



Consistently Deploy F5 Application Delivery and Security Services in Multi-Cloud Environments

What's Inside

- 2 Primary Cloud Scenarios
- 2 Private Cloud Using Software-Defined Architectures
- 3 Dynamic F5 App Services in Container Environments
- 3 Deploy Applications in the Leading Public Clouds
- 4 Application Mobility Across Hybrid and Multi-Cloud Environments
- 4 Consistent App Delivery and Security for Public Cloud
- 5 Integration with SDN Frameworks
- 5 Automation and Orchestration Through Granular Programmability
- 5 Centralized Management and Licensing
- 6 Specifications
- 9 BIG-IP Per-App Virtual Editions
- 9 Get Started Today
- 9 F5 Global Services
- 10 DevCentral

Software-based application delivery services are critical to maintaining the adaptable and secure application infrastructure demanded by enterprises undergoing a digital transformation. F5 can help your transition to 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. In addition, many enterprises have or are planning to deploy applications in multiple clouds, both public and private. Standardizing on F5 app services results in faster cloud deployments, architectural flexibility as needs change, and effective, consistent security across application workload environments.

F5® BIG-IP® virtual editions (VEs) are virtual application delivery controllers (vADCs) that can be deployed on all leading hypervisors and cloud platforms running on commodity servers. BIG-IP VEs deliver all the same market-leading application delivery services—including advanced traffic management, acceleration, DNS, firewall, and access management—that run on F5 purpose-built hardware. VE software images are downloadable and portable between on-premises virtualized data center, public cloud, and private cloud environments. With BIG-IP virtual editions and F5 BIG-IQ® Centralized Management solutions, you can rapidly provision consistent application services across the data center and into the cloud.

Key Benefits

Deploy with increased agility

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

Achieve automation and orchestration in cloud architectures

Automate deployment and configuration or integrate with leading orchestration frameworks—in cloud or software-defined networking (SDN) environments through cloud solution templates, REST APIs, and granular programmability.

Optimize application and security services

Rapidly provision and consolidate application services on your existing servers, unlocking the broadest feature density through flexible licensing models that align to your business needs.

Provide the ultimate in flexibility

Get the most flexible deployment options in the industry, with support across all major virtualization and container platforms for both private and public cloud environments.

Primary Cloud Scenarios

BIG-IP virtual editions (VEs) can be used to deliver a consistent set of advanced application services in all three primary cloud scenarios: private cloud/software-defined data center (SDDC), public cloud, and multi/hybrid cloud as described below.

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, such as 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, 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.




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

Dynamic F5 App Services in Container Environments

Organizations are rapidly adopting containerized environments to develop and package applications. These organizations typically use management and orchestration frameworks to coordinate provisioning and automation of containers. Developers must also be able to easily configure application services in those containerized environments.

F5 Container Connector is a container integration solution that helps developers and system teams manage ingress control and F5 application services in container and Platform as a Service (PaaS) deployments. Container Connector integrates BIG-IP platforms with native PaaS container orchestration and management systems, such as Kubernetes, Red Hat OpenShift, Pivotal Cloud Foundry, and Mesos/Marathon. It provides integrated automation of application performance, routing, and security services during development for a faster time to market.

Deploy Applications in the Leading Public Clouds

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. Utilizing F5 application and security services delivered by BIG-IP VEs provides the following benefits:

- **Faster time to market**—Rapidly provision advanced application services when launching new applications or migrating existing applications to public cloud providers.
- **Consistent levels of availability, performance, and security**—Provide your customers the same user experience while protecting both your revenue and reputation.
- **Flexible licensing models**—Bring your own license to start; pay as you go for just what you need through hourly or monthly utility billing, available in the leading cloud providers, including: Amazon Web Services (AWS), Microsoft Azure, Google Cloud Platform, and IBM Cloud. Alternatively, you can use annual subscription-based licensing or three-year Enterprise Licensing Agreement (ELA) to align with an OpEx consumption model.
- **Integration with public cloud providers**—Dynamically scale out app services through integration with AWS Auto Scaling or protect your apps with an out-of-the-box preconfigured web application firewall (WAF) solution in the Azure Security Center.




Figure 2: Deploy applications in the leading public cloud providers with BIG-IP application and security services.

Application Mobility Across Hybrid and Multi-Cloud Environments

Despite the many benefits of public cloud deployments, enterprises often avoid moving all applications or data to public cloud IaaS providers due to the loss of control, risk, regulatory compliance, and lack of support in legacy application design. Ensuring application mobility across a combination of private and public cloud environments requires a consistent set of application services and security policies. These application and security services provide optimized user experience, protection against L3–7 attacks, bursting to the public cloud for temporary or seasonal spikes in application usage, and global availability and disaster recovery.

Consistent App Delivery and Security for Public Cloud

As more enterprise applications migrate to the public cloud, it's becoming more difficult to maintain network requirements and consistent application policies. In addition, many enterprises are unaware of the number of public cloud apps deployed, the configurations and services for those apps, and how to discover apps in flight. Organizations are turning to interconnection services at colocation providers for direct public cloud access. However, application delivery and security services across public cloud providers is disparate and varied, compared to on-premises and private cloud solutions.

Application Connector is an add-on to the F5 BIG-IP platform, allowing services insertion for public cloud applications. It also acts as a cloud proxy instance for securely connecting public clouds to an organization's application service infrastructure within cloud interconnects (colocations) or data centers. Application Connector performs automatic workload discovery within AWS and Azure public clouds and provides a secure connection back to interconnect services or the data center.




Figure 3: F5 BIG-IP delivers a consistent set of services and policies across private data centers and multiple public cloud providers.

Integration with SDN Frameworks

Software-defined networking (SDN) achieves agility, flexibility, and cost efficiency as it applies to 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 between virtualized networks and traditional network architectures to provide a smooth transition and investment protection.

Automation and Orchestration Through Granular 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.

From the beginning, F5 iRules scripting has provided granular traffic control and visibility, enabling customization, rapid response to errors in application code and security vulnerabilities, and support for new protocols. With the addition of F5 iApps templates, deployment and configuration of application services is automated and goes from weeks to minutes. F5 iControl REST APIs and SDKs provide integration with open source and commercially developed orchestration, cloud (AWS, Azure, Google Cloud Platform, OpenStack), and configuration management (Puppet, Chef, Ansible) systems.

In addition, F5 cloud solution templates enable DevOps teams to quickly deploy common F5 services, including autoscaling traffic management and WAF, in one click. Best of all, these solutions are available on [GitHub](#) and cloud provider marketplaces with full support by F5.

Centralized Management and Licensing

BIG-IP Centralized Management is an intelligent framework for managing F5 application delivery and security solutions. It provides a single pane of glass to manage and deploy F5 devices, including key BIG-IP modules such as BIG-IP Local Traffic Manager™ (LTM), BIG-IP Application Security Manager™ (ASM), BIG-IP Advanced Firewall Manager™ (AFM), BIG-IP Access Policy Manager® (APM), as well as F5 WebSafe™. Use BIG-IP Centralized Management to:

- Track devices.
- Back up images and configurations.
- Monitor dashboards, reporting, and alerting.
- Provide role-based access control (RBAC).
- Manage BIG-IP VE licenses.
- Ensure consistent security and traffic management policies across your infrastructure.

License management of VEs lets you automate large-scale virtual ADC deployments, including per-app VEs, in all clouds via an F5 subscription or ELA licensing. With BIG-IP 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

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

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
Performance	Starting	Maximum*

Software compression

Compression throughput	20 Mbps	4 Gbps
------------------------	---------	--------

DNS

Query response per second	1,000	250,000
---------------------------	-------	---------

BIG-IP APM/SWG

BIG-IP APM access sessions	500	10,000
BIG-IP APM concurrent users	50	2,500
Secure web gateway sessions	100	100

Figure 4: 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.

* 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

Performance	Maximum*
L7 requests per second	2 million
L4 connections per second	625,000
L4 throughput	40 Gbps

SSL	
SSL RSA TPS (2K keys)	9,700
SSL throughput (RSA)	6.9 Gbps
SSL ECC TPS	47.5K
SSL Throughput (ECC)	14.5 Gbps

Figure 5: BIG-IP LTM VE performance on Dell PowerEdge R620 with Intel Xeon CPU E5-2699 v4 @ 2.2GHz and Intel XL710 40G NIC—configured for SR-IOV using VMware ESXi 6.0 hypervisor (KVM for SSL ECC performance). High performance VE licensed for 16 vCPUs, running BIG-IP TMOS v13.0HF2 and later required.

Supported Hypervisors and Linux Distributions

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

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

Figure 6: 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](#).)

	VE High-Perf. SR-IOV	VE High-Perf. Paravirtualized Driver
KVM	●	●
VMware vSphere	●	●

Figure 7: High-performance, VE-supported hypervisors. Note: The high performance paravirtualized driver is used as the default driver for throughput licensed and high performance VE.

* 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.

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.

	 Lab	 25 Mbps	 200 Mbps	 1 Gbps	 3 Gbps	 5 Gbps	 10 Gbps*
Amazon Web Services** and GovCloud	●	●	●	●	●†	●	●
Amazon IC Marketplace		●	●	●			
Microsoft Azure and Government	●	●	●	●	●†		
Google Cloud Platform	●	●	●	●		●	
VMware on IBM Cloud††		●	●	●		●	

Figure 8: F5 BIG-IP VE support for the leading public cloud IaaS providers. For details and a list of validated cloud providers, visit [F5 Ready](#).

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 a high-performance VE license with no throughput limits and allows you to increase the number of vCPUs to increase performance. The maximum vCPUs supported is 24.

Meeting your applications' needs in a dynamic environment has never been easier. The **Good, Better, Best** bundle offerings from F5 provide you with the best value through flexibility to provision additional advanced modules as needed.

BIG-IP Virtual Edition Version Plus licenses provide predictable lifecycle upgrades by offering support for three major software releases.††† In addition to Version Plus bring-your-own-licenses (BYOLs), which are portable between on-premises data centers and supported public clouds, F5 offers pay-as-you-go utility licensing (hourly/monthly), annual subscription-based licenses (1-year or 3-year), and a 3-year Enterprise Licensing Agreement (ELA) for maximum flexibility and yearly budget protection.

* 10 Gbps is a BYOL and applies to non-Internet facing IP traffic only.

** Includes VMware on AWS.

† BYOL only

†† Utility (PAYG) billing only

††† Please refer to SOL15643 on askf5.com to learn more about BIG-IP VE license offerings.

BIG-IP Per-App Virtual Editions

Another option for application teams with app delivery and security needs is the F5 BIG-IP Per-App VE. Per-App VEs deliver local traffic management and advanced web application firewall for each application from development to production deployment. BIG-IP Per-App VEs deliver the same scalable, secure, and customizable services provided by traditional F5 VEs—at a price and with a license model (limited to one VIP and three Virtual Servers) appropriate for supporting individual applications.

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 90-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](#) and [Azure](#) 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:

Datasheets

[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-IP Centralized Management](#)
[Application Connector](#)
[Container Connector](#)

Web pages

[Virtual Editions](#)
[Cloud Computing](#)
[OpenStack](#)
[Cloud Solution Templates](#)
[AWS Marketplace](#)
[Azure Marketplace](#)
[Google Cloud](#)
[IBM Cloud](#)

Case studies

[Alberta Motor Association Boosts Business Agility with F5](#)
[TEPCO SYSTEMS Builds a Cloud Environment for Developers](#)
[Everbridge Manages Traffic and Security Across Global Cloud Providers and Local Data Centers](#)
[Cerner Uses Virtualized F5 Solution to Ensure Availability of Hospital Health Care Systems](#)

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](#)

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

