



BIG-IP Cloud Edition Buying Guide

2019

INTRODUCTION

F5 BIG-IP Cloud Edition was built to help network operations teams and applications teams collaborate more effectively in the rapid delivery of secure, appropriately supported applications. BIG-IP Cloud Edition simplifies and centralizes core device and app services management functions like setup, licensing, upgrades, analytics, and scaling. It then makes it easy for operations teams to focus on defining a self-service catalog of application services that developers can access on demand via API calls. These services are defined, updated, and deployed for each individual application in contrast to the traditional, consolidated model in which a single Application Delivery Controller (ADC) supports multiple applications.

As well as bringing a new level of architectural flexibility to enterprise-class application delivery and security services, BIG-IP Cloud Edition also has several how-to-buy options. Designed to give you financial flexibility to match service flexibility, BIG-IP Cloud Edition is available with subscription, utility, and enterprise license options, as well as a traditional perpetual purchase option. While having options is a good thing, it comes with the responsibility to choose the best one for your enterprise. That's why we've put together this handy guide to help you navigate the options and find the optimal commercial model for your organization.

F5 BIG-IP Cloud Edition: A brief architectural overview

BIG-IP Cloud Edition builds a centrally managed, two-tier architecture in AWS, Microsoft Azure, or VMware with groups of specially licensed BIG-IP instances dedicated to serving individual applications. Unlike a traditional BIG-IP device, which is usually deployed to support multiple applications, a BIG-IP Per-App Virtual Edition (VE) supports only a single application. BIG-IP Per-App VE instances are deployed, provisioned, and scaled by F5 BIG-IQ Centralized Management, which provides management, visibility, and licensing services across all instances—no matter where they are.

BIG-IP CLOUD EDITION ARCHITECTURE

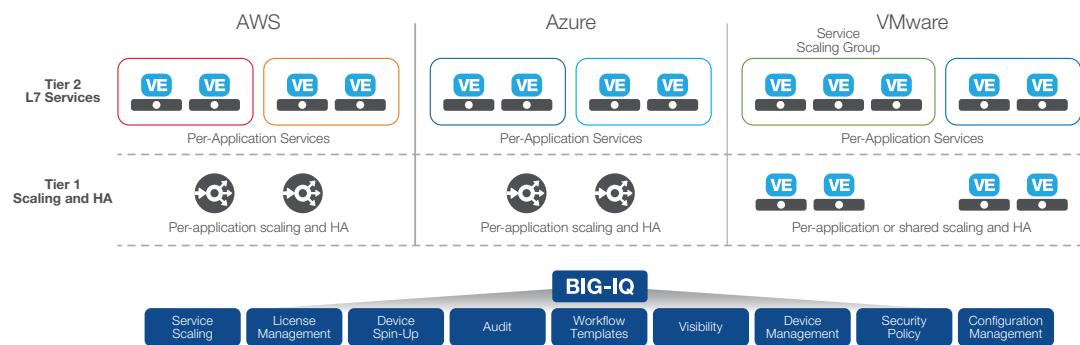


Figure 1: BIG-IP Cloud Edition builds a centrally managed, two-tier architecture in AWS, Azure, or VMware with groups of BIG-IP Per-App VE instances.

BIG-IP Per-App Virtual Edition

A BIG-IP Per-App VE is a specially licensed, full-featured BIG-IP instance that has been designed to provide dedicated services for a single application.

Each BIG-IP Per-App VE instance comes with:

- One virtual IP address
- Three virtual servers (a combination of a virtual address and a listening port)
- 25 Mbps or 200 Mbps throughput

BIG-IP LTM

F5 BIG-IP Local Traffic Manager (LTM) software delivers industry-leading application traffic management, including advanced load balancing, rate shaping, content routing, SSL management, and complete control of the application layer traffic in both directions.

BIG-IP Advanced WAF

Adding to the capabilities of BIG-IP LTM software, F5 Advanced WAF offers all the features of a traditional web application firewall, plus enhanced protection in the form of layer 7 DDoS mitigation, advanced bot detection, and API security management.

BIG-IP Service Scalars

Managing traffic with BIG-IQ Centralized Management

BIG-IQ powers the deployment, scaling, management, and licensing of the BIG-IP instances that provide the application delivery and security services that keep your applications secure, fast, and available. BIG-IQ dramatically reduces the management overhead of even the most complex application delivery configurations and makes managing multiple BIG-IP devices easy.

Cloud platforms

BIG-IP Cloud Edition is available on AWS, Azure, and in VMware. Separate management of the environments is not required, and applications can be deployed in either location using the same tools and templates. If you're deploying on AWS or Azure, remember to consider the cost of Amazon EC2 instances and Azure Compute instances respectively on which to run your BIG-IP Per-App VE instances in addition to the cost of BIG-IP licensing.

BIG-IQ DRAMATICALLY
REDUCES THE
MANAGEMENT
OVERHEAD OF EVEN
THE MOST COMPLEX
APPLICATION DELIVERY
CONFIGURATIONS.

Sizing and Capacity Planning

While you should refer to more technical documentation for your BIG-IP Cloud Edition design, it's important to understand some architectural details when choosing your solution.

How many BIG-IP Per-App VE instances will you need?

There are two critical limits on BIG-IP Per-App VE instances:

- Objects
- Throughput

Unlike a more traditional deployment, BIG-IP Per-App VE instances will generally be deployed in an all-active configuration with the tier-one traffic management device taking care of high availability and scaling. This means that you can generally get more real throughput per provisioned VE instance than in a more hardware-centric, active-standby high-availability (HA) pair. BIG-IP Per-App VE instances are available in 25 Mbps and 200 Mbps throughput licenses and are designed to scale out using service scaling groups. You will need to have a base estimate of the required throughput for each application for which you plan to manage traffic.

Next, you need to decide whether the 25 Mbps or 200 Mbps license is appropriate. If you have larger throughput requirements per application, then the 200 Mbps license makes more sense because you will have fewer overall devices. If you have smaller or more fine-grained requirements, then the 25 Mbps license is appropriate.




Figure 2: BIG-IP Per-App VE instances are available in 200 Mbps and 25 Mbps throughput licenses and are designed to scale out using service scaling groups.

You can mix and match license types within your environment, but a specific application will only be serviced by one license type.

For each app, determine:

- Total throughput required
- Likely growth
- Volatility




Figure 3: Provision your base requirement with capacity for 20 minutes of maximum expected growth so that services can flex with demand while still having some capacity to handle spikes.

When thinking of volatility, there are a few variables to consider. First, the triggers for scaling events are based on the throughput and/or CPU thresholds of the busiest device in a service scaling group, taken over a five-minute average. A new BIG-IP Cloud Edition will take about 10 minutes to become active after startup. F5 recommends provisioning your base requirement with capacity for 20 minutes of maximum expected growth. That way, services can flex with demand while still having some capacity to handle spikes.

Although sizing can be complex, the good news is that with BIG-IP Cloud Edition, you can flex each application's instances on demand, so getting things perfect—or building in huge spare capacity— isn't essential.

Application	Throughput	Volatility	25M	200M	License
Intranet	100	Low	5	0	LTM
Web Shop Dev	5	Low	1	0	LTM
Web Shop Test	20	Med	2	0	LTM + WAF
Marketing Platform	800	High	0	6	LTM + WAF

Table 1: A possible scenario with several applications to service, each having different throughput requirements and volatility.

Auto-scale options for service scaling groups

Service scaling capability is implemented differently between environments. In public cloud platforms like AWS, scaling in and out is handled by a simple, built-in load balancer such as AWS ELB.

In VMware scaling, layer 4 DDoS and firewall functions are provided by special BIG-IP VE instances. These VE instances are not configured to provide layer 7 application delivery and security services, but simply distribute traffic to the per-app VE instances and provide network layer access control and DDoS mitigation.

The auto-scale instances are designed to be high throughput, low complexity, and shareable between multiple applications.

Existing or new BIG-IP hardware platforms can also be used as auto-scale instances without the purchase of additional licenses, assuming the existing equipment has an active LTM license.

Sizing and planning for BIG-IQ

BIG-IQ high availability and backup

Since BIG-IP Cloud Edition essentially routes all control plane activities through the BIG-IQ management layer—BIG-IQ handles real-time monitoring and scale-in and -out events and manages subscription licenses—it becomes a critical part of the delivery system and therefore will usually need to be highly available.

Your planning should, therefore, include an active-standby BIG-IQ pair, with the appropriate license for the number of BIG-IP instances under management.

BIG-IQ licensing

You will need to plan to license the BIG-IQ component for the maximum number of BIG-IP instances under management. This will depend on what combination of BIG-IP instance sizes and types you are planning to run.

At the end of this exercise you should have established the following key requirements:

	BIG-IP Per-App VE 25 LTM	BIG-IP Per- App VE 25 LTM + WAF	BIG-IP Per- App VE 200 LTM	BIG-IP Per- App VE 200 LTM + WAF	BIG-IQ Devices Licensed
Count					

Table 2: Key requirements for your environment.

Purchase Options

After you have established your needs, the next step is working out the commercial model that suits you best. F5 offers BIG-IP Cloud Edition with three purchase options:

- Subscription
- Perpetual licensing
- Enterprise License Agreement (ELA)

In all purchase models, F5 BIG-IQ License Manager (LM) allows you on-demand, self-service licensing of BIG-IP instances, enabling you to build, scale, and destroy application delivery services dynamically without contacting F5 for licensing of BIG-IP instances.

Subscription licensing

Subscription Licensing Offer for BIG-IP VE is an auto-renewal agreement for BIG-IP VE licenses, available in one-year, two-year, and three-year periods. Within the agreement, organizations choose local BIG-IP VE licenses at subscription initiation. The initial order of subscription licenses is populated in BIG-IQ LM and is available for immediate deployment. Subscription Licensing Offer for BIG-IP VE enables IT departments to self-manage their lifecycles. License instantiation or revocation can be done via BIG-IQ LM, either directly or via REST API. Net-new licenses from the BIG-IP VE subscription list can be instantiated at any time. Activity reports are provided monthly to F5 via BIG-IQ API upload or email. F5 Premium (24x7) support and software updates are included in the subscription.

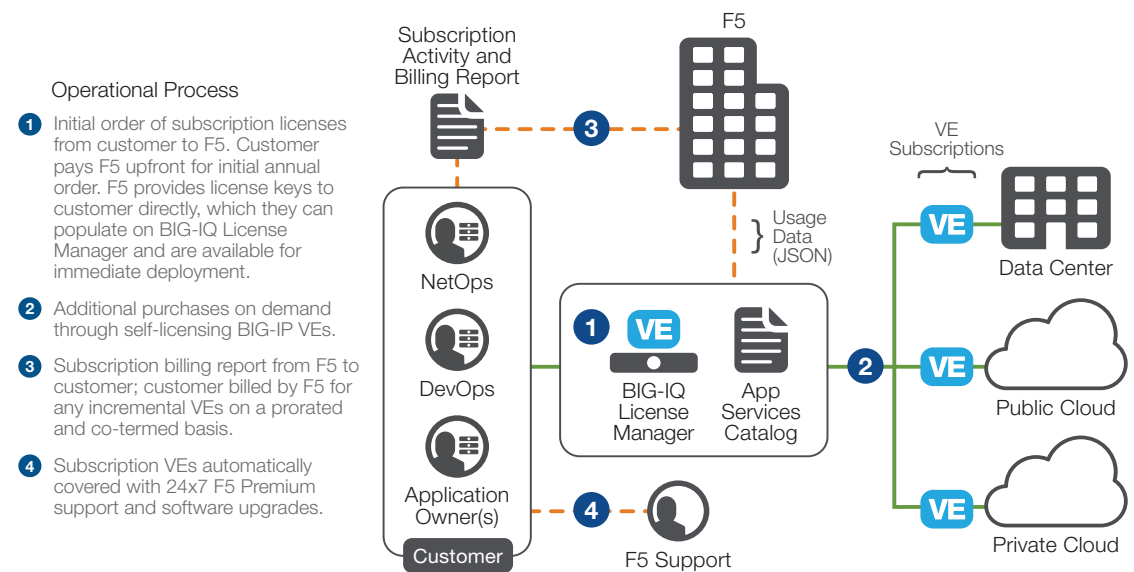


Figure 4: Subscription Licensing Offer for BIG-IP VE is an auto-renewal agreement for BIG-IP VE licenses, available in one-year, two-year, and three-year periods.

Enterprise Licensing Agreement

The F5 ELA is an annual spend commitment that allows you to deploy services from the F5 catalog wherever and whenever you need. This helps you dynamically create or scale IT infrastructure to meet business needs. With the three-year license subscription, you can commit to a minimum annual spend, over a three-year period, for application and security services—with room to grow. The annual fee is adjusted for subsequent years based on the prior year’s usage. If you exceed your yearly budget, you won’t be penalized with a retroactive charge-back.

Perpetual Licenses

Perpetual pool licenses are also available for BIG-IP Cloud Edition. Pool licenses are purchased for a fixed number of devices but are not permanently tied to a specific device. Pool licenses are managed by BIG-IQ and can be used for scaling via the service scaling groups, and can be used and recovered in the same manner as other licenses. The significant difference is that pool license numbers are fixed and cannot be scaled dynamically without the purchase of further pool license packs. Support is not automatically included but is available to purchase as an additional item.

	“Instant” access to additional licenses	Fixed yearly costs	Support included	Initial \$ commitment
Subscription	•		•	\$\$
ELA	•	•	•	\$\$\$
Perpetual		•		\$

Table 3: Comparison of subscription purchase options.

A note on BIG-IQ licensing: BIG-IQ is currently available only as a perpetual license; however, it is capable of managing both the BIG-IP Per-App VE solution and your existing hardware or software BIG-IP instances.

The architectural flexibility of BIG-IP Cloud Edition can help network operations teams and applications teams rapidly deliver secure, appropriately supported applications. Designed to give you financial flexibility to match service flexibility, BIG-IP Cloud Edition is available with subscription, utility, and enterprise license options, as well as traditional perpetual purchase options.

To learn more about the enterprise-class application delivery and security services offered by BIG-IP Cloud Edition—and the purchase options that suit you best—contact your sales rep.



US Headquarters: 401 Elliott Ave W, Seattle, WA 98119 | 888-882-4447 // Americas: info@f5.com // Asia-Pacific: apacinfo@f5.com // Europe/Middle East/Africa: emeainfo@f5.com // Japan: f5j-info@f5.com

©2019 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of the respective owners with no endorsement or affiliation, expressed or implied, claimed by F5. GUIDE-CLOUD-297422327 | 12.18