

By Gregory Coward, Solution Architect, [F5 Networks](https://www.f5.com/) (<https://www.f5.com/>)

This document shows how to install and configure the [F5 BIG-IP](https://www.f5.com/products/big-ip-services) (<https://www.f5.com/products/big-ip-services>) Application Delivery Controller (ADC) before you [integrate the ADC with GKE On-Prem](/gke-on-prem/docs/how-to/installation/install) (</gke-on-prem/docs/how-to/installation/install>). If you're interested in installing F5 BIG-IP ADC using manual load-balancing mode on GKE On-Prem, see [Installing F5 BIG-IP ADC for GKE On-Prem using manual load balancing](/solutions/partners/installing-f5-big-ip-adc-for-gke-on-prem-using-manual-load-balancing) (</solutions/partners/installing-f5-big-ip-adc-for-gke-on-prem-using-manual-load-balancing>).

F5 is a leading provider of ADC services. The F5 BIG-IP platform provides various services to help you enhance the security, availability, and performance of your apps. These services include L7 load balancing, network firewalling, [web application firewalling \(WAF\)](https://wikipedia.org/wiki/Web_application_firewall).

([https://wikipedia.org/wiki/Web\\_application\\_firewall](https://wikipedia.org/wiki/Web_application_firewall)), and DNS services. For GKE On-Prem, the BIG-IP provides external access and L3/4 load-balancing services.

Among its various services and features, BIG-IP provides [container ingress services \(CIS\)](https://clouddocs.f5.com/containers/v2/) (<https://clouddocs.f5.com/containers/v2/>). CIS provides platform-native integrations for BIG-IP devices from platform as a service (PaaS) providers like Kubernetes. This integration makes it possible to dynamically allocate BIG-IP L4-L7 services in container orchestration environments.

[Anthos](/anthos/) (</anthos/>) uses a version of CIS to automatically provision L4 load-balancing services on the BIG-IP platform while providing external access to apps deployed on GKE On-Prem clusters.

After you configure BIG-IP, you can integrate with GKE On-Prem for ingress as part of the [default installation](/gke-on-prem/docs/how-to/installation/requirements) (</gke-on-prem/docs/how-to/installation/requirements>).

**Note:** Currently, when deployed as part of the [default GKE On-Prem installation](/gke-on-prem/docs/how-to/installation/requirements) (</gke-on-prem/docs/how-to/installation/requirements>), only a single CIS controller can be deployed. To use more controllers, consider deploying the BIG-IP using [manual load-balancing mode](/gke-on-prem/docs/how-to/installation/manual-lb) (</gke-on-prem/docs/how-to/installation/manual-lb>).

- Learn about BIG-IP high availability.
- Set up BIG-IP.
- Configure BIG-IP before you deploy it to GKE On-Prem.

This tutorial uses the following billable components of Google Cloud:

- [Anthos \(/anthos/pricing\)](/anthos/pricing)

To generate a cost estimate based on your projected usage, use the [pricing calculator \(/products/calculator\)](/products/calculator). New Google Cloud users might be eligible for a [free trial \(/free-trial\)](/free-trial).

1. GKE On-Prem automatically installs F5 BIG-IP container ingress services (CIS). Each version of GKE On-Prem relies on a [different version of CIS \(/gke-on-prem/docs/how-to/load-balance-basic\)](/gke-on-prem/docs/how-to/load-balance-basic). Ensure that the F5 BIG-IP load balancer version supports the F5 BIG-IP CIS version that comes with GKE On-Prem by consulting with the [F5 BIG-IP Controller/Platform compatibility matrix \(https://clouddocs.f5.com/containers/v2/releases\\_and\\_versioning.html#connector-compatibility\)](https://clouddocs.f5.com/containers/v2/releases_and_versioning.html#connector-compatibility).
2. Obtain an F5 BIG-IP Application Delivery Controller and license. The F5 BIG-IP ADC is available in various hardware platforms and virtual editions. Regardless of the platform you use, the solution is supported, and the following configuration process is applicable.

There are three types of licenses for F5 BIG-IP.

License type	Production	Evaluation	Demonstration
Expiration date	No	Yes, 45 days	Yes, 30 days

License type	Production	Evaluation	Demonstration
Throughput limitations	Up to 40 Gbps on GKE On-Prem	25 Mbps - 10 Gbps	1 Mbps
Use case	Production	Proof of concept, demonstration	Proof of concept, demonstration

- If you have a production license, you can use that license.
- For the purpose of this tutorial, you can [request a free trial license key](https://www.f5.com/trials) (<https://www.f5.com/trials>).
- If your throughput requirements are greater than 1 Mbps provided by the trial license, you can request an evaluation license from F5:
  - [Request a BIG-IP evaluation license](https://www.google.com/url?sa=D&q=https%3A%2F%2Ff5.com%2Fproducts%2Ftrials%2Fbig-ip-evaluation-request%2Fpopup%2Ftrue%3Fsbtitle%3DBIG-IP%2BEvaluation%2BTrial%2BRequest%26utm_source%3Df5com%26utm_medium%3Dweb%26utm_campaign%3Dhtb-trial-page) ([https://www.google.com/url?sa=D&q=https%3A%2F%2Ff5.com%2Fproducts%2Ftrials%2Fbig-ip-evaluation-request%2Fpopup%2Ftrue%3Fsbtitle%3DBIG-IP%2BEvaluation%2BTrial%2BRequest%26utm\\_source%3Df5com%26utm\\_medium%3Dweb%26utm\\_campaign%3Dhtb-trial-page](https://www.google.com/url?sa=D&q=https%3A%2F%2Ff5.com%2Fproducts%2Ftrials%2Fbig-ip-evaluation-request%2Fpopup%2Ftrue%3Fsbtitle%3DBIG-IP%2BEvaluation%2BTrial%2BRequest%26utm_source%3Df5com%26utm_medium%3Dweb%26utm_campaign%3Dhtb-trial-page))
  - If you use Anthos, [Send email to F5 to request a BIG-IP evaluation license](mailto:googleteam@f5.com) (<mailto:googleteam@f5.com>).
- If the BIG-IP system contains an evaluation or demonstration license, the BIG-IP system stops processing traffic when the license expires.

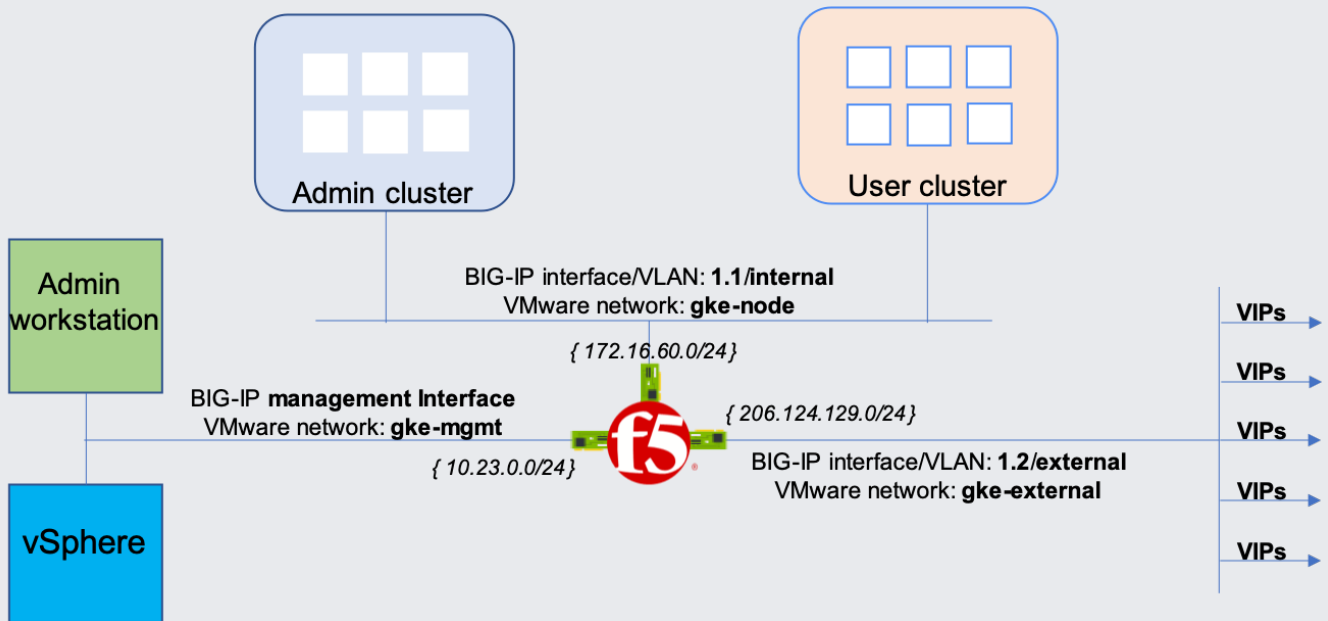
3. [Activate a license key for BIG-IP](https://www.f5.com/products/get-f5) (<https://www.f5.com/products/get-f5>).

4. Make sure your environment meets the following minimum system requirements:

- 8 vCPUs that aren't shared between other hosts on that system
- 16 GB memory that isn't shared between other hosts on that system

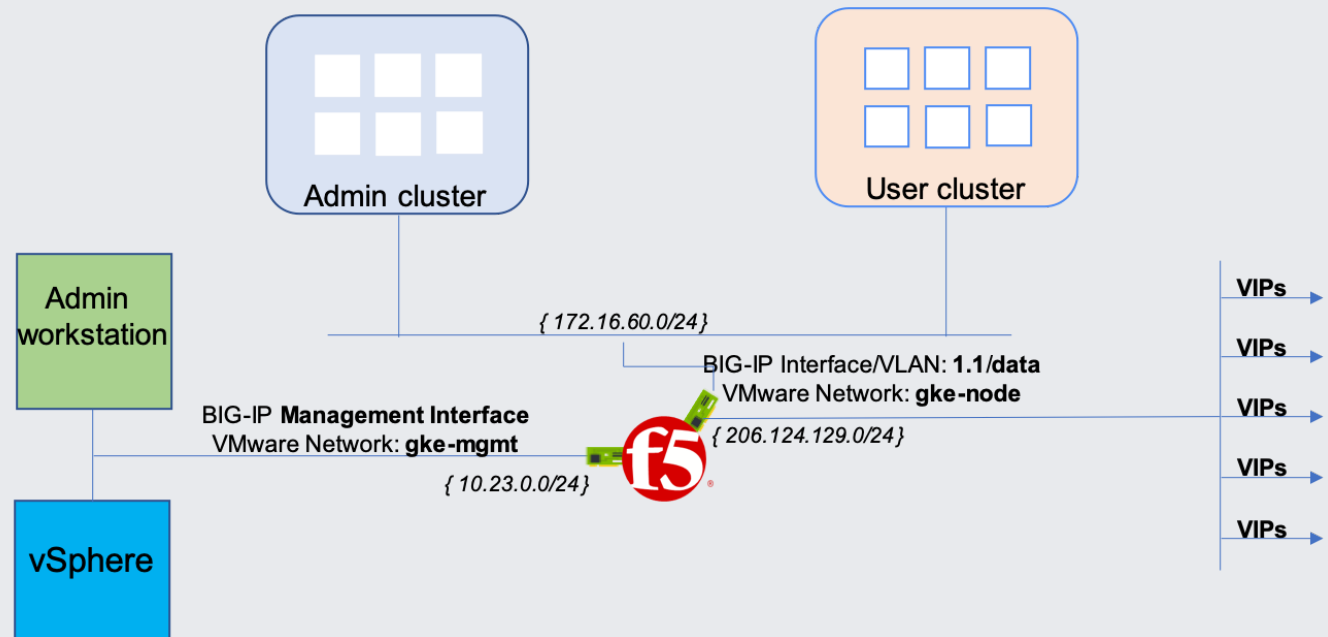
There are two common scenarios to deploy BIG-IP ADC to GKE On-Prem clusters. Because the BIG-IP acts as a proxy for external access to the clusters, it's common to deploy a BIG-IP with

three or more interfaces, as illustrated in the following diagram.



In the preceding diagram, separate interfaces serve internal *private*- and external *public*-facing traffic independently. This architecture provides better visibility for monitoring and troubleshooting, and increased throughput.

Although uncommon, you can deploy BIG-IP in a two-armed mode, where only one interface serves data plane traffic, as illustrated in the following diagram.



While the preceding configuration isn't considered a best practice, if you're integrating into an existing environment with a pre-defined network architecture, you might need this type of

configuration.

To ensure availability of your applications, it is considered a best practice for the BIG-IP to be deployed in either an active-standby or active-active device service cluster. The BIG-IP uses device service clustering

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-system-device-service-clustering-administration-13-1-0/1.html#guid-8766a4e7-4d1f-4bf8-afbe-8037ae89fd08](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-system-device-service-clustering-administration-13-1-0/1.html#guid-8766a4e7-4d1f-4bf8-afbe-8037ae89fd08))

(DCS) to synchronize configuration data between cluster group members and to provide automatic failover if the active device fails.

You can configure the BIG-IP for DSC as part of the initial setup

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/big-ip-system-initial-configuration-14-0-0/01.html#guid-5cfb2400-eea3-4ba6-858d-8dd8db68b2cb](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/big-ip-system-initial-configuration-14-0-0/01.html#guid-5cfb2400-eea3-4ba6-858d-8dd8db68b2cb))

or anytime thereafter by running the Config Sync/HA utility

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-system-device-service-clustering-administration-13-1-0/1.html#guid-8766a4e7-4d1f-4bf8-afbe-8037ae89fd08](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-system-device-service-clustering-administration-13-1-0/1.html#guid-8766a4e7-4d1f-4bf8-afbe-8037ae89fd08))

. With that said, there are important considerations for the DSC and how the BIG-IP interacts with GKE On-Prem.

As noted earlier, the BIG-IP uses F5 Container Integration Services (CIS) to integrate with GKE On-Prem. CIS deploys a lightweight controller into the environment to monitor the partition it manages for configuration changes. If CIS discovers changes, the connector reapplies its own configuration to the BIG-IP system.

F5 doesn't typically recommend making configuration changes to objects in any partition managed by F5 CIS by any other means, including by syncing configuration from another device or service group. However, because GKE On-Prem uses NodePort (</gke-on-prem/docs/concepts/networking>) mode only, you can deploy the BIG-IP systems in an active-standby cluster and take advantage of F5's native configuration auto-sync capabilities.

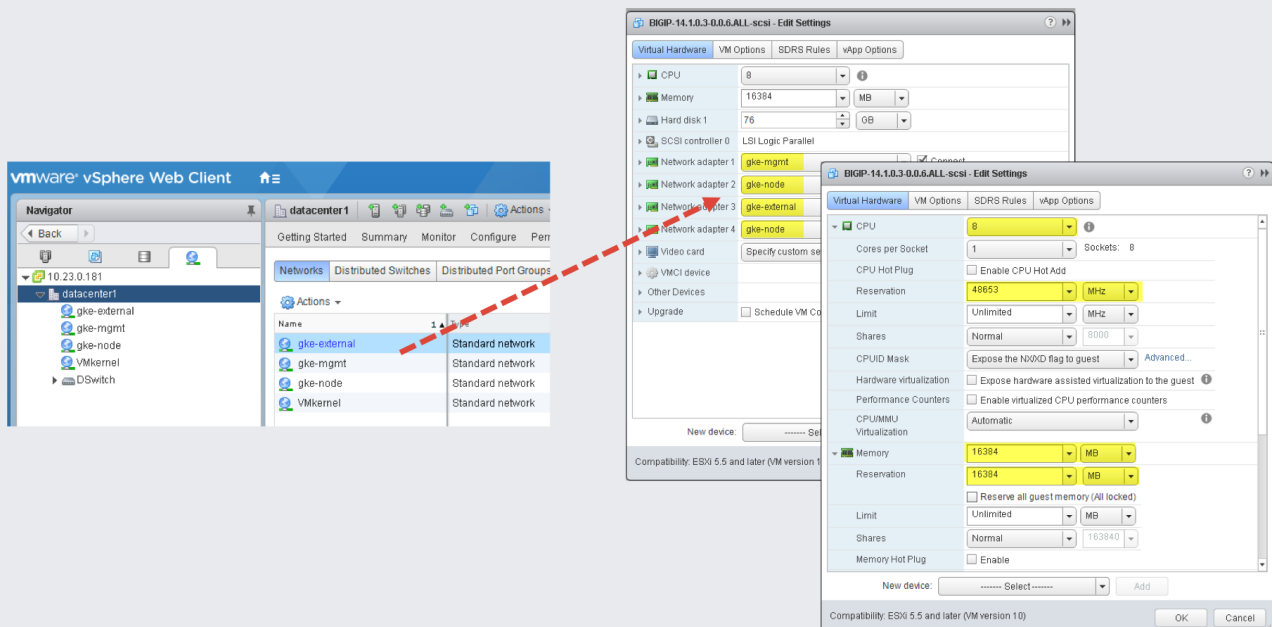
BIG-IP systems use two types of IP addresses, floating and non-floating. A non-floating IP address always remains with the BIG-IP device where it was defined. Floating IP addresses are used and are owned by whichever BIG-IP is the active device. In the event of a failover, these addresses float from the formerly active to the newly active BIG-IP.

While F5 typically recommends deploying one controller instance per BIG-IP device to further enhance availability, you can deploy a single CIS instance pointing to the BIG-IP cluster's floating self-IP address. When changes to the GKE On-Prem infrastructure occur and are updated on the active BIG-IP by the controller), the standby device receives updates automatically from the configuration synchronization.

## 1. Follow the instructions to set up BIG-IP virtual edition deployed on VMware ESXi 6.5

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-ve-setup-vmware-esxi-13-1-0/3.html](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-ve-setup-vmware-esxi-13-1-0/3.html))

The OVF template requires configuring four interfaces. The fourth interface is designated for HA heartbeat traffic between BIG-IP pairs. For this three-arm deployment, assign the internal network `gke-node`.



## 2. After the VM boots, use the F5 BIG-IP's setup utility.

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/big-ip-system-initial-configuration-14-0-0/01.html#guid-5cfb2400-eea3-4ba6-858d-8dd8db68b2cb](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/big-ip-system-initial-configuration-14-0-0/01.html#guid-5cfb2400-eea3-4ba6-858d-8dd8db68b2cb))

for initial configuration. The setup utility walks you through the following configuration tasks:

- a. From a network accessible workstation on which you configured the `gke-mgm` interface, go to the following URL `https://management_IP_address` where `management_IP_address` is the address you configured for your device.
  - b. When prompted, enter the default username as `admin` and the password as `admin`.
  - c. Click **Log in**.
3. To install a license, in the **Base Registration Key** field, enter your key. The type of license dictates the BIG-IP's services and bandwidth limits.
  4. To enhance performance when working with GKE clusters, set the **Management (MGMT)** plane provisioning to **Large**.
  5. To provide L3/4 load balancing to the GKE On-Prem environment, set the **Local Traffic (LTM)** module to **Nominal**.

Module	Provisioning	License Status
<input checked="" type="checkbox"/> Management (MGMT)	Large	N/A
<input type="checkbox"/> Carrier Grade NAT (CGNAT)	Disabled	Licensed
<input checked="" type="checkbox"/> Local Traffic (LTM)	<input checked="" type="checkbox"/> Nominal	Licensed

6. In the **Host and User Information** window, you provide the hostname, FQDN of the BIG-IP system, and update the `admin` and `root` account passwords.
7. In the **Networking** window, you walk through configuring the BIG-IP's basic networking. The utility creates the internal, `gke-node` and external, `gke-external1` interfaces, VLANs, and self-IP addresses.

Network >> VLANs: VLAN List

Name	Application	Tag	Untagged Interfaces	Tagged Interfaces	Partition / Path
<input type="checkbox"/> internal		4094	1.1		Common
<input type="checkbox"/> external		4093	1.2		Common

The fourth interface deployed by VMware is left unconfigured.

Network >> Self IPs

Name	Application	IP Address	Netmask	VLAN / Tunnel	Traffic Group	Partition / Path
<input type="checkbox"/> 172.16.60.10		172.16.60.10	255.255.255.0	internal	traffic-group-local-only	Common
<input type="checkbox"/> 208.124.129.25		208.124.129.25	255.255.255.0	external	traffic-group-local-only	Common

After the Setup utility completes, you have a functioning BIG-IP with a management plane interface attached to the `gke-mgmt` VMware network and two data plane interfaces attached to VMware networks, `gke-node` and `gke-external`.

Before you deploy GKE On-Prem, more configuration of BIG-IP is required.

- Create an administrative partition

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-user-account-administration-12-0-0/3.html](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-12-0-0/3.html))

for each admin and user cluster you intend to expose and access. Initially, you define two partitions: one for the admin cluster, and one for the first user cluster. Don't use cluster partitions for anything else. Each of the clusters must have a partition that is for the sole use of that cluster.

The existing Administrator role provides enough permissions for use with GKE On-Prem. For more information, see User roles

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-user-account-administration-11-6-0/3.html#taskid](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-11-6-0/3.html#taskid))

. You can also learn how to create additional users

([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/bigip-user-account-administration-13-1-0/4.html](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-user-account-administration-13-1-0/4.html))

1. Use an SSH utility, such as Putty (<https://www.putty.org/>) or Cygwin (<https://www.cygwin.com/>) to establish an SSH connection to the BIG-IP's management interface.
2. Allocate more memory to the `rest.javad` process:



You don't need access to tmsh ( [Traffic Management Shell](https://clouddocs.f5.com/cli/tmsh-reference/latest/general/tmsh.html)

(<https://clouddocs.f5.com/cli/tmsh-reference/latest/general/tmsh.html>)) to run this command.

The `restjavad` process provides control-plane access to the BIG-IP system by using an HTTP REST API.

3. Restart the `restjavad` process:

4. If you're using BIG-IP 13.1 or a higher version, you can skip this step. Allocate more memory to Tomcat:

5. Restart the Tomcat service:

As noted earlier, CIS is deployed to the GKE On-Prem environment to monitor configuration changes and update the BIG-IP system, effectively overwriting the relevant user partition. As a result of this integration, manual modifications made to a virtual server within a CIS-managed cluster are overwritten.

While not typically recommended, there can be situations, such as using custom protocol profiles or persistence settings, where you might use BIG-IP system functionality that isn't natively supported by the CIS controller. To accomplish this, you can whitelist a CIS-configured virtual server. Whitelisting tells the Controller to merge its configuration into the existing virtual server instead of overwriting it.

To whitelist a particular CIS-managed virtual server, complete the following steps:

1. Establish an SSH connection to the BIG-IP's management interface.

2. Access the TMOS command-line shell:

3. Change to your *partition-name* user partition.

4. Whitelist a virtual server:

Replace the following:

- *virtual-server-name*: the name of the virtual server you want to whitelist.

At this point, your BIG-IP is configured and ready to provide L3/4 load-balancing services for your GKE On-Prem environment. You can now proceed with [deploying GKE On-Prem](#) (/gke-on-prem/docs/how-to/installation/getting-started).

- Try out other Google Cloud features for yourself. Have a look at our [tutorials](#) (/docs/tutorials).
- Learn more about F5 [BIG-IP Application Services](#) (<https://www.f5.com/products/big-ip-services>).
- Learn more about F5 [Container Ingress Services](#) (<https://clouddocs.f5.com/containers/v2/kubernetes/>).
- Learn more about BIG-IP configurations and capabilities:
  - [Certificate profiles](#) ([https://techdocs.f5.com/kb/en-us/products/big-ip\\_ltm/manuals/product/ltm-profiles-reference-13-1-0/6.html#guid-cc146765-9237-474b-9faf-0e18a208cb23](https://techdocs.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/ltm-profiles-reference-13-1-0/6.html#guid-cc146765-9237-474b-9faf-0e18a208cb23))

- WAF policies (<https://support.f5.com/csp/article/K85426947>)