



BIG-IP Virtual Edition

- 2 Primary Cloud Scenarios
- 2 Private Cloud Using Software-Defined Architectures
- 3 Deploy Applications in and Across Public Cloud Environments
- 4 Application Portability Across Hybrid and Multi-Cloud Environments
- 5 Colocation Deployments with Direct Connect to Public Cloud
- 6 Integration with SDN Frameworks
- 6 Achieve Hardware-Comparable Performance with Software
- 7 Dynamic App Services for Container Environment
- 8 Automation, Orchestration, and Programmability
- 8 Centralized Management of BIG-IP VE
- 9 Specifications
- 12 F5 BIG-IP Virtual Editions: Simplified Licensing and Choices
- 13 Transitioning to BIG-IP Next Virtual Edition
- 14 Get Started Today

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 and easing your transition from hardware to software. VEs 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 VEs, simplifying cloud migrations. VEs 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 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, reduce application time to market, and 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.




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, Flexible Consumption Program (FCP), pay-as-you-go, or on a perpetual basis.





Figure 2: BIG-IP VEs 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.


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



COLOCATION DEPLOYMENTS WITH DIRECT CONNECT TO PUBLIC CLOUD

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, both for 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 the low latency, highly responsive user experiences their data center delivers.

BIG-IP Virtual Edition (VE) 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 VE 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 NICs 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 NICs.
- **Accelerated cryptographic and compression processing**—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 systems teams manage front-door ingress control and advanced application delivery and security services for container and Platform as a Service (PaaS) deployments. CIS integrates BIG-IP VE with native container environments and orchestration systems, including Kubernetes and Red Hat OpenShift. That integration enables 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:** Enables 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-IP 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-IP 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-IP'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-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: 1 vCPU, 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

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.

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](https://support.f5.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	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	100,000
SSL Throughput (ECC)	37 Gbps

SSL with Intel QAT	Maximum*
SSL RSA TPS (2K keys)	95,000
SSL throughput (RSA)	60 Gbps
SSL ECC TPS	59,000
SSL Throughput (ECC)	46 Gbps
BIG-IP DNS	Maximum*
Query responses per second	1.8 million

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.

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 16 vCPUs (for ECC) and 20 vCPUs (for RSA), running BIG-IP TMOS v14.1.0.3 and later.

Figure 9: High-Performance VE (8vCPU/16GB RAM) with enabled 5 Intel FPGA PAC N3000 SmartNIC.

BIG-IP VE FOR SMARTNICS

	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 Gbps	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](https://ask.f5.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](https://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
- **Flexible Consumption Program (FCP)**—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.

Transitioning to BIG-IP Next Virtual Edition

BIG-IP Next is the next generation BIG-IP software built to offer greater automation capabilities, scalability, and ease-of-use for organizations running applications on-premises, in the cloud, or out at the edge. At its core, it's still the same BIG-IP that F5 customers know and trust, simply designed and rearchitected for the future. Powerful declarative APIs are the foundation of BIG-IP Next's API-first design, making it faster and easier for DevOps, NetOps, and other BIG-IP-reliant teams to manage and automate their BIG-IP deployments. A completely rearchitected and modern software layer also provides the basis for significantly improved control plane scale, reduced cloud footprint, rapid instance upgrades, and much more.

BIG-IP Next will be available in a Virtual Edition form factor for deployment across public and private cloud environments beginning in 2022. The new software framework will require comparatively fewer physical resources to operate than the current BIG-IP Virtual Edition—helping lower cloud costs and energy consumption. Further optimizations within the new architecture also enable BIG-IP Next Virtual Edition to be spun up in a shorter timespan in support of more dynamic application environments.

For more information about BIG-IP Next Virtual Edition contact an [F5 sales representative](#).

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 Advanced WAF](#)
- [BIG-IP Access Policy Manager](#)
- [BIG-IP Carrier-Grade NAT](#)
- [BIG-IP Policy Enforcement Manager](#)
- [BIG-IP 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](#)
- [Journeys Migration Utility](#)

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](#)
- [Automate BIG-IP VE deployment with F5 Cloud Solution Templates](#)





F5 Distributed Cloud App Stack

Deploy and Operate a Fleet of Applications and Infrastructure

F5® Distributed Cloud App Stack is a SaaS-based offering that deploys, secures, and operates a fleet of applications across the heterogeneous infrastructure in private, public, telecommunications, and edge clouds. It can scale to a large number of clusters and locations with centralized orchestration, observability, and operations to reduce the complexity of managing a fleet of distributed clusters.

Using a distributed control plane running in F5's global network infrastructure, Distributed Cloud App Stack delivers a logically centralized cloud that can be managed using industry-standard Kubernetes APIs (application programming interfaces). This control plane removes the overhead of many individually managed Kubernetes clusters and allows customers to automate application deployment, scaling, security, and operations across their entire deployment as a “unified cloud.”

Distributed Cloud App Stack enables enterprises to leverage best-of-breed cloud services. You can deploy applications across multiple clouds to access key services via a cloud-agnostic distributed application management platform.

Key Benefits of Distributed Cloud App Stack

Operational efficiency via fleet management

Get SaaS-based lifecycle management of applications and infrastructure services across a fleet of distributed sites, along with intent-based policy and configuration.

Immersive applications at the edge or in the network

Move performance-sensitive applications (or portions of apps) to edge sites or the network edge to improve latency and performance. Leverage AI/ML applications on the edge and GPU as a Service in the network.

Security for distributed apps and data

Verify identity across distributed sites that drive policy and encrypt data (at rest and in transit), along with the innovative Blindfold capability that secures secrets and keys.

Seamless scalability

Deploy a purpose-built control plane across F5's globally distributed infrastructure that scales to a large number of application clusters.




Figure 1: Cloud-native stack for managing distributed workloads

Use cases across clouds, networks, and edge sites

Edge cloud

- Edge application management
- Deployment of AI/ML applications on fleet of edge sites

Multi-cloud

- Multi-cloud application management
- Cloud-native compute infrastructure
- GPU as a Service

Network cloud

- Network edge applications

VERTICALS

F5 Distributed Cloud Services help enterprises across numerous industry verticals distribute cloud-native applications spanning multiple public and private clouds and edge sites—without having to rebuild their software infrastructure.



E-commerce



Healthcare



Finance



Manufacturing



High-tech



Retail



Telecommunications



Automotive



Gaming



Energy

Distributed Cloud App Stack Features

Distributed Cloud App Stack delivers a complete range of services to automate application and infrastructure deployment, scaling, security, and lifecycle management across a vast number of distributed sites.

APPLICATION SERVICES

Distributed application management

Deploy and orchestrate applications across a segment or the entire fleet of clusters.

Centralized, SaaS-based management of distributed applications with a single pane of glass and rich observability. Deploy AI/ML applications across a fleet of edge sites.

Continuous delivery and verification

Continuously verify application and infrastructure health to generate alerts and automatically take actions when anomalies are detected. Integrates with existing CI/CD pipelines and development workflow.

Identity and secrets management

Uniform identity for applications across different multi-cloud and/or edge environments. Manage identity lifecycle for each application instance via automatic certificate rotation, increasing security and operational efficiency.

Container security and isolation

Container-level security services to protect and isolate applications and infrastructure from errant or malicious containers.

Observability

Single pane of glass visibility from application to infrastructure across heterogeneous edge and cloud deployments. Provides granular status of application deployments, infrastructure health, and security/connectivity performance.




Figure 2: A complete range of networking and security services to connect and secure applications

INFRASTRUCTURE SERVICES

Optimized operating system

A consistent and efficient operating system for deployment across heterogeneous cloud and edge locations, optimized to run virtual machines, containers, and Lambda workloads.

Clustering

Scale compute and storage resources seamlessly by clustering multiple nodes together to match service needs.

Distributed storage

Software-defined storage across nodes and/or clusters. Consistent volumes across heterogeneous storage types.

Multi-tenancy

Run third-party and/or multiple business lines' applications while providing complete isolation of compute, network, and storage resources. Provides the ability to run heterogeneous workloads (containers, VMs) across different namespaces.

Managed Kubernetes environment

Orchestrate and scale infrastructure resources in an automated fashion. Self-healing and progressive rollout with health monitoring for automated rollback.

Distributed infrastructure management

Manage globally distributed and/or heterogeneous infrastructure as a fleet. SaaS-based management with a single pane of glass and rich observability. Supports vGPU resource sharing across applications, namespaces, and tenants.

About F5 Distributed Cloud Services

F5 Distributed Cloud Services are SaaS-based security, networking, and application management services that can be deployed across multi-cloud, on-premises, and edge locations.

For more information on Distributed Cloud App Stack, visit f5.com. Interested in talking to an F5 sales specialist? Contact sales@f5.com today.

