malicious attempts to overcome SSL/TLS and compromise private keys, user passwords, and other sensitive information. It provides full SSL/TLS termination, and decrypts and re-encrypts terminated traffic—allowing complete inspection and mitigation of concealed, malicious threats. When Advanced WAF is combined with BIG-IP LTM, organizations also gain comprehensive SSL/TLS DDoS mitigation and SSL/TLS offload protection to secure against SSL/TLS attacks including SSL floods, POODLE, Heartbleed, and various memory-cracking tools.

Identify anomalous behavior

With Advanced WAF, IT can easily detect traffic that does not conform to normal behavior and evades usual volumetric protections—such as an uncommon increase or decrease in latency or the transactions rate. Advanced WAF can identify and uniquely block excessive failures to authenticate IP addresses generating a high volume of login attempts, as well as other anomalies in the typical traffic pattern. These include sessions opened at high rates or requesting too much traffic. Behavioral analytics and machine learning in Advanced WAF automatically monitor client and server traffic for anomalies in a continuous feedback loop.

Patch vulnerabilities immediately

Advanced WAF integrates with leading web application vulnerability scanners to allow you to easily manage assessments, discover vulnerabilities, and apply specific policies from a single location. These unique capabilities facilitate near-instantaneous mitigation of application assessment results, ensuring protection while developers correct vulnerable code—patching in minutes instead of weeks or months. With Advanced WAF, administrators can import testing results from DAST scanners, including scanners from WhiteHat, IBM, and QualysGuard, and layer a vulnerability-driven policy (received from F5 scanner integrations) on top of a current rapid deployment or SharePoint policy. When combined with WhiteHat Sentinel, Advanced WAF also detects and reports recent website changes to the scanner. This ensures scanning of otherwise overlooked URLs and parameters, and the application of specific policies—enabling organizations to secure their applications immediately after updating.

Advanced WAF DAST support helps IT deliver next-generation website security using simple, accurate, automated services. These services protect assets in a dynamic threat environment with more comprehensive assessments, zero false positives, and more manual and automated virtual patches than any other WAF solution.

Enforce geolocation-based blocking

Attacks are increasing from a variety of global sources. Advanced WAF enables you to block these attacks based on geolocation: states, countries, or regions. Administrators can easily select allowed or disallowed geolocations for strong policy enforcement and attack protection. Geolocation-based blocking also protects against anomalous traffic patterns from specific countries or regions, and enables traffic throttling based on location. Advanced WAF geolocation-based protection can be applied to a CAPTCHA challenge and to protect RAM cache and other resources from DDoS attacks.

Inspect SMTP and FTP

Advanced WAF enables SMTP and FTP security checks to protect against spam, viral attacks, directory harvesting, and fraud. Using default settings, administrators can easily configure security profiles to inspect FTP and SMTP traffic for network vulnerabilities and protocol compliance. Default settings can also be used to trigger alarms or block requests for violations.

SMTP security checks enable validation of incoming mail using several criteria, while disallowing or allowing common call methods used to attack mail servers. Additionally, administrators can set rate limits on the number of incoming messages, create allowlists and denylists, and validate DNS SPF records. FTP violations can be triggered for anonymous, passive, or active requests; specific FTP commands; command line length; and excessive login attempts. Administrators can use default SMTP/FTP settings for easy setup or customize profiles to address specific risks and more effectively ensure protocol compliance.

Streamline Learning, Deployment, and Management

Organizations want to turn on protections immediately without extensive security expertise. F5 Advanced WAF simplifies and automates configuration and policy deployment with pre-built security policies that provide out-of-the-box protection for common applications such as Microsoft Outlook Web Access, Lotus Domino Mail Server, Oracle E-Business Financials, and Microsoft SharePoint. The validated policies also serve as a starting point for more advanced policy creation. This allows even novice users to rapidly deploy policies and immediately secure applications with little-to-zero configuration time needed.

Unified learning and dynamic policy building

At the heart of Advanced WAF is the unified learning and dynamic policy builder engine, which automates policy creation and tuning for increased operational efficiency and scalability. The policy builder engine automatically builds security policies around security violations, advanced statistics, and heuristics over time. It also understands expected behavior to affect more accurate traffic filtering.

By examining hundreds or thousands of requests and responses, the policy builder engine populates the security policy with legitimate elements more precisely than other WAFs. Dynamically generated policies are initially put into staging, then automatically moved from staging and enforced as they meet the rule thresholds for stabilization.
The policy builder engine supports automatic policy adaptation and learning following the occurrence of violations or as new parameters are observed. Policy maintenance is simplified by a GUI with a single-page view of all learning suggestions. One-click actions allow you to browse, search, accept, and ignore potential suggestions for policy adjustments, hardening policies with ease.

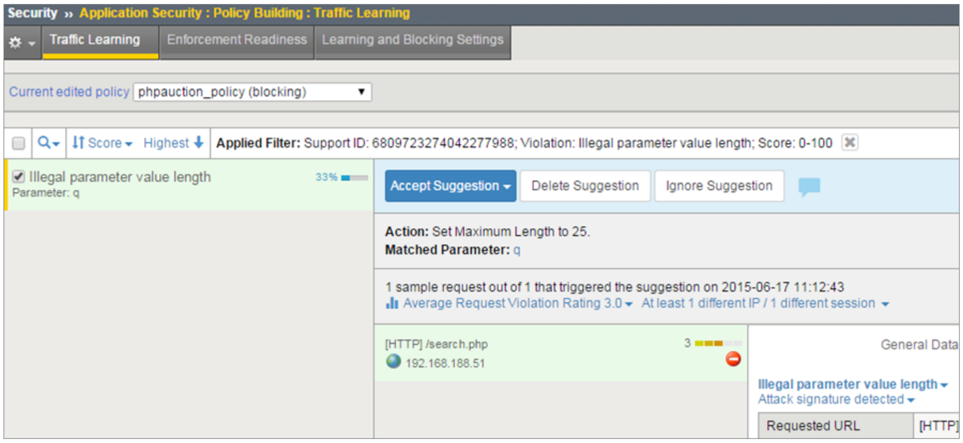

Current edited policy: phpauction_policy (blocking)

Applied Filter: Support ID: 6809723274042277988; Violation: Illegal parameter value length; Score: 0-100

☑ Illegal parameter value length 33%
Parameter: q

Action: Set Maximum Length to 25.
Matched Parameter: q

1 sample request out of 1 that triggered the suggestion on 2015-06-17 11:12:43
Average Request Violation Rating 3.0 At least 1 different IP / 1 different session

[HTTP] /search.php 3
192.168.188.51

Illegal parameter value length
Attack signature detected

Requested URL	[HTTP]

The enhanced learning GUI offers a single-page view of all learning suggestions.

Centralized management and monitoring

When you are deploying multiple Advanced WAF devices, F5 BIG-IQ® Centralized Management centralizes administration across your entire F5 infrastructure. Administrators get a consolidated view of all F5 devices, which helps to manage better relationships between devices, reduce IT overhead, and minimize configuration errors.

Advanced WAF provides an open API that supports easy integration to cloud/aaS virtual platforms and third-party policy management solutions. Engineers can fully configure and manage Advanced WAF policies from a programmatic interface that supports all policy management tasks, including login configuration, learning, semi-automatic tuning, utilization queries, and health monitoring. The Advanced WAF REST API exposes the entire range of Advanced WAF policy entities to support open models of WAF as a Service.

Leverage Rich, Actionable Reporting

F5 Advanced WAF provides powerful reporting capabilities that allow you to easily analyze incoming requests, track trends in violations through event correlation, generate security reports, evaluate possible attacks, and make informed security decisions. For security experts or generalists, Advanced WAF provides clear, discernable information with comprehensive visibility into attacks and changes in the threat landscape.

The Advanced WAF overview screen displays active security policies, security events and attacks, anomaly statistics, and networking and traffic statistics. Information can be saved or sent as an email attachment. Monitoring capabilities show how the application is being accessed and how it is behaving. The unique REST API supports easy integrations with higher-level SIEM or management services. Advanced WAF also offers predefined and customizable dashboards, charts, reports, and stats—highlighting DoS and brute force attacks, web scraping and IP enforcement, session tracking status, and more.

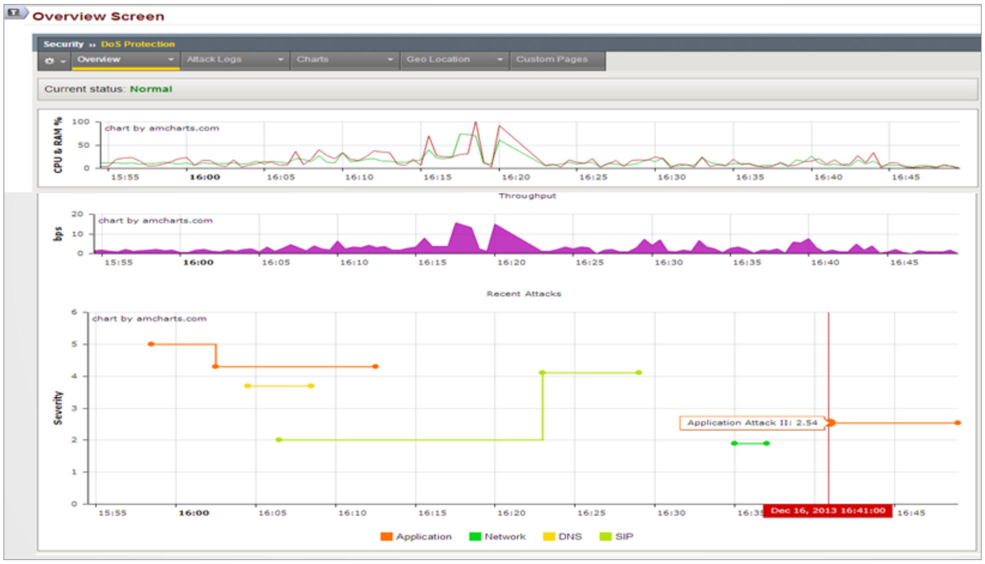

Current status: Normal

The security overview screen provides an easy view of what is happening on your system.

In-depth forensic analysis and database security

For deeper threat analysis, Advanced WAF integrates with high-speed indexing and search solutions like Splunk. These solutions offer deeper visibility into attack and traffic trends, long-term data aggregation, and identification of unanticipated threats before exposure occurs. Advanced WAF also supports database reporting for a real-time view into database activity and SQL statements generated by front-end users. Indexing and search solutions combine with Advanced WAF to provide richer forensic information for increased security effectiveness when mitigating threats.

Maintain compliance with industry and regulatory mandates

Advanced WAF makes it easy for organizations to understand and maintain regulatory compliance. Built-in security protection, logging and reporting, and remote auditing help organizations comply with industry security standards (including PCI DSS, HIPAA, BASEL II, FFIEC, SOX)—cost-effectively and without multiple appliances, application changes, or rewrites. With PCI reporting, Advanced WAF lists required security measures, determines if compliance is being met, and details necessary steps to becoming compliant.

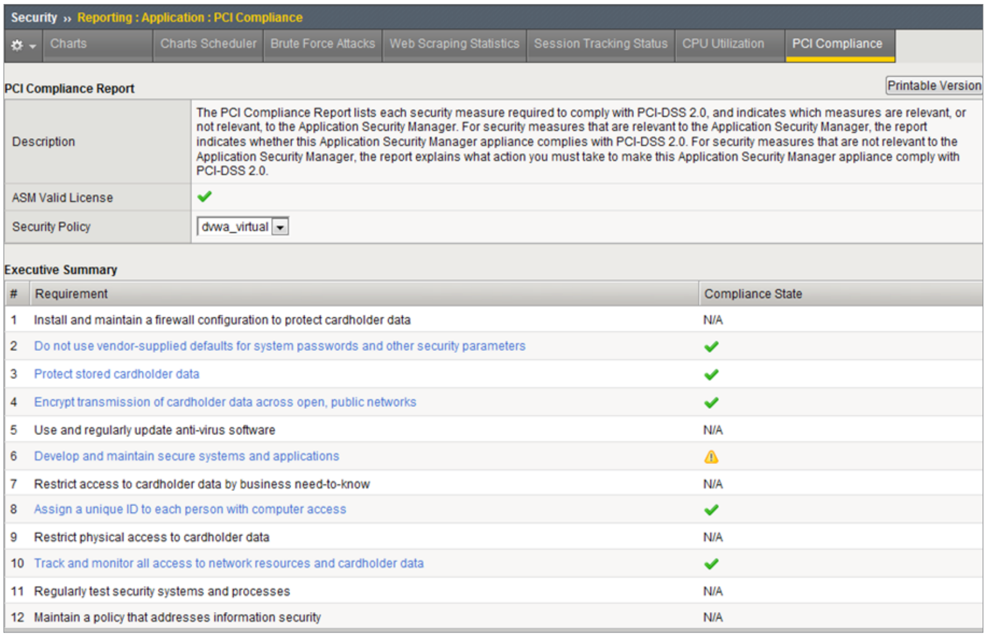

PCI Compliance Report

Description	The PCI Compliance Report lists each security measure required to comply with PCI-DSS 2.0, and indicates which measures are relevant, or not relevant, to the Application Security Manager. For security measures that are relevant to the Application Security Manager, the report indicates whether this Application Security Manager appliance complies with PCI-DSS 2.0. For security measures that are not relevant to the Application Security Manager, the report explains what action you must take to make this Application Security Manager appliance comply with PCI-DSS 2.0.
ASM Valid License	✓
Security Policy	dvwa_virtual

Executive Summary

#	Requirement	Compliance State
1	Install and maintain a firewall configuration to protect cardholder data	N/A
2	Do not use vendor-supplied defaults for system passwords and other security parameters	✓
3	Protect stored cardholder data	✓
4	Encrypt transmission of cardholder data across open, public networks	✓
5	Use and regularly update anti-virus software	N/A
6	Develop and maintain secure systems and applications	⚠
7	Restrict access to cardholder data by business need-to-know	N/A
8	Assign a unique ID to each person with computer access	✓
9	Restrict physical access to cardholder data	N/A
10	Track and monitor all access to network resources and cardholder data	✓
11	Regularly test security systems and processes	N/A
12	Maintain a policy that addresses information security	N/A

Maintain compliance with industry and regulatory mandates.

Meet Complex Deployment Requirements

The explosion of the Internet of Things (IoT) has caused a tremendous impact on organizations. The number of web-facing applications that must be managed and secured has jumped dramatically. In addition, the increasing focus toward hybrid application deployment means that business apps now reside in multiple settings—data center, private cloud, and public cloud. As a result of these changes, new requirements are necessary for securing apps and transitioning WAF services from the data center to the cloud.

Hybrid WAF deployment models

F5 Advanced WAF offers flexible options that allow administrators to easily deploy firewall services close to the application. Administrators can also transition hardened security policies from data center appliances to Advanced WAF Virtual Edition (VE) in virtual and private cloud environments. Advanced WAF offers the same quality of protection and scalability with an appliance and software edition. Policies and iRules can seamlessly move between hardware devices and virtual appliances without manual updates.

F5's WAF technology supports application security in any environment, whether deployed on F5 hardware, as a virtual edition, or as a wholly managed cloud-based service.

F5 Silverline Web Application Firewall is built on F5 Advanced WAF, but is provided via the Silverline cloud-based application services platform and wholly deployed, set up, and managed by the highly specialized experts in the F5 Security Operations Center (SOC). With 24x7x365 expert support to protect web applications and data (and enable compliance with industry security standards), the Silverline Web Application Firewall service provides application protection without the need for capital investment and security expertise.

Running multiple instances of Advanced WAF

Advanced WAF uses F5 ScaleN® with F5 Virtual Clustered Multiprocessing™ (vCMP) to provide the most cost-effective application security implementation for managing large-scale deployments—whether you are a managed service provider offering WAFs as a service or simply managing a large number of Advanced WAF devices.

With Advanced WAF and vCMP-enabled systems, administrators can easily consolidate multiple firewalls onto a single device and allocate Advanced WAF resources in a more flexible and isolated manner for different customers, groups, applications, and services. vCMP enables you to run multiple instances of Advanced WAF on a single F5 platform with high-density firewall isolation through a combination of hardware and software. Guest firewalls can be clustered for easier administration and maintenance, and to ensure consistency throughout the firewall infrastructure. vCMP allows you to consolidate and better manage your security infrastructure, ensuring efficiencies and meeting service-level agreements (SLAs) with a dynamic, flexible WAF service infrastructure.

F5 Security Services

IT managers need a consolidated network and web application firewall solution to defend against multi-layered attacks, such as network and layer 7 events. F5 Advanced WAF, together with F5 Web Fraud Protection, F5 BIG-IP Advanced Firewall Manager™ (AFM), and F5 BIG-IP DNS, covers the threat spectrum—mitigating L3–L7 attacks, providing client-side fraud protection, and safeguarding the DNS infrastructure. When used with F5 Access Manager® (AM), Advanced WAF provides context-aware, policy-based access with simplified authentication, authorization, and accounting (AAA) management for web applications. As a component of F5's comprehensive security services, Advanced WAF benefits from other F5 modules to enable data center security, extensive application protection, and access management capabilities.

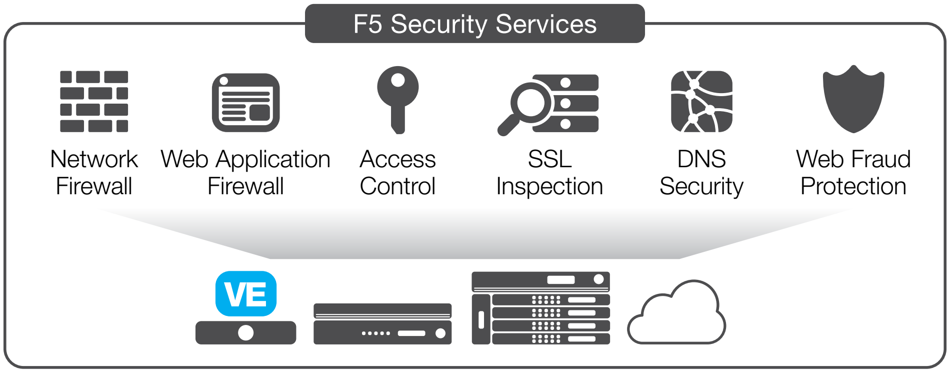

Advanced WAF, together with other BIG-IP modules, consolidates application protection and access management onto a single high-performing security platform.

F5 Advanced WAF Features and Specifications

Web Application Firewall

Deployment

Rapid deployment wizard with self-help hints	Yes
Unified learning and policy builder	Yes—with manual and automated policy building
Policy staging	Yes
Route domain support	Yes
VE, appliance, or managed service	Yes—managed services require Silverline License

WAF Security

Application layer encryption	Yes
Brute Force mitigation	Yes
Credential Stuffing protection	Yes
Behavioral denial-of-service (DoS) protection	Yes—protection for all applications
L7 DoS and DDoS detection including: HASH DoS, Slowloris, floods, Keep-Dead, XML bomb	Yes
Web scraping prevention	Yes
OWASP Top 10 prevention	Yes
Automated attack defense and bot detection	Yes
Advanced protections against threats including: Web injections, data leakage, session hijacking, HPP attacks, buffer overflows, shellshock	Yes
Mobile bot protection	Yes—with the F5 Anti-Bot Mobile SDK
Geolocation blocking	Yes
IP intelligence reputation services	Yes—with F5 IP Intelligence Services
SSL termination with re-encryption	Yes
Security incident and violation correlation	Yes
Client-side certification support	Yes
Client authentication	LDAP, RADIUS; more methods available with F5 Access Manager
Database security	Yes—with Oracle Database firewall
Response checking	Yes
Violation risk scoring	Yes
Web service encryption and decryption	Yes—and with signature validation
Device-ID detection and finger printing	Yes
Live signature updates	Yes
WebSocket traffic filtering	Yes
IP shunning (layer 3 denylisting in HW)	Requires F5 BIG-IP AFM license

Reporting and Analytics

Customizable charts and reports	Yes
Security overview report	Yes—drill down capabilities to granular details
Combined network and application attack report	Yes—with combined F5 BIG-IP AFM and F5 WAF deployment
WAF health monitoring	Yes
Compliance support PCI-DSS, HIPAA, SOX, Basel II	PCI-DSS, HIPAA, SOX, Basel II
Central management and reporting with role-based access control	Yes—requires F5 BIG-IQ Centralized Management
Automatic policy sync between WAF devices	Yes

Other

iRules and fast cache integration	Yes
SNMP reporting	Yes
REST API	Yes
ICAP support	Yes
DAST integration	Yes—WhiteHat, QualysGuard, and IBM
Fraud protection	Yes—requires F5 Fraud Protection Service license
SSL acceleration	Yes—core to the BIG-IP platform

BIG-IP Platform and TMOS support

Multi-tenancy	Yes—with F5 vCMP
High availability	Yes—active-passive or active-active
64-bit OS support	Yes
Application acceleration	Yes—requires F5 BIG-IP LTM
TCP optimization	Yes
Advanced rate shaping and QoS	Yes
F5 IPv6 Gateway™	Yes
IP port filtering	Yes
VLAN support	Yes
Secure SSL certificates from access	Yes
Integrates with BIG-IP AFM and F5 AM for complete data center security with identity and access management	Yes

F5 Advanced WAF

F5 Advanced WAF is available as a standalone solution or as an add-on module for BIG-IP Local Traffic Manager (LTM) on any F5 platform, and on BIG-IP LTM Virtual Edition (VE). F5 Access Manager (AM) is available as an add-on module to the Advanced WAF standalone appliance. F5 AM Lite (with 10 free user licenses) is included with any Advanced WAF standalone purchase. For detailed physical specifications, please refer to the BIG-IP System Hardware Data Sheet.

BIG-IP Platforms

Only F5's next-generation, cloud-ready application delivery controller (ADC) platform provides DevOps-like agility with the scale, security depth, and investment protection needed for both established and emerging apps. The new F5 BIG-IP iSeries appliances deliver quick and easy programmability, ecosystem-friendly orchestration, and record-breaking, software-defined hardware performance. As a result, customers can accelerate private clouds and secure critical data at scale while lowering TCO and future-proofing their application infrastructures. F5 solutions can be rapidly deployed via integrations with open source configuration management tools and orchestration systems.

In addition to iSeries, F5 offers the VIPRION® modular chassis and blade systems designed specifically for performance and for true on-demand linear scalability without business disruption. VIPRION systems leverage F5's ScaleN clustering technology so you can add blades without reconfiguring or rebooting.

Virtual editions of F5 software run on commodity servers and support the range of hypervisors and performance requirements. These virtual editions provide agility, mobility, and fast deployment of app services in software-defined data centers and cloud environments. See the F5 System Hardware, VIPRION, and Virtual Edition data sheets for more details. For information about specific module support for each platform, see the latest release notes on AskF5. For the full list of supported hypervisors, refer to the VE Supported Hypervisors Matrix.

F5 platforms can be managed via a single pane of glass with BIG-IQ Centralized Management.

BIG-IP iSeries Appliance

VIPRION Chassis

BIG-IP Virtual Editions

Virtual Editions

F5 Advanced WAF Virtual Edition (VE) can help you meet the needs of your virtualized environment by scaling to 20 cores/vCPUs.

F5 Advanced WAF VE

Hypervisors Supported:	· VMware vSphere Hypervisor 4.0, 4.1, 5.0, and 5.1 and vCloud Director 1.5 · Citrix XenServer 5.6 and 6.0 · Microsoft Hyper-V for Windows Server 2008 R2 and 2012 · KVM – Linux Kernel 2.6.32 (RHEL 6.2/6.3, CentOS 6.2/6.3)

Advanced WAF VE is also available as an Amazon Machine Image for use within Amazon Web Services.

amazon web services | Partner Network

F5 Global Services

F5 Global Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services can help ensure your applications are always secure, fast, and reliable. For more information about F5 Global Services, contact consulting@f5.com or visit f5.com/support.

More Information

To learn more about F5 Advanced WAF, visit f5.com to find these and other resources.

Additional Resources

F5 Advanced WAF Overview

Advanced Application Threats Require an Advanced WAF

F5 Labs 2018 Application Protection Report

eBooks

Bots Mean Business

Credential Stuffing | A Security Epidemic

OWASP Top 10 and Beyond

Report

Gartner Web Application Firewall Magic Quadrant, 2018

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

Americas
info@f5.com

Asia-Pacific
apacinfo@f5.com

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

Japan
f5j-info@f5.com

DATA SHEET

BIG-IP Access Policy Manager

WHAT'S INSIDE

Simple, Secure, and Seamless Access to Any Application, Anywhere

Applications are gateways to your critical and sensitive data. Simple, secure access to your applications is paramount, but application access today is extremely complex. Apps can be hosted anywhere—in the public cloud, in a private cloud, on-premises, or in a data center. Ensuring users have secure, authenticated access anytime, anywhere, to only the applications they are authorized to access is now a significant challenge. There are different application access methods to deal with these complexities. There are various sources for authorized user identity, as well as dealing with applications that require modern or traditional authentication and authorization methods, single sign-on (SSO), federation, and more, in addition to the user access experience to support and consider.

With digital transformation touching every part of an enterprise today, native cloud and Software as a Service (SaaS) applications are now the enterprise application standard. Many organizations, though, find that they are unable or unwilling to migrate all of their applications to the cloud. There may be mission-critical classic or custom applications that should not or cannot support being migrated to the public cloud or be easily replaced by a SaaS application. Applications are being hosted in a variety of locations, with differing and many times disparate authentication and authorization methods that are unable to communicate with each other and can't work seamlessly across existing SSO or federated identity, that are unable to support the newest identity means like Identity as a Service (IDaaS), and are not equipped to support multi-factor authentication (MFA).

F5® BIG-IP® Access Policy Manager® (APM) is a secure, flexible, high-performance access management proxy solution managing global access to your network, the cloud, applications, and application programming interfaces (APIs). Through a single management interface, BIG-IP APM consolidates remote, mobile, network, virtual, and web access. With BIG-IP APM, you can create, enforce, and centralize simple, dynamic, intelligent application access policies for all of your apps, regardless of where or how they are hosted.

ICSA labs
CERTIFIED
SSL VPN
FIREWALL - CORPORATE

KEY BENEFITS

Simplify access to all apps
Bridge secure access to on-premises and cloud apps with a single login via SSO. It even works for applications unable to support modern authentication such as Security Assertion Markup Language (SAML), or OAuth and OpenID Connect (OIDC).

Zero trust application access
Identity Aware Proxy (IAP) delivers a zero trust model validation for application access based on identity-awareness and granular context, securing every app access request without the need of a VPN.

Secure web access
Control access to web-based applications and web content centralizing authentication, authorization, and endpoint inspection via web app proxy.

Centralize and manage access control
Consolidate management of remote, mobile, network, virtual, and web access in a single control interface with adaptive identity federation, SSO, and MFA via dynamically enforced, context-based and identity-aware policies.

Streamlined authentication and authorization
Adaptive identity federation, SSO, and MFA employing SAML, OAuth, and OIDC for a seamless and secure user experience across all apps.

BRIDGING SECURE APPLICATION ACCESS

Modern authentication and authorization protocols—including Secure Assertion Markup Language (SAML), and OAuth with OpenID Connect (OIDC)—reduce user dependency on passwords, increase security, and improve user experience and productivity. However, not all applications support modern authentication and authorization protocols. Many applications, such as classic applications or custom-built applications, support classic authentication and authorization methods, such as Kerberos, NT LAN Manager (NTLM), RADIUS, header-based, and more. This further complicates application access and security. The need to support different, disparate protocols unable to share user authentication and authorization information inhibits the use of SSO and MFA. That in turn negatively impacts user experience and application security. It also makes it difficult to adapt modern corporate password policy of periodic password changes, and increases organizational costs as multiple access methods become necessary.

BIG-IP APM serves as a bridge between modern and classic authentication and authorization protocols and methods. For applications which are unable to support modern authentication and authorization protocols, like SAML and OAuth with OIDC, but which do support classic authentication methods, BIG-IP APM converts user credentials to the appropriate authentication standard supported by the application. BIG-IP APM ensures that users or organizations can use SSO to access any application anywhere—regardless of its location (on-premises, in a data center, in a private cloud, or in the public cloud as a native cloud or SaaS application), or whether or not it supports modern or classic authentication and authorization. This helps decrease the number of passwords users have to create, remember, and use, helping to stem the tide of credential-based attacks. It enables compliance with modern corporate policies of periodic password changes to combat stolen credentials. It also decreases the cost to organizations of having to purchase and maintain separate access solutions for applications hosted on-premises, in a data center, and in a private cloud, versus native cloud and SaaS apps.

BIG-IP APM supports identity federation and SSO options by supporting connections initiated by both SAML identity providers (IdP) and service providers (SP) leveraging SAML 2.0. F5 supports JWE token consumption, enabling BIG-IP APM to consume encrypted JWT tokens from SAML IdP vendors. This allows BIG-IP APM to maintain secrecy between the issuer or the token and the recipient. BIG-IP APM empowers administrators to centrally enable and disable user authorized access to any identity-enabled applications, regardless of where they are hosted, saving time and boosting administrative productivity.

Support for OAuth 2.0 open-standard for authorization enables BIG-IP APM to serve as a client, as an authorization delegate for SaaS applications, and can enhance protection for and authorization of APIs for web services.

BIG-IP APM supports Proof Key for Code Exchance (PKCE), a more secure authorization flow based on OAuth 2.0 that improves security for all OAuth clients (including mobile and public).

KEY BENEFITS (CONT.)

Defend your weakest links
Protect against data loss, malware, and rogue device access with comprehensive, continuous endpoint integrity and security checks.

Protect APIs
Enable secure authentication for REST APIs and integrate OpenAPI or "swagger" files to ensure appropriate authentication actions while saving time and cost.

Do it all at scale
Support all users easily, quickly, and cost-effectively with no performance trade-offs for security, even in the most demanding environments.

This helps prevent attacks like authorization code interception by enabling dynamic secret tokens. PKCE enables apps to make direct requests to the token provider, adding an extra layer of security for public apps.

SUPPORT FOR IDAAS

With support for SSO and Kerberos ticketing across multiple domains, BIG-IP APM enables additional types of authentication, such as U.S. Federal Government Common Access Cards (CAC) and the use of IDaaS—such as Azure Active Directory (Azure AD), Okta, and others—to access all applications regardless of location or modern authentication and authorization support. For instance, users can be automatically signed on to back-end applications and services that are part of a Kerberos realm. This provides a seamless authentication flow once a user has been authenticated through a supported user-authentication mechanism. BIG-IP APM also supports smart cards with credential providers, so users can connect their devices to their network before signing in.

SUPPORT FOR MFA

Through F5's extensive partner ecosystem, BIG-IP APM also integrates with most leading MFA solutions, including those from Cisco Duo, Okta, Azure AD, and others. By integrating with your existing MFA solution, BIG-IP APM enables adaptive authentication, allowing various forms of single-, two-, or multi-factor authentication to be employed based on user identity, context, and application access. In addition, to help you deploy MFA, BIG-IP APM includes one-time password (OTP) authentication via email or SMS.

After the user has logged into an application, an additional means of authentication may be required to ensure secure access to mission-critical or particularly sensitive applications and files. This is commonly referred to as step-up authentication. BIG-IP APM supports step-up authentication for single- and multi-factor authentication. Any session variable may be used to trigger step-up authentication, and you can use additional authentication capabilities or select from our partner offerings. In addition, any session variable may be part of access policy branching (such as URL branching) per request policy. Step-up authentication policies may be based on applications, secure portions of applications, sensitive web URIs, extending sessions, or any session variable.

Many authentication solutions use application coding, separate web server agents, or specialized proxies that present significant management, cost, and scalability issues. With AAA control, BIG-IP APM enables you to apply customized access policies across many applications and gain centralized visibility of your authorization environment. You can consolidate your AAA infrastructure, eliminate redundant tiers, and simplify management to reduce capital and operating expenses.

ZERO TRUST APPLICATION ACCESS

Many organizations—possibly including yours—are rapidly moving toward adoption of a zero trust security architecture. The pillars of a zero trust security architecture are identity and context.

A zero trust approach to security means adopting a mindset that attackers have already infiltrated your network and are lurking, waiting for an opportunity to launch an attack. It eliminates the idea of a trusted insider within a defined network perimeter, assuming, at best, a limited secure network perimeter. It encourages never trusting users, even if they've already been authenticated, authorized, and granted access to applications and resources. A zero trust security approach applies least privilege rights to user access, allowing users to access only those applications and resources they are authorized for, and restricting their access to a single application or resource at a time.

Identity- and context-awareness are also what define Identity Aware Proxy (IAP). IAP enables secure access to specific applications by leveraging a fine-grained approach to user authentication and authorization. IAP enables only per-request application access, which is very different than the broad network access approach of VPNs that apply session-based access, which is not a zero trust approach. With this approach, VPN becomes optional to access applications. IAP enables the creation and enforcement of granular application access policies based on contextual attributes, such as user identity, device integrity, and user location. IAP relies on application-level access controls, not network-layer rules. Configured policies reflect user and application intent and context. IAP requires a strong root of trusted identity to verify users, and to stringently enforce what they are authorized to access.

Identity Aware Proxy is key to both a zero trust security architecture and to F5 BIG-IP APM. BIG-IP APM and F5 Access Guard deliver Identity Aware Proxy using a zero trust validation model on every application access request. Providing authenticated and authorized users secure access to specific applications, it leverages F5 best-in-class access proxy. BIG-IP APM centralizes user identity and authorization. Authorization is based on the principles of least privileged access.

Through IAP, BIG-IP APM examines, terminates, or authorizes application access requests. Policies within BIG-IP APM can be created to:

- Verify user identity
- Check device type and posture
- Validate user authorization
- Confirm application integrity and sensitivity
- Confirm time and date accessibility

- Limit or halt access if the user’s location or their device posture is deemed incorrect, inappropriate, or insecure
- Request additional forms of authentication—including multi-factor authentication (MFA)—if the user’s location or the sensitive nature of the applications or its data warrant it
- And more

Data from user and entity behavior analytics (UEBA) and other API-driven risk engines can be integrated seamlessly adding another level of security and application access control.

BIG-IP APM checks user device security posture via F5 Access Guard, a browser extension that coordinates with BIG-IP APM. However, BIG-IP APM and F5 Access Guard go beyond simply checking device integrity at authentication to deliver continuous, ongoing device posture checks, ensuring that user devices not only meet but adhere to endpoint security policies throughout application access. If BIG-IP APM detects any change in device integrity, it can either limit or stop application access, halting potential attacks before they can even be launched.

A guided configuration workflow allows organizations to host web applications protected by Identity Aware Proxy on a webtop, providing users a single catalog of their applications. It offers a seamless user experience, as users can access applications, regardless of where they are hosted. It also simplifies the administrative workflow, enabling administrators to easily pick, choose, and modify the applications accessible by a specific user group.

BIG-IP APM, through IAP, also simplifies application access for remote or home-based workers and better enables and secures application accessibility, and optionally eliminates the need for VPNs.

ROBUST ENDPOINT SECURITY

BIG-IP APM inspects and assesses users’ endpoint devices before authentication and throughout the user’s application access with F5 Access Guard. F5 Access Guard examines device security posture and determines if the device is part of the corporate domain. Based on the results, BIG-IP APM will apply dynamic access control lists (ACLs) to deploy context-based security. BIG-IP APM and F5 Access Guard include preconfigured, integrated endpoint inspection checks, including checks for OS type, antivirus software, firewall, file, process, registry value validation and comparison (Windows only), as well as device MAC address, CPU ID, and HDD ID. For mobile devices running iOS or Android, BIG-IP APM’s endpoint inspection checks the mobile device UDID and jailbroken or rooted status.

RISK-BASED ACCESS USING THIRD-PARTY RISK ENGINES (HTTP CONNECTOR)

Many organizations have deployed third-party user and entity behavior analytics (UEBA) or risk engines. The ability to leverage an existing UEBA or risk engine to infuse real-time analytics and risk data within their access control policies can help those organizations ensure that access to networks, clouds, applications, and even APIs, are regulated based on a risk profile. It is also important to address risk-based access to networks, clouds, apps, and APIs that is triggered by a variety of relevant variables.

Through its HTTP Connector, BIG-IP APM integrates with third-party UEBA and risk engines, leveraging their risk assessment via REST APIs as part of its policy-based access controls. This enables risk-based access to networks, clouds, apps, and APIs, further enhancing BIG-IP APM's zero trust IAP solution. BIG-IP APM's HTTP Connector leverages user group, domain, and network-based triggers to increase the enforceability of risk-based access. Risk-based access enhances security, providing greater visibility and analytics to determine whether to grant or deny access to your networks, cloud, applications, and APIs.

INTELLIGENT INTEGRATION WITH IDENTITY AND ACCESS MANAGEMENT (IAM)

F5 partners with leading on-premises and cloud-based identity and access management (IAM) vendors, such as Microsoft, Okta, and Ping Identity. This integration enables local and remote user SSO via SAML, OAuth or FIDO2 (U2F) to applications based on premises or in a data center. For organizations that do not wish to replicate their user credential store in the cloud with IDaaS or cloud-based IAM offerings, working with its partners, F5 and BIG-IP APM work to help these organizations maintain control of on-premises user credentials. This is accomplished by creating a bridge between the IAM vendor's offering and the local authentication services. This bridge, or identity provider chain, leverages SAML to federate user identity.

UNIFYING ACCESS FROM ANY DEVICE

BIG-IP APM is positioned between your applications and your users, providing a strategic application access control point. It protects your public-facing applications by providing granular policy for identity- and context-aware user access, while consolidating your access infrastructure. It secures remote and mobile access to applications, networks, and clouds via SSL VPN or zero trust application access. BIG-IP APM converges and consolidates all access—network, cloud, application, and API—within a single management interface. It also enables and simplifies the creation of easy to manage dynamic access policies.

BIG-IP APM includes a dynamic web-based application portal or webtop. The BIG-IP APM webtop shows only the applications authorized for and available to a user based on their identity and context—regardless of where the applications are hosted—on-premises, in a data center, in a private cloud, in a public cloud, or offered as a service.

BIG-IP APM enables Datagram Transport Layer Security (DTLS) mode, supporting DTLS 2.0 for remote connections that secure and tunnel delay-sensitive applications. It supports IPsec encryption for traffic between branch offices or data centers. Per-app VPN via an application tunnel through BIG-IP APM enables access to a specific application without the security risk of opening a full network access tunnel.

F5 BIG-APM enables secure access to applications, networks, and clouds via the BIG-IP Edge Client and F5 Access. The BIG-IP Edge Client is available for Apple MacOS, Microsoft Windows, Linux platforms, Chromebooks, and includes support for Windows on ARM64 devices. F5 Access is an optional mobile client for ensuring secure access from mobile devices supporting Apple iOS and Google Android, and is available for download from the Apple App Store or Google Play.

BIG-IP Edge Client and F5 Access integrate with leading mobile device management (MDM) and enterprise mobility management (EMM) solutions—including VMware Horizon ONE (AirWatch), Microsoft Intune, and IBM MaaS360—to perform device security and integrity checks and to deliver per-app VPN access without user intervention. Context-aware policies are assigned based on a device's security state. These policies enable, modify, or disable application, network, and cloud access from the device. Hardware attributes may be mapped to a user's role to enable additional access control decision points. A browser cache cleaner automatically removes any sensitive data at the end of a user's session.

Biometrics, such as fingerprint access, are supported to open and access the F5 Edge Client. This simplifies access, since a user will no longer need to create, remember, and input a username/password credential to access the Edge Client. It also makes accessing the Edge Client more secure, as users reuse passwords or create simple username/password pairs, making them easier for attackers to hack.

BIG-IP APM also supports server authentication via Client Certificate Constrained Delegation (C3D). By employing C3D, BIG-IP APM addresses certificate-based authentication, limiting the need for and use of credentials. With C3D, organizations can implement stronger encryption protocols and the latest key exchanges, as well as employ client certificate authentication, enable end-to-end encryption in reverse proxy environments, leverage Perfect Forward Secrecy (PFS), and validate client certificates using Online Certificate Status Protocol (OCSP).

SEAMLESS ACCESS TO ALL APPLICATIONS

As organizations focus on reducing user friction and increasing agility, their need to provide seamless access to all applications becomes a priority. BIG-IP APM enables organizations to reduce friction for users to remote access (SSL VPN). It also reduces friction for web applications, as well. BIG-IP APM supports SSO across both remote access and web applications with a single login for either Apple Macs or Microsoft Windows devices (via Windows Hello For Business). Organizations are able to support the user login via U2F tokens (such as Yubico keys) or password-less FIDO2 via the F5 Edge Client to reduce user friction and increase application access security.

STREAMLINE VIRTUAL APPLICATION ACCESS

Virtual desktop and application deployments must scale to meet the needs of thousands of users and hundreds of connections per second. BIG-IP APM serves as a gateway for virtual application environments. It includes native support for Microsoft Remote Desktop Protocol (RDP), native secure web proxy support for Citrix XenApp and XenDesktop, and security proxy access for VMware Horizon. Administrators can control the delivery and security components of enterprise virtualization solutions via BIG-IP APM's unified access, security, and policy management. These scalable, high-performance capabilities simplify user access and control in hosted virtual desktop environments. BIG-IP APM delivers simple, broad virtual application and desktop support.

BIG-IP APM supports two-factor authentication via RSA SecureID and RADIUS through the native client for VMware End User Computing (EUC) deployments. BIG-IP APM supports Citrix Virtual Apps and Desktops and Citrix StoreFront. BIG-IP APM, when integrated with the Microsoft RDP protocol, enables the remote desktop access needed to install client-side components or run Java. It allows Microsoft RDP to be available for use on new platforms, such as Apple iOS and Google Android devices. It also enables native RDP clients on non-Windows platforms such as Mac OS and Linux, where previously only a Java-based client was supported. BIG-IP APM's Microsoft RDP support works with any Microsoft, Apple, or Google web browser, or RDP app installed.

PROTECTING APIs

APIs are the connective tissue in modern application architectures. Attackers are leveraging APIs to launch attacks, because they are ripe for exploitation: Many organizations expose APIs to the public and their supply chain partners or they inadvertently leave them unprotected. While attackers are exploiting APIs to launch attacks, organizations can ensure API security via authentication, especially if it's adaptable and protected by consistent, flexible authentication and authorization policies. BIG-IP APM enables secure authentication for REST APIs. It also ensures appropriate authorization actions are taken. BIG-IP APM integrates existing OpenAPI, or "swagger" files, saving time, human resources, and cost when developing API protection policies, while ensuring accurate API protection policies are in place.

SECURING CREDENTIALS

User credentials are like the keys to the kingdom: All an attacker has to do is steal one set of user credentials, and they can enjoy unfettered access to your organization's network, clouds, and apps.

BIG-IP APM's credential protection, as part of an optional license of BIG-IP DataSafe™, secures credentials from theft and reuse. It protects against Man-in-the-Browser (MitB) attacks with real-time, adaptable login encryption, and encrypts user credentials entered into its webtop. BIG-IP APM, in conjunction with BIG-IP DataSafe, renders the credentials unreadable and unusable, even in the unlikely event an attacker successfully steals them. BIG-IP APM also ensures login security for all applications associated via federation.

F5 DISTRIBUTED BOT DEFENSE–BUILT INTO THE BIG-IP PLATFORM

Bots cause significant financial pain through scraping that slows performance, scalping and inventory hoarding that frustrate loyal customers, enumerating gift card codes to steal balances, creating fake accounts to commit fraud, and credential stuffing—the testing of stolen credentials—that leads to account takeovers.

Today's advanced persistent bots are more sophisticated than ever, evading many standard bot defenses available within WAFs. Criminals will retool bots to bypass defenses within hours, utilize millions of valid IP addresses, rapidly solve CAPTCHAs, mimic human behavior, and introduce subtle randomness.

To stay ahead of attackers, F5® Distributed Cloud Bot Defense uses rich, client-side signal collection, industry-leading code obfuscation, aggregate telemetry collection, and AI for unparalleled long-term efficacy and near-zero false positives while maintaining access for good bots. And because F5 defends the most targeted sites on the web—including those of the world's largest banks, retailers, and airlines, F5 is ready when these attacks target your organization.

Deploy Distributed Cloud Bot Defense directly from your BIG-IP or through a connector that's right for your application, with support services tailored to your needs, from self-service to managed service.

VISUAL POLICY EDITOR (VPE)

Through its advanced graphical Visual Policy Editor (VPE), BIG-IP APM makes designing and managing granular access control policies on an individual or group basis fast and simple. With VPE, you can efficiently create and edit entire dynamic access policies in just a few clicks. BIG-IP APM's VPE can define rules per URL path. By centralizing and simplifying the management of contextual policies, you can efficiently direct fine-grained user access to applications, networks, and clouds.

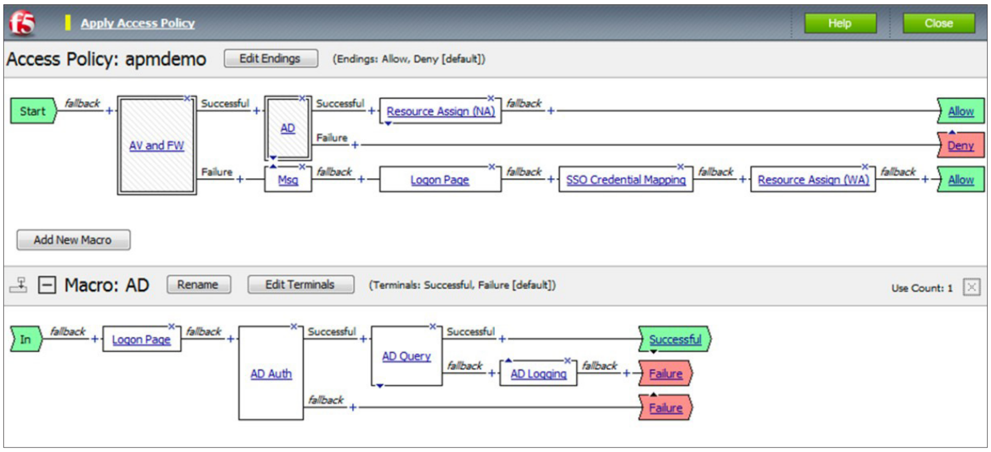

Figure 1: The BIG-IP APM advanced VPE makes it fast and easy to create, modify, and manage granular application-, user-, network/cloud-, and vulnerability context-based access policies.

BIG-IP APM lets you design access policies for authentication and authorization, as well as endpoint security checks, enforcing user compliance with corporate policies and industry regulations. One access profile may be defined for all connections coming from any device, or you can create multiple access profiles for different access methods from various devices.

BIG-IP APM enforces access authentication using ACLs and authorizes users with dynamically applied layer 4 and layer 7 ACLs on a session. Both L4 and L7 ACLs are supported based on endpoint posture as a policy enforcement point. Individual and group access to approved applications and networks is allowed by BIG-IP APM using dynamic, per-session L7 (HTTP) ACLs. The VPE in BIG-IP APM can be used to quickly and easily create, modify, and manage ACLs.

ACCESS GUIDED CONFIGURATION (AGC)

BIG-IP APM includes an Access Guided Configuration (AGC) capability that simplifies the deployment and management of application access. The AGC guides your administrator through a step-by-step process of setting up and deploying BIG-IP APM, saving you and your administrator deployment time and cost. BIG-IP APM's AGC also allows your administrator to quickly, simply onboard and operationally manage classic mission-critical applications, such as SAP ERP and Oracle PeopleSoft, to Azure AD. This simplified guided access eliminates numerous steps previously required in Azure AD to bridge the access gap between applications supporting modern authentication, and apps that support classic authentication methods, greatly reducing administrative overhead involved in modernizing those applications.

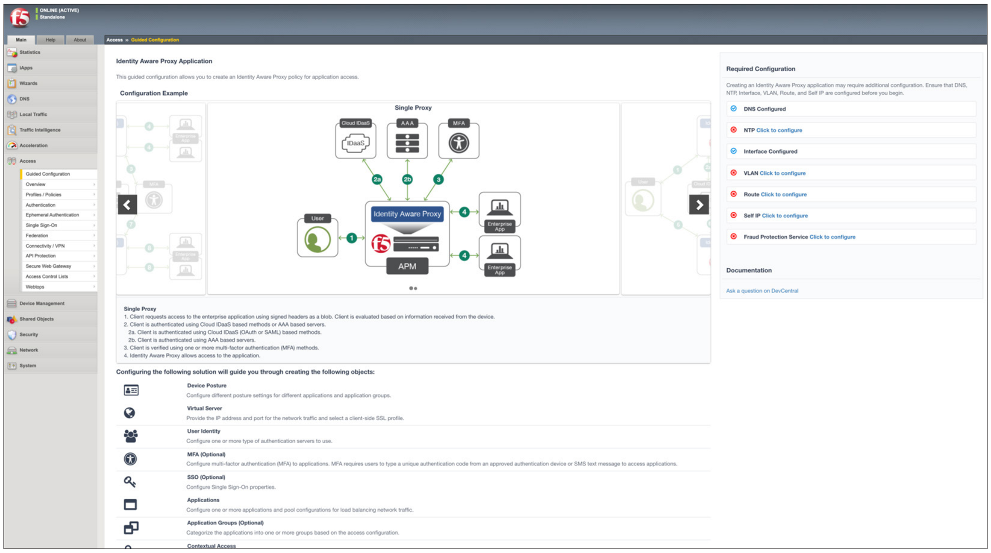

Identity Aware Proxy Application

This guided configuration allows you to create an Identity Aware Proxy policy for application access.

Configuration Example

Single Proxy
1. Client requests access to the enterprise application using signed headers as a blob. Client is evaluated based on information received from the device.
2. Client is authenticated using Cloud IDaaS based methods or AAA based servers.
 2a. Client is authenticated using Cloud IDaaS (OAuth or SAML) based methods.
 2b. Client is authenticated using AAA based servers.
3. Client is verified using one or more multi-factor authentication (MFA) methods.
4. Identity Aware Proxy allows access to the application.

Configuring the following solution will guide you through creating the following objects:

Device Posture
Configure different posture settings for different applications and application groups.

Virtual Server
Provide the IP address and port for the network traffic and select a client-side SSL profile.

User Identity
Configure one or more type of authentication servers to use.

MFA (Optional)
Configure multi-factor authentication (MFA) to applications. MFA requires users to type a unique authentication code from an approved authentication device or SMS text message to access applications.

SSO (Optional)
Configure Single Sign-On properties.

Applications
Configure one or more applications and pool configurations for load balancing network traffic.

Application Groups (Optional)
Categorize the applications into one or more groups based on the access configuration.

Contextual Access

Required Configuration

Creating an Identity Aware Proxy application may require additional configuration. Ensure that DNS, NTP, Interface, VLAN, Route, and Self IP are configured before you begin.

- ✓ DNS Configured
- ⊗ NTP Click to configure
- ✓ Interface Configured
- ⊗ VLAN Click to configure
- ⊗ Route Click to configure
- ⊗ Self IP Click to configure
- ⊗ Fraud Protection Service Click to configure

Documentation

Ask a question on DevCentral

Figure 2: BIG-IP APM's Access Guided Configuration saves deployment time and cost.

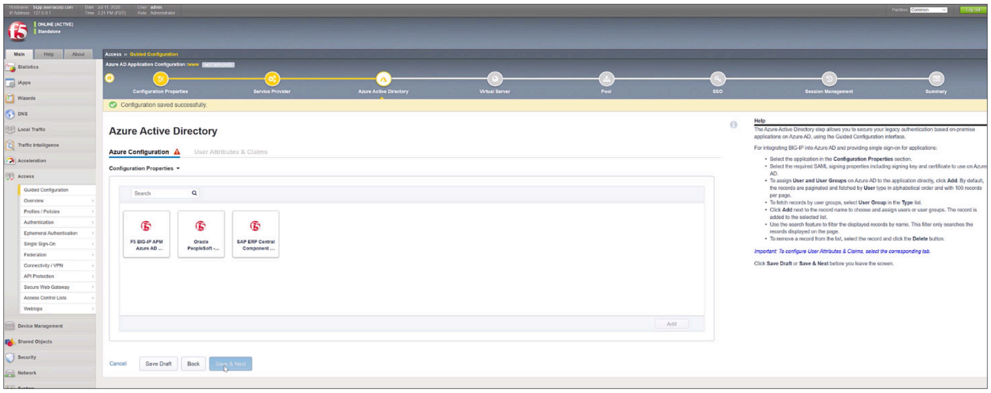

Configuration saved successfully.

Azure Active Directory

Help

The Azure Active Directory step allows you to secure your legacy authentication based on-premise applications on Azure AD, using the Guided Configuration interface.

For integrating BIG-IP into Azure AD and providing single sign-on for applications:

- Select the application in the **Configuration Properties** section.
- Select the required SAML signing properties including signing key and certificate to use on Azure AD.
- To assign **User and User Groups** on Azure AD to the application directly, click **Add**. By default, the records are paginated and fetched by **User** type in alphabetical order and with 100 records per page.
- To fetch records by user groups, select **User Group** in the **Type** list.
- Click **Add** next to the record name to choose and assign users or user groups. The record is added to the selected list.
- Use the search feature to filter the displayed records by name. This filter only searches the records displayed on the page.
- To remove a record from the list, select the record and click the **Delete** button.

Important: To configure User Attributes & Claims, select the corresponding tab.

Click **Save Draft** or **Save & Next** before you leave the screen.

Figure 3: F5 BIG-IP APM's Access Guided Configuration enables quick, simple onboarding and management of custom applications and classic applications, such as SAP ERP and Oracle PeopleSoft, with Azure AD.

CENTRALIZE ACCESS POLICY MANAGEMENT

If you have multiple BIG-IP APM deployments, F5 BIG-IQ Centralized Management® will help you efficiently manage them. It can manage policies for up to 100 BIG-IP APM instances, enabling you to import, compare, edit, and update granular access policies across multiple user devices.

With BIG-IQ Centralized Management and BIG-IP APM, you can import configurations from a master "source" BIG-IP APM instance, simplifying access policy distribution. You may also edit device- or location-specific objects directly on BIG-IQ Centralized Management and have them propagate throughout your BIG-IP APM deployment. You can easily view the differences between current and proposed access configurations.

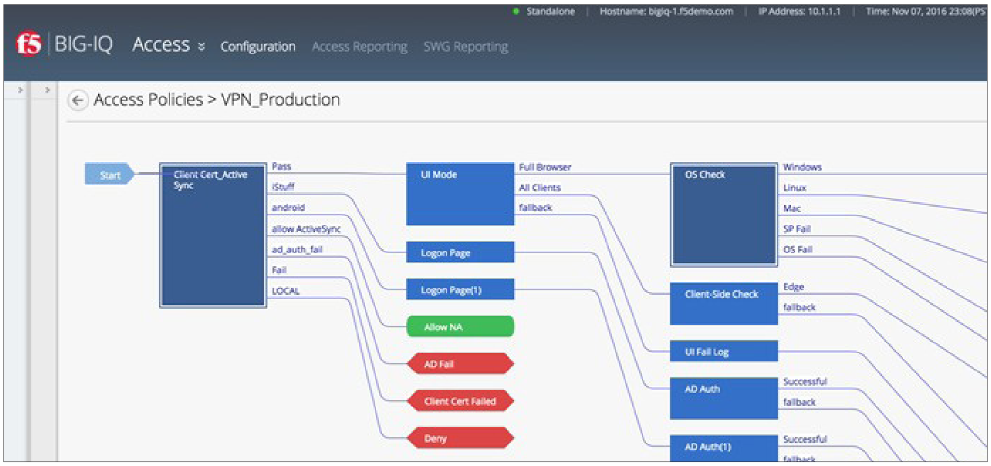

Figure 4: BIG-IQ Centralized Management enables the import, comparison, editing, and updating of access policies across multiple devices from a single interface.

ENHANCE VISIBILITY AND REPORTING

An in-depth view of logs and events provides access policy session details. With reports available through BIG-IQ Centralized Management, BIG-IP APM helps you gain greater visibility into application access and traffic trends, aggregate data for long-term forensics, accelerate incident responses, and identify issues and unanticipated problems before users can experience them.

BIG-IP APM can customize reports with granular data and statistics for intelligent reporting and analysis. Examples include detailed session reports by:

- Access failures
- Users
- Resources accessed
- Group usage
- IP geolocation

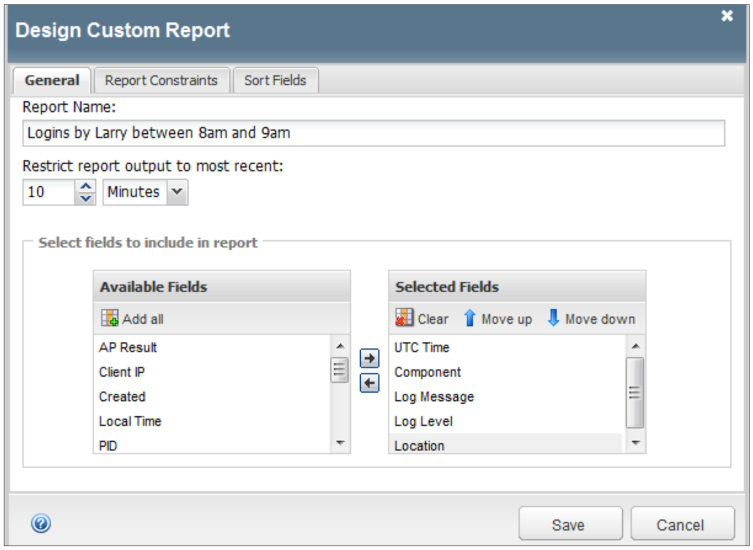

Design Custom Report

General | Report Constraints | Sort Fields

Report Name: Logins by Larry between 8am and 9am

Restrict report output to most recent: 10 Minutes

Select fields to include in report

Figure 5: Custom reports provide granular data and statistics for intelligent analysis.

BIG-IP APM integrates with BIG-IQ Centralized Management to provide enhanced visibility through access reports and logs. It delivers analytical reports and logs based on devices and groups, so you can increase your insight into user access and analysis. It also helps you take quick action if required, including the termination of specific access sessions. In addition, it provides a CSV export of BIG-IP APM report data, so it's accessible for customized reports. BIG-IQ Centralized Management's customized dashboard view helps you to better envision trends and relationship contexts more easily. This improves your response time should issues arise. Through this holistic view of application and network access, you can better understand the effectiveness of the access policies you've established, locate and address weak points, and enhance your responses to issues and concerns.

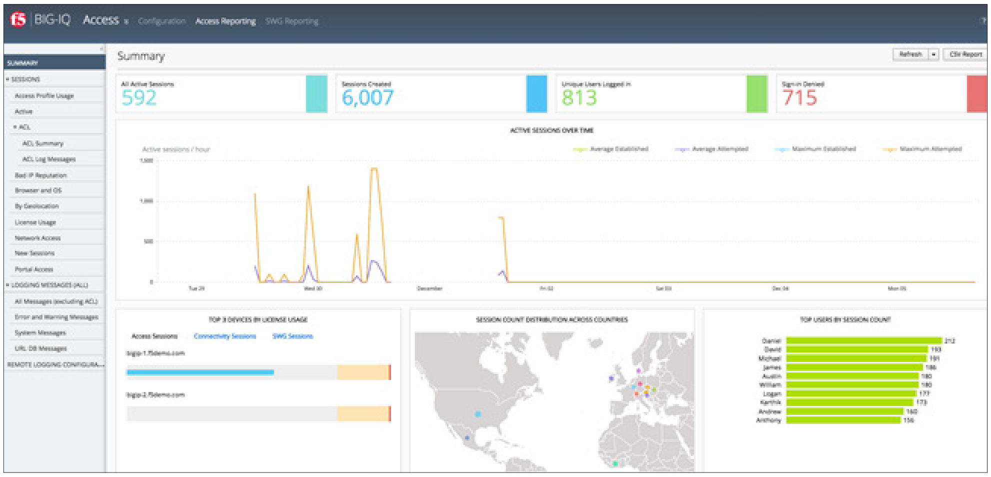

Summary

All Active Sessions 592

Sessions Created 6,007

Unique Users Logged In 813

Sign-in Denied 715

Figure 6: The BIG-IQ Centralized Management comprehensive dashboard for BIG-IP APM helps you better view trends and relationship contexts.

In addition to the access dashboard available through BIG-IQ Centralized Management for BIG-IP APM, the access policy dashboard on the BIG-IP system provides a fast overview of access health. You can view the default template of active sessions, network access throughput, new sessions, and network access connections, or create customized views using the dashboard windows chooser. By dragging and dropping the desired statistics onto the windowpane, you gain a real-time understanding of access health.

UNPARALLELED FLEXIBILITY, HIGH PERFORMANCE, AND SCALABILITY

BIG-IP APM delivers flexible application, network, and cloud access, keeping your users productive and enabling your organization to scale quickly and cost-effectively.

BIG-IP APM can be deployed a variety of ways to address your specific access needs. BIG-IP APM may be:

- Deployed as an add-on module for BIG-IP LTM to protect public-facing applications
- Delivered as a standalone BIG-IP appliance or as standalone F5 VIPRION® chassis
- Included with a BIG-IP LTM Virtual Edition (VE) to deliver flexible application access in virtualized environments
- Run on high-end Virtual Editions and high-performance Virtual Editions
- Offered on a Turbo SSL platform

In addition to being licensed for these platforms, BIG-IP APM may also be licensed as the Best bundle in F5's Good-Better-Best offering, as part of F5 Enterprise Licensing Agreement (ELA) for BIG-IP VEs, and subscription licensing models.

BIG-IP APM is available on a chassis platform and on all BIG-IP appliances. It supports the F5 Virtual Clustered Multiprocessing™ (vCMP) environment. The vCMP hypervisor provides the ability to run multiple instances of BIG-IP APM, resulting in multi-tenancy and effective separation. With vCMP, network administrators can virtualize while achieving a higher level of redundancy and control.

BIG-IP APM offers SSL offload at network speeds and supports up to 3,000 logins per second. For organizations with an ever-growing base of web application users, this solution scales quickly and cost-effectively.

BIG-IP APM use is based on two types of user sessions: access sessions and concurrent connection use (CCU) sessions. Access sessions apply to authentication sessions, IAP, VDI, and similar situations. CCU is applicable for network access, such as full VPN access, application tunnels, or web access. The BIG-IP platform and the VIPRION platform—both of which support BIG-IP APM—handle exponentially more access sessions than CCU sessions in use cases such as authentication, SAML, SSO, and forward proxy. This means that if you intend to use BIG-IP APM for authentication, VDI, and the like, the number of sessions supported on VIPRION can be up to 2 million, and the BIG-IP platform can support up to 1 million.

BIG-IP APM Features

Whether running as a standalone or a bundled BIG-IP platform module or on a VIPRION chassis blade, BIG-IP APM is based on the intelligent, modular F5 TMOS® operating system that delivers insight, flexibility, and control to help you better enable application, network, and cloud access.

BIG-IP APM FEATURES INCLUDE:

- Granular access policy enforcement
- Creating and managing identity- and context-aware policies
- Policy routing
- Support for Identity Aware Proxy (IAP) enabling zero trust application access
- Context-based authorization with dynamic L4/L7 ACLs
- SAML 2.0 identity federation support
- Support for OAuth 2.0 authorization protocol
- Simplified identity federation for applications with multi-valued attributes
- SSO support for classic authentication (Kerberos, header- based, etc.), credential caching, OAuth 2.0, SAML 2.0, and FIDO2 (U2F)
- Integrates with third-party SSO solutions
- Credential caching and proxy for SSO
- Bridging modern authentication and authorization (SAML, OAuth/OIDC) and classic authentication and authorization methods
- OIDC Proof Key for Code Exchange (PKCE) support
- Support for SAML-based authentication using BIG-IP Edge Client and F5 Access for Android and for iOS
- Support for SAML-artifact binding
- Support for SAML ECP profile support
- AAA server authentication and high-availability
- Step-up authentication support
- JSON Web Encryption support for public-facing clients
- Multi- factor authentication (MFA) via one-time password (OTP)
- Seamless integration with third-party MFA solutions
- DTLS 2.0 mode for delivering and securing applications
- SSL VPN remote access

BIG-IP APM FEATURES INCLUDE (CONT.):

- Always connected access
- Establish an always-on VPN tunnel (with Windows OS login and BIG-IP Edge Client for Windows)
- Broad client platform support (see F5 BIG-IP APM Client Compatibility Matrices for each BIG-IP release)
- Robust web browser support (see F5 BIG-IP APM Client Compatibility Matrices for each release)
- Continuous endpoint integrity and security checks
- Support for endpoint security and VPN without web browser plug-ins
- Site-to-site IPsec encryption
- Application tunnels
- Dynamic “webtops,” based on user identity
- Integration with leading IAM vendor products (Microsoft, Okta, Ping Identity)
- Authentication methods: form, certificate, Kerberos SSO, SecurID, basic, RSA token, smart card, N-factor
- User credential protection
- API protection and authorization
- Risk-based access leveraging third-party UEBA and risk engines (HTTP Connector)
- Support for Identity-as-a-Service (IDaaS), including Azure Active Directory and Okta
- Visual Policy Editor (VPE) and Access Guided Configuration (AGC)
- IP geolocation agent (in VPE)
- Windows machine certificate support
- Windows Credential Manager integration
- External logon page support
- Access control support to BIG-IP LTM virtual server
- Scales up to 2 million concurrent access sessions
- BIG-IP Edge Client and F5 Access integrate with VMware Horizon ONE (AirWatch), Microsoft Intune and IBM MaaS360
- BIG-IP Edge Client integration with Windows on ARM64
- Export and import of access policies via BIG-IQ Centralized Management
- Configurable timeouts
- Health check monitor for RADIUS accounting
- Landing URI variable support
- DNS cache/proxy support
- Supports Google reCAPTCHA v2 for authentication and contextual authentication
- IPv6 ready
- Style sheets for customized logon page
- Centralized advanced reporting with Splunk
- vCMP
- F5 iRules® scripting language
- Full proxy
- BIG-IP APM and BIG-IP ASM layering

F5 BIG-IP Platforms

Please refer to the BIG-IP System Hardware, VIPRION, and Virtual Edition data sheets for more details. For information about specific module support for each platform, see the latest release notes on AskF5. For the full list of supported hypervisors, refer to the VE Supported Hypervisors Matrix. F5 platforms can be managed via a single pane of glass with BIG-IQ Centralized Management.

BIG-IP iSeries Appliances

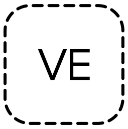

BIG-IP Virtual Editions

VIPRION Chassis

F5 Global Services

F5 Global Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it’s providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services can help ensure your applications are always secure, fast, and reliable. For more information about F5 Global Services, contact consulting@f5.com or visit f5.com/support.

To learn more about BIG-IP APM, visit f5.com/apm.

DC0422 | DS-PROJ-SEC-883309267

DATA SHEET

F5 BIG-IP AFM–SERVICE PROVIDER SECURITY PLATFORM

WHAT'S INSIDE

DEFEND THE NETWORK CORE AND EDGE, AND PROTECT SERVICES

Mobile and fixed line service providers rely on their networks and data centers to drive their service-based revenues. Given their critical nature, these networks have become a major target for attack. While service providers are busy mitigating simple attacks, hackers are using more sophisticated, evolving strikes to disrupt service or steal data.

F5 BIG-IP Advanced Firewall Manager (AFM) is a high-performance, full-proxy network security solution designed to protect networks and data centers against incoming threats that enter the network on the most widely deployed protocols. Built on F5's industry-leading Application Delivery Controller (ADC), BIG-IP AFM gives service providers a scalable, subscriber-aware platform that delivers the flexibility, performance, and control needed to mitigate aggressive distributed denial-of-service (DDoS) and protocol attacks before they overwhelm and degrade services.

BIG-IP AFM's unique application-centric design enables greater effectiveness in guarding against targeted network infrastructure-level attacks. It tracks the state of network sessions, maintains deep subscriber and application awareness, and uniquely mitigates attacks based on more granular details than traditional firewalls. With BIG-IP AFM, organizations receive protection from more than 100 attack signatures—more hardware-based signatures than any other leading firewall vendor—along with unsurpassed programmability, interoperability, and visibility into threat conditions.

ICSA labs CERTIFIED IPSEC

ICSA labs CERTIFIED FIREWALL - CORPORATE

KEY BENEFITS

Ensure services availability
Secure the network edge and core from DDoS and protocol threats with in-depth rules customization, and increased performance and scalability.

Protect with full proxy capabilities
Inspect all incoming subscriber connections and server-to-client responses, and mitigate threats based on security and protocol parameters before forwarding them.

Inspect SSL sessions
Decrypt SSL traffic to identify potentially hidden attacks—at high rates and with high throughput.

Automate security deployment
Simplify configuration with security policies oriented around services and protocols and an efficient rules and policy GUI.

Scale to meet network demand
Meet demands for higher bandwidth usage and concurrency rates with F5's proven virtual software editions and hardware systems to flexibly ensure performance while under attack.

Consistent protection for containerized applications
Protect container-based applications regardless of platform or location with attack detection and mitigation services to mitigate attacks and risks.

Flexible automation options for ease of integration into operations
Extensive integration with third-party and public cloud automation tools to speed BIG-IP AFM into production.

Actionable reporting and visibility
Easily understand your security status with rich telemetry that can be customizable into reports and charts to provide insight to all event types and enable effective forensic analysis.

Reduced Operational Complexity
Single platform to consolidate and deliver Firewall, CGNAT, DNS, protocol protection and deep packet inspection to reduce operational complexity and costs.

SUBSCRIBER SERVICES AND APPLICATION-CENTRIC SECURITY POLICES

FULL PROXY SECURITY

Unlike traditional firewalls, BIG-IP AFM is built on the full-proxy architecture. Incoming connections are fully terminated, inspected for possible security threats, and only then forwarded to the server—assuming no threats are present.

With these full-proxy capabilities, BIG-IP AFM has in-depth understanding of the most commonly used inbound protocols such as HTTP/S, DNS, Diameter, GTP, SSH, ICMP, and TCP, and supports a rich set of services that expand beyond traditional stateful firewall capabilities. Additionally, this security enables deep visibility into connections, allowing data to be manipulated and modified before it's sent to servers or otherwise.

In the reverse direction, server-to-client communication is also proxied. BIG-IP AFM can scrub return data for sensitive information—for instance, protocol response codes that could divulge network information for reconnaissance attacks—and private data, such as credit card or Social Security numbers. The full-proxy design enables termination of SSL, enforcement of security policies, east-west firewall capabilities, and other performance-related services—helping organizations address challenges in volatility inside and outside of the data center.

APPLICATION-CENTRIC SECURITY POLICIES

Gone are the days of mapping applications to zones or scouring through spreadsheets of firewall policies to distinguish attacks on specific applications or to identify the IP address for a particular application server.

Unlike most network security solutions, BIG-IP AFM security policies are logically aligned with the applications in specific traffic flows—streamlining security operations and heightening security effectiveness. However, similar to web application firewall solutions, BIG-IP AFM attaches network security policies to application objects. Details about the application parameters, including server addressing, SSL offload, and access policies, can be grouped together with security parameters, including policies, SSL inspection, and logging. This includes information on which layer 7 protocols are permitted for specific application port access. F5's app-centric approach provides increased efficiency in addressing app concerns and more accuracy in threat detection and policy effectiveness.

Further, since the configuration for an application is unified with an associated network security policy, deprovisioning of applications is also streamlined. When an application is deprovisioned, the obsolete security rules are simultaneously deprovisioned. BIG-IP AFM helps ensure the effectiveness of application deployment and simplifies policy assurance above rigid zone-based or segment-based constructs.

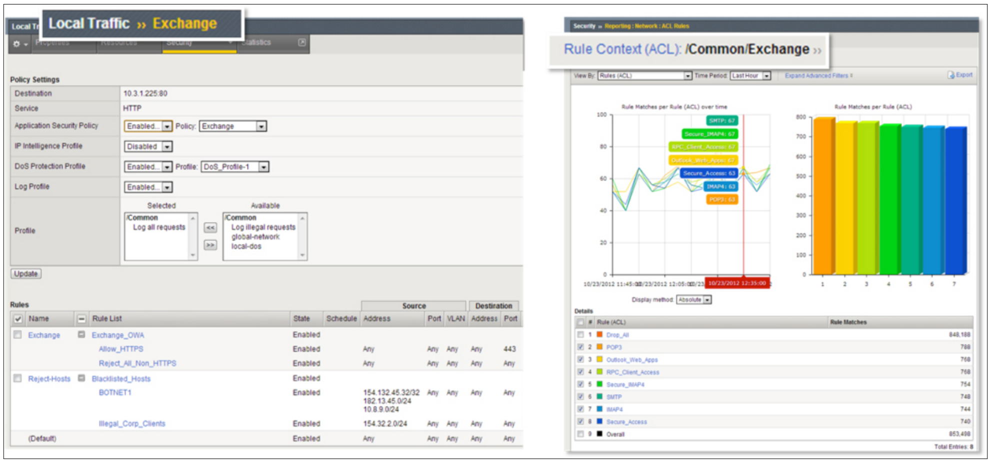

Figure 1: BIG-IP AFM orients firewall policies around the application itself–streamlining security operations.

Local Traffic » Exchange

Policy Settings

Destination	10.3.1.225:80
Service	HTTP
Application Security Policy	Enabled... Policy: Exchange
IP Intelligence Profile	Disabled
DoS Protection Profile	Enabled... Profile: DoS_Profile-1
Log Profile	Enabled...
Profile	Selected: /Common Log all requests Available: /Common Log illegal requests global-network local-dos

Rules

						Source			Destination	
☑	Name	⊟	Rule List	State	Schedule	Address	Port	VLAN	Address	Port
□	Exchange	⊟	Exchange_OWA	Enabled						
			Allow_HTTPS	Enabled		Any	Any	Any	Any	443
			Reject_All_Non_HTTPS	Enabled		Any	Any	Any	Any	Any
□	Reject-Hosts	⊟	Blacklisted_Hosts	Enabled						
			BOTNET1	Enabled		154.132.45.32/32 182.13.45.0/24 10.8.9.0/24	Any	Any	Any	Any
			Illegal_Corp_Clients	Enabled		154.32.2.0/24	Any	Any	Any	Any
	(Default)			Enabled		Any	Any	Any	Any	Any

Security » Reporting : Network : ACL Rules

Rule Context (ACL): /Common/Exchange »

View By: Rules (ACL) Time Period: Last Hour

Display method: Absolute

Details

□	#	Rule (ACL)	Rule Matches
□	1	🟧 Drop_All	848,188
☑	2	🟧 POP3	788
☑	3	🟨 Outlook_Web_Apps	768
☑	4	🟩 RPC_Client_Access	768
☑	5	🟩 Secure_IMAP4	754
☑	6	🟩 SMTP	748
☑	7	🟦 IMAP4	744
☑	8	🟦 Secure_Access	740
□	9	⬛ Overall	853,498

Total Entries: 8

NETWORK DDoS PROTECTION

The full-proxy architecture of BIG-IP AFM helps to ensure the application infrastructure is protected using advanced capabilities to mitigate DoS and DDoS attacks. The out-of-the-box functionality includes a comprehensive set of signatures that enable organizations to defend against, track, and report a breadth of well-known network DDoS attacks and methodologies. Admins can automatically or manually set DDoS threshold values. Furthermore, you can configure packet limits, percentage increases for thresholds, and set absolute rate limits of packets used in attack vectors. Using DoS profiles, BIG-IP AFM performs a variety of checks and mitigates a multitude of attacks, including flood, sweep, teardrop, and smurf attacks, while protecting protocols like SIP and DNS.

BIG-IP AFM also helps to ensure clean pipes for inbound traffic. Using remotely triggered black hole filtering (RTBH), BIG-IP AFM stops attack traffic even before it leaves the ISP network realm. When activated, BIG-IP AFM automatically broadcasts malicious IPs to upstream routers to enforce denylisting through participating ISP routers, ensuring that only good traffic is routed to the data center network and applications within. RTBH functionality leverages the BIG-IP AFM IP shun category denylist that uniquely identifies and blocks malicious L3–L7 attack sources in hardware until feed lists are updated. BIG-IP AFM can also signal and redirect traffic to F5 Silverline for DDoS Protection service. Silverline’s DDoS defense can be either reactive or proactive hybrid DDoS defense—ensuring always-up services by rerouting attacks away from the data center for cloud-based mitigation.

BIG-IP AFM USES HARDWARE-BASED DDOS MITIGATION THAT SCALES TO PREVENT HIGH-VOLUME, TARGETED, NETWORK FLOOD ATTACKS—WHILE ALLOWING LEGITIMATE TRAFFIC TO FLOW WITHOUT COMPROMISING PERFORMANCE.

BIG-IP AFM offers more granularity and visibility into traffic and DDoS attacks than most solutions, with detailed logging and reporting of attack detection and mitigation. It also delivers increased SYN cookie protections, per-server granular DDoS policies, IP reputation intelligence, and custom allowlist and denylist support. BIG-IP AFM uses hardware-based DDoS mitigation that scales to prevent high-volume, targeted, network flood attacks—while allowing legitimate traffic to flow without compromising performance.

Dynamic Denylisting

Attackers are continually changing their tactics in order to circumvent security controls. One of the more useful evasive tactics used by attackers is changing different IP addresses. By changing their IP addresses and associated domain names they can avoid being blocked and continue their attacks.

BIG-IP AFM provides IP denylisting capabilities that help organizations to minimize enforcement time of dynamic security controls that guard against known malicious IPs. IP denylisting complements existing IP intelligence services. It facilitates more immediate filtering of malicious traffic until intelligence feeds containing denylisted IP addresses are updated. Up to 100,000 entries can be denylisted almost instantaneously to enable temporary, immediate blocking (or allowlisting) of malicious IPs. IP denylisting reduces time-to-enforcement and increases speed of mitigation based on real-time intelligence from BIG-IP AFM, other BIG-IP modules, and third-party monitoring systems.

IP Intelligence

Organizations today are exposed to a variety of potentially malicious attacks from rapidly changing IP addresses. A major advantage in your network protection scheme is the ability to anticipate, detect, and respond to threats before they hit the data center. BIG-IP AFM integrates with **F5 IP Intelligence Services** for stronger context-based security that strategically guards against evolving threats at the earliest point in the traffic flow.

IP Intelligence Services minimizes the threat window and enhances BIG-IP AFM DDoS and network defense with up-to-date network threat intelligence for stronger, context-based security. It maintains information on more than one million malicious URL and IP addresses, and can effectively block connections to and from those addresses. To minimize the threat window and keep an organization's data and reputation safe, the IP Intelligence Services database of addresses is refreshed every five minutes from the cloud. Administrators can assign default classes and behaviors to feed lists, allowing more control for each IP intelligence category by specifying response actions and default logging for each policy. IP Intelligence Services reduces risk and increases data center efficiency—eliminating the effort to process bad traffic.

IN-DEPTH INFRASTRUCTURE PROTECTION

SSH Channel Protection

BIG-IP AFM uniquely controls operations in the SSH channel and helps prevent data breaches, malware distribution, and compliance failures. When deployed in front of SSH servers, BIG-IP AFM acts as a man-in-the-middle SSH proxy—filtering SSH traffic, and controlling access to files, databases, and system information by limiting task users can perform. Unlike leading firewalls, SSH policies limit permissible actions per user or per virtual server to strengthen security on SSH channels—tracking usage and preventing misuse of SSH channels by employees and contractors and stopping east-west attacks that move throughout the infrastructure. Additionally, BIG-IP AFM prevents SSH sessions from remaining open indefinitely and ensures effective and continuous SSH key management for tighter security and compliance.

BIG-IP AFM BENEFITS FROM THE EXTENSIBILITY OF F5 IRULES, ALLOWING ADMINISTRATORS TO EXPAND FUNCTIONALITY AND DEPLOY CUSTOM RULES THAT PROTECT AGAINST COMPLEX AND MULTI-LEVEL ATTACKS.

Unsurpassed Flexibility and Extensibility

Rapid response is vital in minimizing risk imposed by uncommon attacks. Many firewalls fail to secure the perimeter when faced with less common attacks like Heartbleed. As a component of the F5 BIG-IP® platform, BIG-IP AFM benefits from the extensibility of F5 iRules®, allowing administrators to expand functionality and deploy custom rules that protect against complex and multi-level attacks.

F5 iRules is a scripting language with open APIs that can operate directly on payloads in the data plane. With iRules, administrators can create custom rules to mitigate uncommon, highly sophisticated DDoS attacks that may not be covered by the BIG-IP AFM packaged solution. The scope of iRules commands provides deep visibility into packets, especially IP/TCP header fields, enabling effective L2–L4 DDoS signatures and flow control via iRule signatures. iRules benefits from BIG-IP AFM anti-DDoS support, which distinguishes between good and bad traffic based on signature(s) and takes action to block, drop, log, redirect, or stop traffic for inspection based on signature matching.

With iRules customization, capabilities including IP intelligence, geolocation features, and statistical sub-sampling can also be applied. iRules has been leveraged by the F5 DevCentral™ community of over 250,000 users, collaborating and creating custom rules that mitigate fewer common threats. These rules are shared to enable other administrators to flexibly expand the functionality of BIG-IP AFM deployments.

PROTECTING NETWORK RESOURCES WITHOUT COMPROMISING FLEXIBILITY AND CONTROLLING COSTS IS A CONSTANT BATTLE. BIG-IP AFM WITH ITS UNMATCHED SUBSCRIBER AGGREGATION CAPACITY ENABLES SERVICE PROVIDERS TO PROTECT NETWORKS AND SUBSCRIBERS WHILE CONSOLIDATING NETWORK INFRASTRUCTURE WITHOUT COMPROMISING FLEXIBILITY.

N6/SGi Firewall

Mobility service providers face not only intense competitive pressure, but also continual attacks on their infrastructure designed to impact subscriber services. Protecting network resources without compromising flexibility and controlling costs is a constant battle. BIG-IP AFM with its unmatched subscriber aggregation capacity enables service providers to protect networks and subscribers while consolidating network infrastructure without compromising flexibility.

In mobile networks, BIG-IP AFM forms the basis of the F5 S/Gi firewall solution. Deployed at the Gi interface of 3G networks and the SGi interface of 4G/LTE networks, the S/Gi firewall solution enforces network perimeters, protects the mobility infrastructure and subscribers. It gives service providers the scalability and flexibility for advanced protocol and service enforcement. The S/Gi firewall solution takes advantage of F5's intelligent services framework, meaning service providers can consolidate additional network and security functions such as carrier-grade NAT (CGNAT)and subscriber traffic visibility—all on a single platform. The F5 iSeries and VIPRION platforms have built-in Telco-grade dependability to ensure consistent service operation under heavy workloads. For encryption/decryption, both the iSeries and VIPRION are FIPS-140 certified to ensure the highest level of protection for subscriber traffic.

CGNAT

The worldwide proliferation of wireless and Internet-enabled devices has led to the rapid depletion of IPv4 addresses. All of the five Regional Internet Registries (RIR) has exhausted its IPv4 allocations, and final pool exhaustion happened in November 2020; meanwhile, IPv6 adoption continues to grow. Service providers need a solution that will help them manage IPv4 address depletion and increase network optimization by seamlessly migrating to IPv6.

F5 BIG-IP Carrier-Grade NAT (CGNAT) offers a broad set of tools that enables service providers to successfully migrate to IPv6 while continuing to support and interoperate with existing IPv4 devices and content. BIG-IP CGNAT offers service providers tunneling solutions with Dual-Stack Lite capabilities as well as native network address translation solutions, such as NAT44 and NAT64. It provides carrier-grade scalability by offering a very high number of IP address translations, very fast NAT translation setup rates, and high throughput. Granular, high-speed Netflow logging brings visibility and control to data analytics. Data collection per protected object is configurable, and can be collected for specific time intervals to simplify analysis and reduce the time to problem resolution.

DNS Security

DNS servers are critical to any mobile or fixed line network operator. Security is paramount to operations, as the DNS protocol serves as the basis for internet infrastructure mapping of the web domains that subscribers access. BIG-IP® DNS delivers an intelligent and scalable DNS infrastructure that gives mobile users faster access and service response. This makes it easy for service providers to optimize, monetize, and secure their DNS infrastructures. F5 DNS provides carrier-grade, high-performance LDNS caching and resolving, and is a hyper-scale authoritative DNS solution that handle business growth and sudden demand spikes.

BIG-IP DNS DELIVERS AN INTELLIGENT AND SCALABLE DNS INFRASTRUCTURE THAT GIVES MOBILE USERS FASTER ACCESS AND SERVICE RESPONSE. THIS MAKES IT EASY FOR SERVICE PROVIDERS TO OPTIMIZE, MONETIZE, AND SECURE THEIR DNS INFRASTRUCTURES.

BIG-IP AFM shields the DNS infrastructure from malicious attacks designed to reduce DNS and service performance, launched by infected subscribers from undesired DNS queries and responses. F5’s intelligent protection of DNS services inspects and validates traffic adherence to the DNS protocol while dropping invalid requests or refusing to accept unsolicited responses. BIG-IP AFM is an ICSA Labs certified network firewall with DDoS threshold alerting that hyper-scales across many devices using IP Anycast for DDoS absorption. It mitigates threats by blocking access to malicious IP domains.

Intrusion Prevention Security

Service provider networks operate a number of protocols that enable fixed and mobile subscriber connectivity. These protocols have weak or non-existent built-in security and can be exploited by an attacker to steal information or impact services. Protecting them is a must in order to deliver an experience that meets customers’ expectations in a competitive marketplace.

BIG-IP AFM Intrusion Prevention System (IPS) delivers deep packet inspection and visibility for incoming network traffic. BIG-IP AFM IPS engine performs Layer 5-7 traffic inspection for security incidents, protocol/application violations and exploits to take appropriate action for prevention. It reviews traffic for adherence to 25+ protocol standards and matches against hundreds of known attack signatures and exploits.

For Mobility service providers, BIG-IP AFM IPS performs traffic inspection and protocol adherence for SS7, Diameter (FS.19 compliance checks), HTTP/2, GTP (FS.20 compliance checks), SCTP and SIP to ensure that these application servers are not attacked or exploited.

The Central management interface for IPS leverages the industry standard domain specific language (SNORT) to provide policies and signatures to validate traffic against. This also enables the ingestion and utilization of existing SNORT policies and signatures from other sources to ensure consistent security.

BIG-IP AFM IPS is architected upon the F5 BIG-IP platform, and designed for native multi-threaded processing, for ease of scalability—enabling it to handle traffic spikes or planned growth without compromising services. The multi-threaded architecture ensures high-availability and enables hit-less upgrades to enable non-stop business operations and risk exposure.

BIG-IP AFM IPS's architecture supports rapid visibility of threats via a BIG-IQ dashboard. High-speed logging for near real-time visibility and control is also supported to third-party SIEM platforms. To further help service providers keep up with evolving attacks and hacks, BIG-IP AFM IPS has incorporated a traffic learning capability that deeply monitors traffic and auto-develops policies. The Traffic Learning capability suggests rules based on traffic patterns that can be accepted or denied manually or automatically to easily add protection without expertise or overhead.

WHEN BIG-IP AFM IS DEPLOYED IN A SERVICE PROVIDER'S NETWORK, IT FEATURES KEY DIFFERENCES WHEN COMPARED TO TRADITIONAL NETWORK FIREWALLS, MAKING IT MORE EFFECTIVE FOR IOT SERVICES.

IoT Protection

The Internet of things (IoT) comes in all shapes and sizes: from three-ton automobiles to clothing to under-the-skin blood sugar monitors and even entire homes. With IoT, communication across the internet is enabled by services from service provider mobile and fixed-line networks. Managing IoT traffic does not come without its challenges. Threats to service provider networks and data centers must be addressed in order for IoT services to be successful.

When BIG-IP AFM is deployed in a service provider's network, it features key differences when compared to traditional network firewalls, making it more effective for IoT services. BIG-IP AFM provides device-aware, application-centric security policies. This allows service providers to offer IoT security services without the need to host the IoT application in their data centers, or directly manage the IoT application.

BIG-IP AFM as an IoT firewall mitigates threats by stopping DDoS and application-layer attacks which may disrupt the integrity and availability of the service provider's network. BIG-IP AFM also ensures that devices are only connecting to 'safe' locations and prevents devices from connecting to unknown services. This reduces the chances of devices being compromised through malware and blocks malicious 'ThingBot' C&C (command and control) communication to stop devices from being exploited remotely.

IPSEC

The explosion of applications within the data center, in the cloud, and out at the edge infrastructure creates a challenge in protecting the data from theft or hack attacks. Enabling and managing encryption and decryption on each application is prohibitive due to operational and infrastructure costs. Yet, encrypting and decrypting traffic between different application hosts and clients spread across networks and remote locations is critical to protecting data and ensuring application integrity.

BIG-IP AFM integrated IPSEC encryption and decryption enables secure data connectivity between hosts and networks to eliminate risk of data theft and application attacks. AFM supports IPsec ESP and AH connections to create secure connections for data transmission.

Protection for Container-based Apps

Application development today is based on a container infrastructure that enables faster application development with greater scale-out capabilities. Container based applications can easily grow and can become distributed; but container-based applications, and the underlying infrastructure, need to be protected just like traditional applications.

BIG-IP AFM Virtual Edition (VE) supports running in both public and private cloud environments and provides protection that readily secures container-based applications by off-loading the “North/South” decryption and encryption of traffic to and from container-based application environments. BIG-IP AFM also monitors and provides visibility for non-volumetric attacks on container-based applications and enables fast development of protection via iRules.

WITH ADVANCED LOGGING AND INTELLIGENT THREAT REPORTING CAPABILITIES, BIG-IP AFM LOGS MILLIONS OF RECORDS IN REAL TIME, PROVIDING GRANULAR VISIBILITY INTO DDOS ATTACKS FOR IN-DEPTH ANALYSIS OF SECURITY EVENTS.

Deep Visibility and Reporting

IT and security teams struggle to collect enough threat intelligence, and analyze enough data, to implement the right security measures. BIG-IP AFM gives organizations deep insight into attacks and mitigation techniques, enabling them to make more informed decisions that increase overall security effectiveness.

With advanced logging and intelligent threat reporting capabilities, BIG-IP AFM logs millions of records in real time, providing granular visibility into DDoS attacks for in-depth analysis of security events. BIG-IP AFM reports provide clear, concise, and actionable information highlighting attacks and trends with drill-down and page-view capabilities. These offer comprehensive details into attacks, threat progression, and firewall BIG-IP AFM health.

With BIG-IP AFM, organizations can also benefit from F5 Analytics, a module of the BIG-IP platform, which combines DDoS reports from BIG-IP ASM and BIG-IP AFM for a single comprehensive view of the entire threat field. F5 Analytics, previously known as the Application Visibility and Reporting module, allows administrators to view and analyze metrics gathered about the network and servers as well as the applications themselves. Additionally, BIG-IP AFM uses SNMP and JSON reporting to easily communicate DDoS attack details and other key events to higher-level monitoring and forensics systems. These systems offer greater analysis that strengthens the organization's overall security posture.

Automation and Centralized Policy Management

Large organizations face a growing challenge in managing a consistent and effective security posture across an ever-expanding number of firewall devices. Too often, security administrators must independently manage each device, reducing operational scalability and increasing overhead costs.

F5 BIG-IQ CENTRALIZED MANAGEMENT PROVIDES CONSOLIDATED, CUSTOMIZABLE, AND SECURITY-SPECIFIC DASHBOARDS THAT PROVIDE NEAR REAL-TIME VISIBILITY INTO TRAFFIC FROM MULTIPLE BIG-IP AFM INSTANCES BY SOURCE AND TRAFFIC TYPE.

F5 BIG-IQ® Centralized Management enables administrators to easily manage and orchestrate F5 devices and the services they deliver, including the security services of BIG-IP AFM. F5 BIG-IQ Centralized Management provides consolidated, customizable, and security-specific dashboards that provide near real-time visibility into traffic from multiple BIG-IP AFM instances by source and traffic type. Leveraging these insights and proactive alerting capabilities, BIG-IQ makes it easy to react quickly to changes and push updated and optimized configuration and policies to many BIG-IP AFM deployments from a single, unified console. BIG-IQ also leverages simple, yet powerful REST APIs to make integrating with third-party services quick and easy. And, with its tight integration with the F5 Automation Toolchain, injecting automation into app services and infrastructure creation, provisioning, and management, you can take a programmatic approach to working with BIG-IP AFM and other F5 solutions.

BIG-IP AFM is also easily automated into production via integrations with third party integrations. F5 has built solution templates for AWS, Azure, and Google public cloud environments and API integration with third-party automation tools such as Ansible and Terraform.

Increased Scalability, Performance, and Reliability

BIG-IP AFM delivers the scalability and performance to tackle the most demanding firewall requirements with outstanding speed and throughput. A single F5 platform BIG-IP AFM uses F5 ScaleN™ with Virtual Clustered Multiprocessing™ (vCMP) enabled systems to give cloud and communications service providers, as well as enterprises, the most cost-effective approach for managing their large-scale firewall deployments.

With vCMP, administrators can easily consolidate multiple firewalls onto a single device and allocate BIG-IP AFM resources in a more flexible and isolated manner than with firewalls for different customers, groups, applications, and services. vCMP supports high-density firewall isolation and guest firewall clustering for easier administration and maintenance and to ensure consistency throughout the firewall infrastructure.

Consolidated Infrastructure Platform

BIG-IP AFM is a core component of F5's solutions for application protection, which combines network security capabilities with traffic management, application security, and DNS security. These solutions can be consolidated onto a single BIG-IP platform, reducing management complexity and overhead, while offering superior performance and scalability. Building upon BIG-IP® Local Traffic Manager™ (LTM), the consolidated protection delivers deep application fluency for the most widely deployed enterprise applications and service provider protocols. This makes it the ideal platform for security standardization for protecting Internet-facing data center and distributed edge applications.

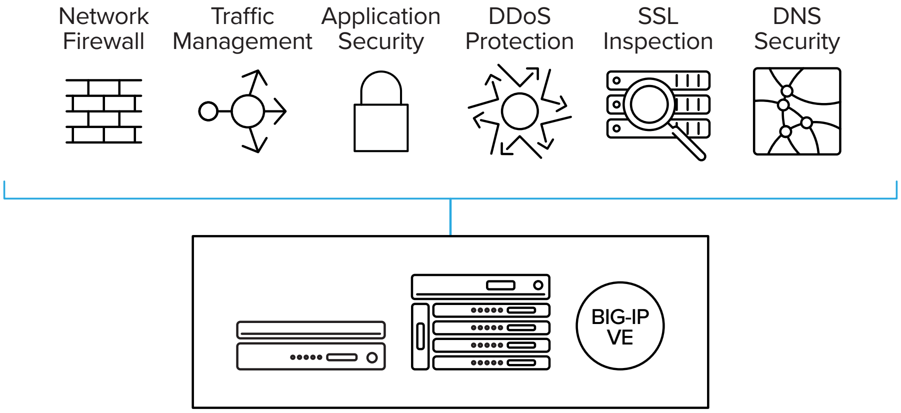

Figure 2: F5's solutions for application protection bring together key network and security functions on a single platform.

F5's solutions for application protection are made up of the following BIG-IP modules:

- **BIG-IP AFM**—This advanced network security solution forms the core of the F5 application protection solution. It provides full SSL visibility at scale, as well as network-layer and session-layer DDoS mitigation.
- **BIG-IP Advanced Web Application Firewall (Advanced WAF)**—Delivers application security, web scraping and bot prevention, and HTTP DDoS mitigation.
- **BIG-IP Local Traffic Manager (LTM)**—Provides advanced traffic management, load balancing, and application delivery.
- **BIG-IP DNS**—Hyperscales and secures the DNS infrastructure during DDoS attacks and keeps global applications online.
- **BIG-IP IPS**—Intrusion Prevention protects infrastructure and protocols and compliance verification.
- **IP Intelligence and Geolocation**—These additional services provide IP reputation and geolocation information for added context-aware security.

FEATURES AND SPECIFICATIONS

BIG-IP AFM is a stateful, full-proxy security solution that provides advanced network protection and capabilities that exceed traditional firewalls.

Protocol anomaly detection	Yes—SYN/ICMP/ACK/UDP/TCP/IP**/DNS/ARP
DoS and DDoS protection	Yes—L3, L4, SSL/TLS, HTTP, Flood, Sweep
Remotely trigger black hole filtering (RTBH)	Yes
SSH Proxy	Yes
Port-misuse protection	Yes
SSL/TLS Reverse proxy	Yes
IP reputation* and geolocations	Yes—including identifying Tor proxies, malware, and command-and-control (C&C) servers
Central management w/RBAC	Yes—with BIG-IQ Centralized Management
SNMP reporting	Yes
DDoS traffic sampling	Yes

* - licensed separately

** - IPv4 and IPv6 supported

BIG-IP AFM Availability

BIG-IP AFM is available with other modules to enable specific infrastructure, protocol and application security use cases, as follows.

NAME	AFM	IPS	LTM	ADV. WAF	APM
AFM's base service provider security platform (protocol protection, visibility and automation, and ICSA-certified firewall)	●				
AFM with Intrusion Prevention System (IPS) for advanced traffic inspection and security	●	●			
AFM with Application Delivery Controller (ADC) for load balancing of inbound subscriber traffic	●		●		
AFM with Advanced Web Application Firewall for protecting layer 7 applications from automated and manual attacks	●		●	●	
AFM with user access management	●		●	●	●
AFM with Access Policy Manager for user access management	●		●	●	●
AFM Add-on for existing LTM platforms (hardware or software)	●				

Note: All BIG-IP AFM licenses include protocol security, routing, and maximum SSL. IP Intelligence and Geolocation are available add-ons for all bundles.

BIG-IP AFM is available as an add-on module for integration with BIG-IP Local Traffic Manager on any BIG-IP platform. For detailed physical specifications, please refer to the BIG-IP System Hardware Data sheet.

PLATFORMS AND SERVICES

BIG-IP LTM Virtual Edition

BIG-IP LTM Virtual Edition (VE) is a version of the BIG-IP system that runs as a virtual machine. BIG-IP AFM can be deployed on a virtual edition. BIG-IP VEs include all features of BIG-IP devices running on the standard F5 TMOS, except as noted in release notes and product documentation. BIG-IP AFM VE’s can be optimized against DDoS attacks and for SSL/TLS processing with Intel SmartNIC and Quick-assist Technology (QAT). Intel SmartNIC includes a FPGA which pre-processes DDoS attacks out of the traffic before it hits the Intel x86 core CPU for processing. For SSL/TLS decryption/encryption, the Intel QAT processor is supported for off-loading this intensive processing from the core CPU for improved performance.

VIPRION Platforms

BIG-IP AFM is also available as an add-on module to BIG-IP Local Traffic Manager on the modular F5 VIPRION® platform. This chassis and blade architecture enable simple scalability as your Application infrastructure grows. The VIPRION platform is NEBS Level 1 and FIPS 140-2 Level 2 compliant to ensure adherence to industry standards and compliance. See the VIPRION Data sheet for details.

BIG-IP Platforms

F5’s next-generation, cloud-ready ADC platform provides DevOps-like agility with the scale, security depth, and investment protection needed for both established and emerging apps.
The new BIG-IP® iSeries appliances deliver quick and easy programmability, ecosystem-friendly orchestration, and record breaking, software-defined hardware performance. The F5 BIG-IP iSeries 15K is designed specifically to meet service providers’ performance requirements in a 1U platform.

As a result, customers can accelerate private clouds and secure critical data at scale while lowering TCO and future-proofing their application infrastructures. F5 solutions can be rapidly deployed via integrations with open source configuration management tools and orchestration systems.

In addition to the BIG-IP iSeries, F5 offers VIPRION modular chassis and blade systems designed specifically for performance and for true on-demand linear scalability without business disruption. A single F5 platform scales to handle up to 576 million concurrent connections, 640 Gbps of throughput, and 8 million connections per second to mitigate even the largest volumetric attacks. VIPRION systems use F5’s ScaleN clustering technology to add blades without reconfiguration or rebooting.

Figure 3: BIG-IP iSeries Appliance, VIPRION Chassis, and BIG-IP Virtual Editions.

Virtual editions of BIG-IP software run on commodity servers and support the range of hypervisors and performance requirements. These virtual editions provide agility, mobility, and fast deployment of app services in software-defined data centers and cloud environments.

See the BIG-IP System Hardware, VIPRION, and Virtual Edition data sheets for more details. For information about specific module support for each platform, see the latest release notes on AskF5. For the full list of supported hypervisors, refer to the VE Supported Hypervisors Matrix.

Business-aligned Licensing

Meeting your applications' needs in a dynamic environment has never been easier. F5 BIG-IP AFM is available via F5's Good, Better, Best licensing that provides you with the flexibility to provision advanced modules on demand, at the best value.

- Provision modules needed to run your applications with F5's Good, Better, Best offerings
- Implement complete application flexibility with the ability to deploy your modules on a virtual or physical platform

BIG-IP AFM VE can be utilized in accordance with business operational needs. F5 VE's can be purchased via Subscription, Enterprise License Agreement, Perpetual for on-premises and/or public multi-cloud architectures. Public cloud 'pay-as-you-go' (PAYG) is also supported.

F5 Professional Services

F5 Global Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Global Services can help ensure your applications are always secure, fast, and reliable. For more information about F5 Global Services, contact consulting@f5.com or visit f5.com/support.

F5 Global Support Services

F5 Global Services helps you deliver your applications with the availability, performance and security users expect. With deep expertise in F5 application services, as well as the ecosystems that they operate in, our teams can help you support new business initiatives including automation, migrating apps to the cloud, ensuring consistent security and deploying new application architectures.

F5 provides the services, support, and training you need to ensure success across the lifecycle of your deployments, from implementation to maintenance to optimization. We help you speed deployment, drive operational efficiencies and get to market as quickly as possible.

F5 offers a range of support packages that provide best in class technical expertise to help solve your challenges. You can take advantage of flexible options including our robust, self-solve knowledge repository, communities and phone. Our worldwide support centers are available 24/7 to provide help wherever and whenever you need it.

Service Provider Essentials

The F5 Service Provider Essentials technical services program (SPE) offers a specialized level of service designed for, and exclusively available to, our service provider customers in almost every country. It provides service provider operations teams with the specialized support they need to effectively manage their F5 application delivery estate. F5 SPE combines high-quality, service provider-specific technical assistance including a dedicated service provider network support team, consulting hours inside the maintenance contract, and high-urgency incident management. Learn more at f5.com/SPE.

Premium Plus

Premium Plus customers receive priority status at F5, with fast, easy access to remote assistance from a dedicated team of senior-level, F5-certified Network Support Engineers and a Service Delivery Manager. You also get proactive support for planned maintenance, advance RMA replacement, software upgrades, and help with F5 iRules scripts. Your Premium Plus team will also work with you to create an IT environment that addresses your business goals. Learn more at www.f5.com/services/support.

MORE INFORMATION

To learn more about BIG-IP AFM and complimentary solutions, visit f5.com to find these and other resources:

Data sheets

BIG-IP Advanced WAF
IP Intelligence Services
Access Policy Manager

Web pages

F5 BIG-IP AFM

Additional resources

Intelligent DNS Firewall for Service Providers
Diameter Security
Near Real-Time DNS Reporting Attack Mitigation Case Study
Carrier-Grade NAT for Service Providers
Key Use Cases for GTP