ORACLE

Deploying F5 BIG-IP Virtual Edition (VE) on Oracle Compute Cloud@Customer or Private Cloud Appliance

Version 1.0

Public

Purpose statement

This solution paper provides comprehensive guidance for deploying and configuring F5 BIG-IP Virtual Edition (VE) on Oracle Compute Cloud@Customer (C3) and Private Cloud Appliance (PCA) platforms. It addresses the growing enterprise need for advanced application delivery, load balancing, and security services within Oracle's edge cloud infrastructure.

Organizations leveraging C3 or PCA can utilize F5 BIG-IP VE to deliver enterprise-grade traffic management, SSL/TLS offloading, and application security while maintaining data sovereignty and regulatory compliance requirements. This document provides IT architects, network engineers, and infrastructure teams with step-by-step deployment procedures, network architecture best practices, and configuration guidelines necessary to successfully implement F5 BIG-IP VE in Oracle edge cloud environments.

Target Audience: Solution architects, network engineers, infrastructure administrators, and IT professionals responsible for application delivery and security on Oracle edge cloud platforms.

Disclaimer

ORACLE

Table of contents

Introduction

F5 Networks and Oracle Cloud Infrastructure deliver integrated Application Delivery Networking through BIG-IP Virtual Edition (VE), a comprehensive platform engineered for speed, availability, and security of business-critical applications. This partnership enables organizations to deploy application services rapidly and securely on Oracle's Edge Cloud platforms: Compute Cloud@Customer (C3) and Private Cloud Appliance (PCA).

This solution paper details the deployment and configuration process for F5 BIG-IP VE on Compute Cloud@Customer and Private Cloud Appliance environments.

Note: This content is provided for informational purposes and self-supported guidance only. Consultancy or other assistance related to the content is not covered under the Oracle Support contract or associated service requests. If you have questions or additional needs, then please reach out to your Oracle Sales contact directly.

Solution Architecture

The following architecture demonstrates a multi-NIC F5 BIG-IP VE deployment on Compute Cloud@Customer or Private Cloud Appliance. This design implements a three-tier subnet model that enables secure, scalable application delivery from Internet-facing clients to backend infrastructure while maintaining strict network segmentation and administrative isolation.

Traffic Flow

User traffic originates from the Internet and is directed to the public IP endpoint (10.122.56.208) mapped to the F5 BIG-IP VE external interface. The F5 virtual server (10.0.1.2) processes incoming requests and load-balances traffic across backend application servers hosted in the Internal subnet. This architecture provides centralized traffic management, SSL/TLS termination, and application-layer security before requests reach internal resources.

IP Address Configuration:

- **Management IP:** 10.0.0.8 (Private) – BIG-IP administrative access.
- **External Virtual Server IP:** 10.0.1.2 (Private) / 10.122.56.208 (Public) – Application traffic endpoint.
- **Internal Self-IP:** 10.0.2.2 – Backend server connectivity.
- **Self-IP Addresses:** Configured for each subnet interface to enable routing between network segments.

Architecture Benefits:

- **Management Isolation:** Dedicated private IP (10.0.0.8) provides secure out-of-band administrative access via SSH (TCP/22) and HTTPS (TCP/443), completely separated from production traffic flows.
- **Service Consolidation:** Single public IP (10.122.56.208) serves multiple applications through port-based routing and virtual server configuration, supporting HTTP (TCP/80), HTTPS (TCP/443), and custom application ports.
- **Cost Efficiency:** Minimal public IP consumption (2 IPs total) while supporting multiple backend applications and services.
- **Scalability:** Architecture supports 20+ additional private IPs for future virtual servers and application expansion.
- **Backend Infrastructure:** Internal subnet (10.0.2.0/24) hosts multi-tier application architecture including:
    - **Web Tier:** Three web servers (10.0.2.3, 10.0.2.4, 10.0.2.5) behind internal load balancing
    - **Application Tier:** Three application servers (App 1-3) for business logic processing
    - **Database Tier:** Three-node database cluster (DB1, DB2, DB3) with cluster endpoint for high availability

Three-Subnet Architecture

Management Subnet (Public Regional) – 10.0.0.0/24

This subnet provides administrative access to the BIG-IP Configuration utility and can be configured as either public or private depending on organizational security requirements.

- **Primary Function:** Administrative access to BIG-IP configuration and monitoring
- **Access Methods:** SSH (TCP/22) for CLI management, HTTPS (TCP/8443) for web-based Configuration utility
- **F5 MGMT VNIC:** Private IP 10.0.0.8, Public IP mapping disabled for production F5 instance
- **Bastion Host:** Windows 2025 Bastion Host (Private IP: 10.0.0.9, Public IP: 10.122.57.61) provides secure jump-box access for administrators connecting from external networks
- **Security Model:** Out-of-band management network isolated from production traffic

External Subnet (Public Regional) – 10.0.1.0/24

This subnet hosts the F5 virtual server configuration that accepts Internet-facing traffic and routes requests to backend applications.

- **Primary Function:** Internet traffic ingress and application delivery
- **F5 External VNIC:** Private IP 10.0.1.2, Public IP 10.122.56.208
- **Virtual Server Configuration:** Virtual Server IP 10.0.1.2 configured to accept HTTP/HTTPS traffic and route to backend pool
- **Backend Pool:** References web servers in Internal subnet (10.0.2.3-5)
- **Public Access:** Single public endpoint (10.122.56.208) accessible via https://my-site.domain.com for production services

Internal Subnet (Private Regional) – 10.0.2.0/24

This subnet hosts the complete application infrastructure stack with no direct Internet connectivity, ensuring backend resources are protected from external threats.

- **Primary Function:** Backend application and data tier hosting
- **F5 Internal VNIC:** Private IP/Self-IP 10.0.2.2 provides routing between F5 and internal resources
- **Web Tier:** Three web servers (10.0.2.3, 10.0.2.4, 10.0.2.5) receive traffic from F5 virtual server
- **Application Tier:** Three application servers (App 1-3) process business logic
- **Database Tier:** Three-node database cluster (DB1, DB2, DB3) with dedicated cluster endpoint and internal load balancing
- **Security Posture:** No direct Internet access; all inbound traffic flows through F5 BIG-IP VE security controls

Network Provisioning and Deployment

Pre-Deployment

All network components, VCN (10.0.0.0/16), subnets, IP addresses, route tables, security lists, and Internet gateway, are provisioned on Compute Cloud@Customer or Private Cloud Appliance before F5 BIG-IP VE instance creation.

VNIC Attachment Sequence:

1. **Primary VNIC (Management):** Automatically created and attached to Management subnet (10.0.0.0/24) during F5 BIG-IP VE instance deployment
2. **Secondary VNICs:** Two additional VNICs must be manually created and attached post-deployment:
    - **External VNIC:** Attached to External subnet (10.0.1.0/24) for Internet-facing virtual servers
    - **Internal VNIC:** Attached to Internal subnet (10.0.2.0/24) for backend server connectivity

Best Practice – Bastion Host Deployment:

Deploy a dedicated bastion host (Windows 2025 or Linux) on the management subnet with a public IP address (10.122.57.61). This enables secure administrative access to the F5 BIG-IP VE management interface from clients operating outside the Compute Cloud@Customer or Private Cloud Appliance environment. The bastion host acts as a secure jump-box, reducing the attack surface by eliminating direct public access to the F5 management interface.

Architecture Diagram

The complete solution architecture, including network topology, IP addressing, traffic flows, and component relationships, is illustrated in the diagram below.

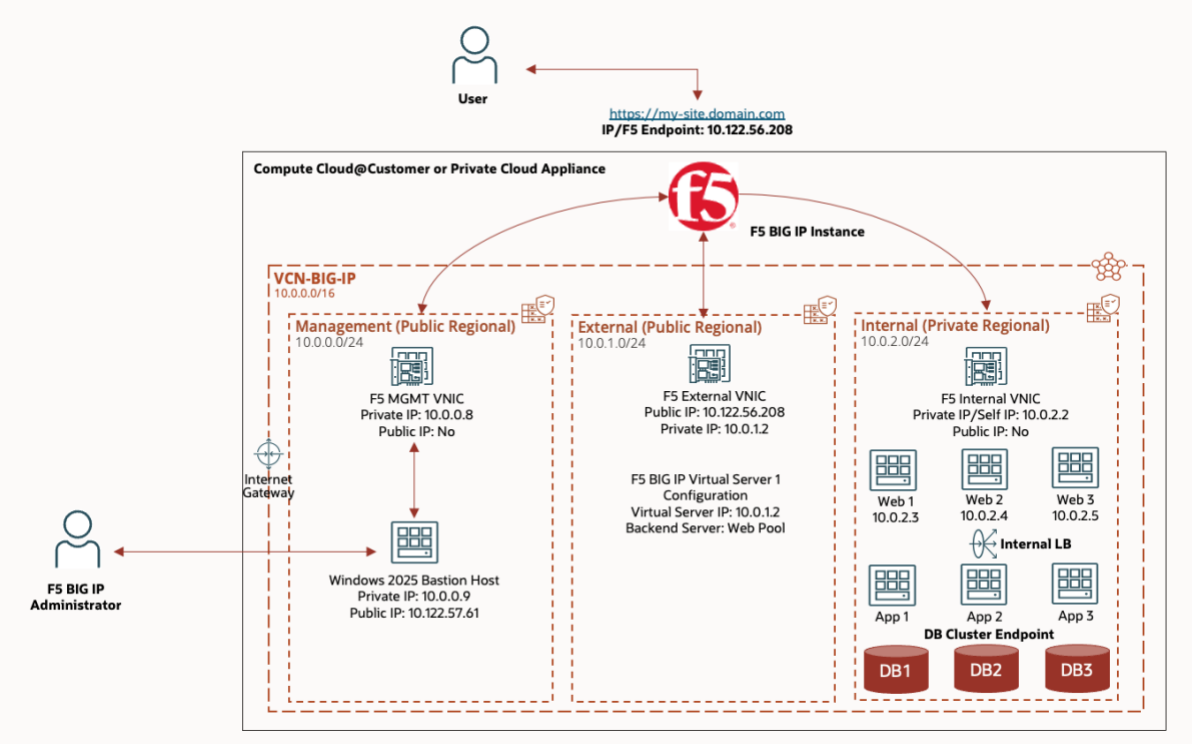

Figure 1. F5 BIG IP Architecture on Compute Cloud@Customer or Private Cloud Appliance

Network Infrastructure Prerequisites

Before deploying F5 BIG-IP on Compute Cloud@Customer or Private Cloud Appliance, the underlying network infrastructure must be configured. The following components are required for F5 BIG-IP operation:

- Virtual Cloud Network (VCN)
- Management, External, and Internal subnets
- Virtual Network Interface Cards (VNICs)

The following sections provide step-by-step instructions for configuring the network infrastructure.

Virtual Cloud Network (VCN)

Create a Virtual Cloud Network on Compute Cloud@Customer or Private Cloud Appliance. If a suitable VCN does not already exist, follow these steps:

1. Navigate to **Networking** > **Virtual Cloud Networks** in the C3 or PCA management interface
2. Click **Create Virtual Cloud Network**
3. Configure the following parameters:
    - **Name:** Provide a descriptive name for the VCN

- **Compartment:** Select the appropriate compartment
- **CIDR Block:** Enter 10.0.0.0/16 (or your preferred address space)

4. Leave all other settings at their default values
5. Click **Create Virtual Cloud Network**

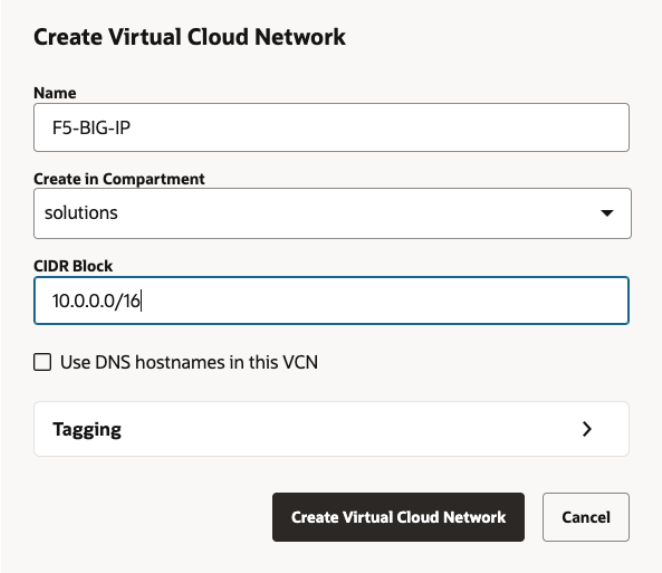

Figure 2. Creating VCN on Compute Cloud@Customer or Private Cloud Appliance

Subnets

This architecture requires three subnets as outlined in the table below. To create the subnets on Compute Cloud@Customer or Private Cloud Appliance, follow the steps listed below:

1. Navigate to **Networking** > **Virtual Cloud Networks**
2. Click on your VCN name
3. Select **Subnets** from the left menu
4. Click **Create Subnet**
5. Configure each subnet according to the specifications below
6. Leave all other settings at their default values
7. Click **Create**

Repeat these steps for each of the three subnets.

NOTE: The Management and External Subnets will need to be configured as public subnet.

Subnet Name	CIDR Block	Public/Private Subnet	Description
Management	10.0.0.0/24	Public Subnet	This subnet is used to access the BIG-IP Configuration utility for configuring your BIG-IP VE instance.
External	10.0.1.0/24	Public Subnet	This subnet is used to deploy the F5 virtual servers to accept internet traffic.
Internal	10.0.2.0/24	Private Subnet	This subnet is used for deployment of the internal servers, such as: Web, Application, Database servers.

Security lists must be configured to permit traffic on ports required for F5 BIG-IP VE management interface, application and operation. Listed below are the IP protocol, source and destination port range needed for F5 BIG-IP VE instance.

Configuration Steps:

- Navigate to **Networking** > **Virtual Cloud Networks**
- Select your VCN
- Click **Security Lists** in the left navigation menu
- Select the default security list for your VCN
- Click **Ingress Rules** in the left menu
- Select all rules using the top checkbox, then click **Edit**
- Configure the following ports for F5 BIG-IP traffic

Source Type	Source CIDR	IP Protocol	Source Port Range	Destination Port Range
CIDR	A range of IP addresses on your network	TCP	All	22
CIDR	A range of IP addresses on your network	ICMP	All	All
CIDR	A range of IP addresses on your network	TCP	All	443 for BIG-IP VE management access
CIDR	A range of IP addresses on your network	TCP	All	4353
CIDR	A range of IP addresses on your network	TCP	All	6699
CIDR	A range of IP addresses on your network	UDP	All	1026

By default, management and external subnets cannot route traffic outside the VPC. Configure an Internet gateway to enable external connectivity. To create an Internet Gateway, follow the steps listed below:

Create Internet Gateway

- Navigate to **Networking** > **Virtual Cloud Networks**
- Select your VCN
- Click **Internet Gateways** (left navigation)
- Verify the correct compartment is selected (same as previous configurations)
- Provide a descriptive name and click **Create Internet Gateway**

Configure Route Table

Associate the Internet gateway with the VCN route table:

- Navigate to **Networking** > **Virtual Cloud Networks**
- Select your VCN
- Click **Route Tables** (left navigation)
- Select the default route table
- Click **Add Route Rule**
- Enter the following configuration:

  a) **Target Type:** Internet Gateway
  b) **Destination CIDR Block:** 0.0.0.0/0
  c) **Compartment:** (retain default)
  d) **Target Internet Gateway:** Select the gateway created in the previous section

- Click **Add Route Rules** to apply

F5 BIG-IP VE Instance Deployment

The steps listed below shows how to deploy F5 BIG-IP VE instances on Compute Cloud@Customer or Private Cloud Appliance.

1. Open a browser, visit the F5 Downloads page, and then log in or register.
2. On the Downloads Overview page, click Find a Download.
3. Under Product Line, click BIG-IP <version>/Virtual Edition, where <version> is the version, you want to download.
4. Under Name, click x.x.x.x_Virtual-Edition, where x.x.x.x is the product container you want to download, and then at the license agreement notification click I Accept.

> **NOTE**: BIG-IP VE 17.5.1.3 and later are supported.

5. Under Filename, click one of the .qcow2.zip image files.
6. Choose the download location closest to you.
7. When the file finishes downloading, unzip the .qcow2.zip file to a local drive.

> **NOTE**: For Windows, use 7-Zip or for Linux or Mac, use unzip.

8. Once extracted, if necessary, extract the .qcow2 from the .tar file with tar xvfz <filename>.tar, using 7-Zip.

Create a storage bucket and pre-authenticated request

9. Create an OCI object storage bucket and then upload the .qcow2 file.
10. In the OCI console, under the **Home menu**, click **Object Storage**.
11. In the Compartment list, select your compartment, and then click **Create Bucket**.
12. In the bucket name box, enter a name, leave all settings with default values, and then click **Create Bucket**.
13. In the center pane, find your bucket, then **Create Pre-Authenticated Request**.
14. In the name box, enter a name, in the expiration date/time box, select a date for expiration, leave all other settings with the default values, and then click **Create Pre-Authenticated Request** and Close.

15. Click your storage bucket name, under Objects, click Upload Object, browse for the .qcow2 file you downloaded in the previous procedure, and then click Upload Object.
16. The OCI Console provides the PRE-AUTHENTICATED REQUEST URL. Next to your object, click …, select Details, and then copy the URL Path for use in the next step.
17. On Compute Cloud@Customer or Private Cloud Appliance, import the image using the **Pre-Authenticated Request** URL of the BIG-IP VE .qcow2 image created above on the previous steps.
    - On Compute Cloud@Customer or Private Cloud Appliance management console, **navigate to Instances**, click on **Customer Images**, then click **Import Image.**
    - Enter the name of the image, ex: F5-BIG-IP
    - Select your compartment
    - Under source type, select **Import from an Object Storage URL**
    - Enter the Pre-Authenticated Request URL of the BIG-IP VE .qcow2 image created above
    - Image type select .QCOW2
    - Launch mode will be automatically selected, which is the **Paravirtualized Mode**
    - Click **Import image**

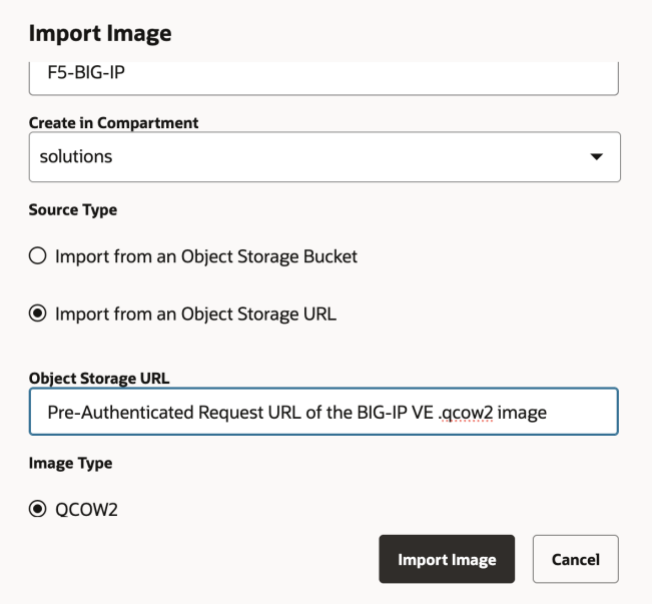

Figure 2. Importing images on Compute Cloud@Customer or Private Cloud Appliance

The import process starts. When the import is complete, the tile next to the image name changes from Importing to Available.

Deploy the BIG-IP VE instance

1. Deploy a BIG-IP VE instance from the custom image created on the Compute Cloud@Customer or Private Cloud Appliance in the previous steps. This instance is referred to as BIG-IP-01
2. Under the **Compute section**, click **Instances**.
3. Click **Create Instance.**
4. In the name box, enter a name of the new instance, on this case, BIG-IP-01

5. Under General, select the **compartment** and the fault domain where the BIG-IP-01 instance will be deployed. On compute Cloud@Customer and Private Cloud Appliance, you can choose the Fault-Domains 1, 2 or 3 or let the system automatically select the best fault domain.
6. Under **Source Image**, select **Custom Images**, the compartment where the customer image was created, then select the F5-BIG-IP custom image created in the previous steps. Click Select Image.
7. Under Shape, select VM.PCAStandard.E5.Flex, then adjust the OCPU and memory to best fit your production environment. Refer to BIG-IP VE Requirements: https://my.f5.com/manage/s/article/K14810
8. Under Boot Volume, specify the boot volume size of 200GB minimum.
9. Under Instance Type, click Virtual Machine.
10. Under Subnet, select the correct VCN and Management interface configured on the previous steps, then click on **Assign Public IP**. This option will assign an IP address from the Compute Cloud@Customer or Private Cloud Appliance pool of IPs already configured on bot platforms, so you can access from your lab network.

> **NOTE**: An SSH key is not required for the F5 BIG-IP instance. SSH access is enabled by default for the root user.

11. Click Create.

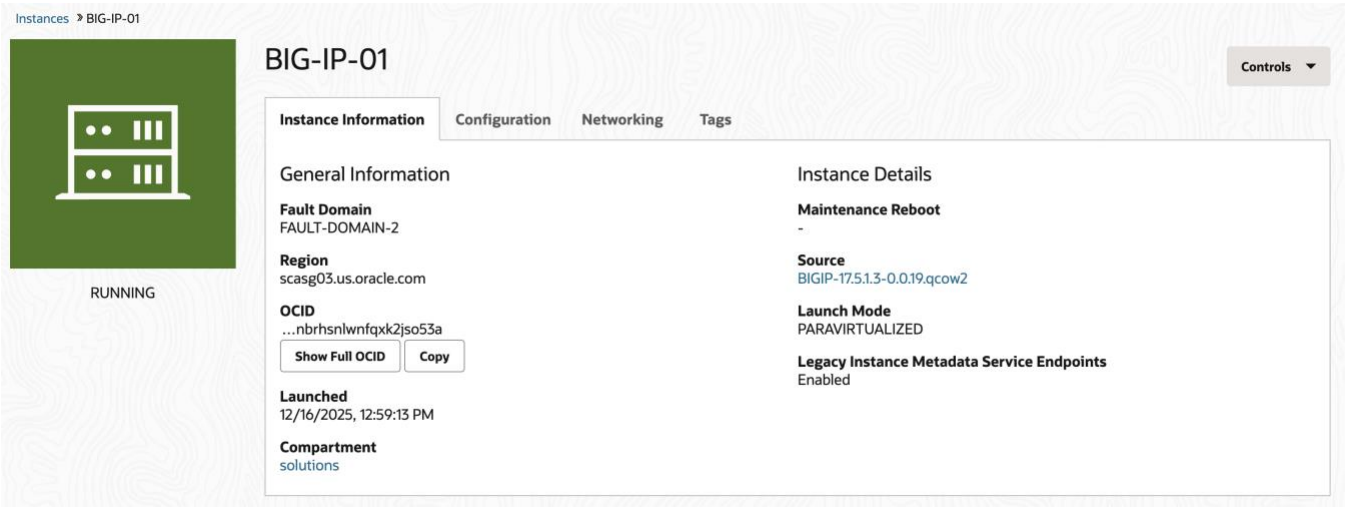

Figure 3. F5 BIG-IP instance running on Compute Cloud@Customer or Private Cloud Appliance

Virtual Network Interface Configuration

In the previous steps, a primary VNIC was attached when the instance was created. This VNIC is dedicated to management traffic and is deployed in the management subnet with the 10.0.0.8/24 private IP without external/public IP attached to it. It is important to mention that for security reasons it is recommended to deploy a bastion host on the management subnet, so the F5 BIG Management interface can be accessed via the bastion host. This will avoid exposing the F5 BIG IP management interface on the public /external subnet.

Three additional VNICs need to be created and attached to the F5 BIG-IP instance to support the remaining subnet traffic requirements. Listed below are the steps to properly create and attach additional VNICs to the F5 BIG-IP instance along with the IP addressing layout needed for this architecture.

- Ensure the BIG-IP VE instance is running on Compute Cloud@Customer or Private Cloud Appliance

- Under the Compute section, click Instances
- Click the F5 BIG-IP instance
- In the left menu, click Attached VNICs.
- Click Create VNIC and complete the information for each VNIC.
- Click Create VNIC for each VNIC.

IMPORTANT: Select the "Skip Source/Destination Check" option for all VNICs below.

Subnet Name	Subnet CIDR	F5 BIG IP VNIC Number	Private IP	Secondary Private IP	Assign public IP address	Description
Management	10.0.0.0/24	1.0	10.0.0.8	No	No	Administrative access to BIG-IP configuration and monitoring
External	10.0.1.0/24	1.1	10.0.1.2	No	Yes 10.122.56.208	Load balancing and application delivery endpoint for all production services
Internal	10.0.2.0/24	1.2	10.0.2.2	10.0.2.202	No	Subnet for all internal servers. Ex: Web, App, DB

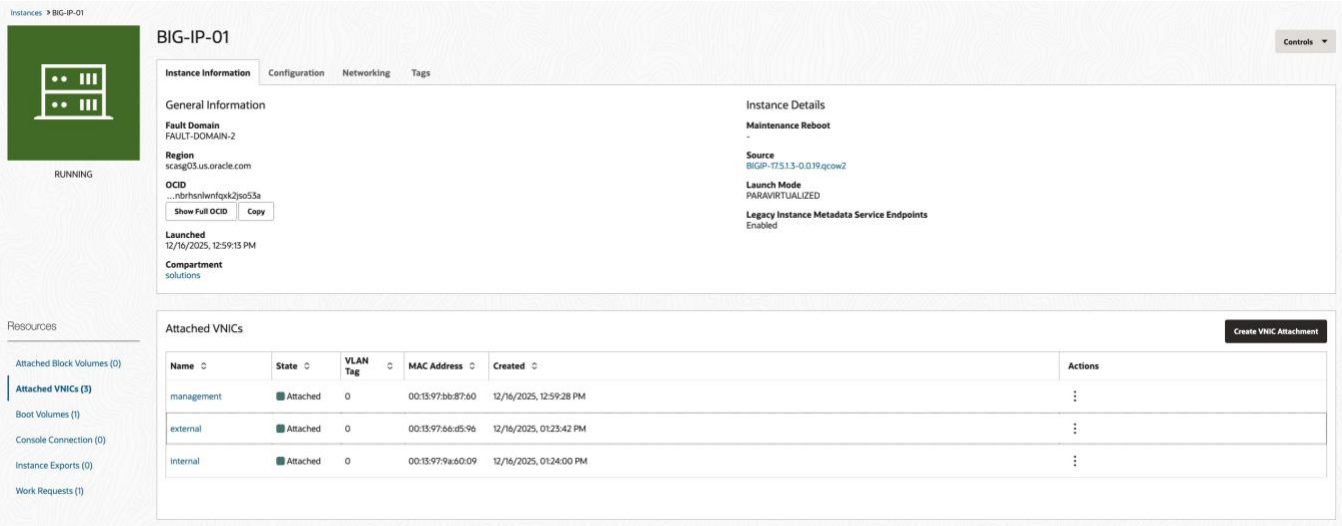

Figure 4. VNICs attached to the F5 BIG-IP instance running on Compute Cloud@Customer or Private Cloud Appliance

Create a secondary private IP for the internal floating self-IP

In the Compute Cloud@Customer or Private Cloud Appliance management interface, create a secondary private IP on the **internal** VNIC.

- Click the top-left menu and under Compute, click Instances.
- Click the F5-BIG-IP instance.
- In the left menu, click Attached VNICs.
- Click the internal VNIC.
- In the right menu, click Assign Secondary Private IP Address
- For Private IP Address, in this example, type 10.0.2.202. Name: **Secondary-Private-IP-Internal-Floating-Self-IP**
- Click Attach IP address.

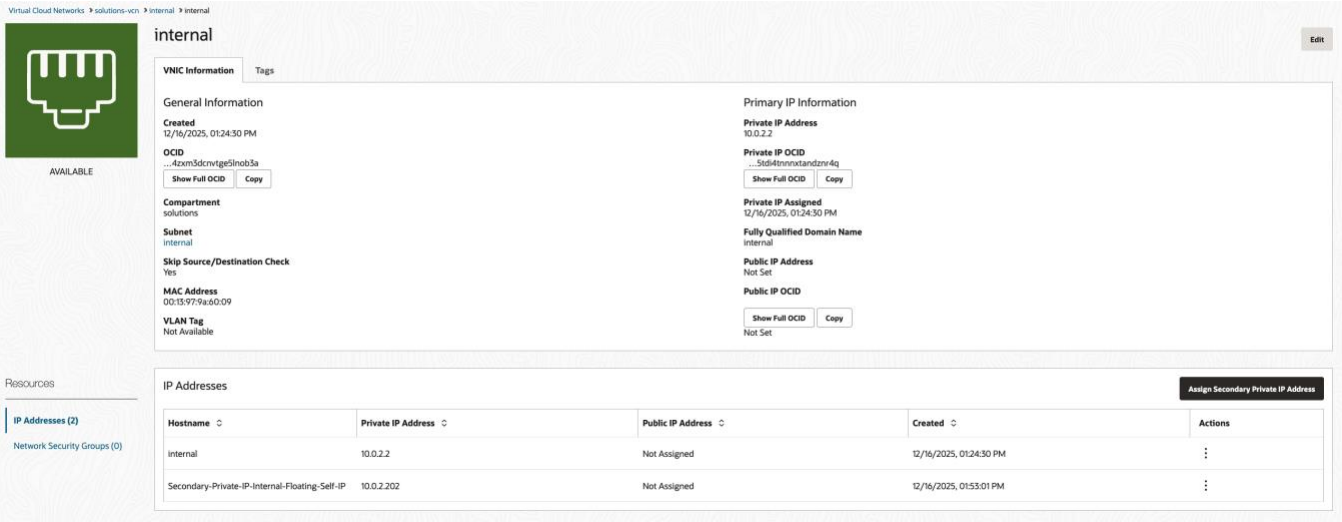

Figure 5. Secondary Private IP addresses configured on the internal VNIC of the F5 BIG-IP instance running on Compute Cloud@Customer or Private Cloud Appliance

- Reboot the F5 BIG-IP Instance. To reboot BIG-IP VE, click your instance, click Reboot, leave the default selection, and then click OK.

IMPORTANT: The BIG-IP VE instance must be rebooted to recognize the VNICs.

IMPORTANT: License BIG-IP VE: You must apply a license before using BIG-IP VE. Contact the F5 sales team to obtain your license, https://www.f5.com/products/get-f5?ls=meta#contactsales

1. Once you have your license, logging with the temporary credentials (root/default) via SSH, then the system will ask to change the password.

```
ssh root@<instance-ip-address>
```

```
Password:
You are required to change your password immediately (root enforced)
Changing password for root.
(current) BIG-IP password:
New BIG-IP password:
Retype new BIG-IP password:

The password for the "admin" user ID has been changed to match the new password for the
"root" user ID.
The password for "admin" user is marked as expired and must be changed the next time the
"admin" user logs in.
Future changes to the "root" password will not affect the password of the "admin" user ID
[root@localhost:NO LICENSE:] config #
```

Open a web browser and log in to the BIG-IP Configuration utility by using **https** with your Primary VNIC's public IP Address, for example: https: <external-ip-address>. The username is **admin** and password for the "admin" user ID has been changed to match the new password for the "root" user ID.

NOTE: You will be asked to change the password of the admin user.

1. On the **Setup Utility Welcome** page, click **Next**.
2. On the **General Properties** page, click **Activate**.
3. In the **Base Registration key** field, enter the case-sensitive **registration key** from F5.

For **Activation Method**, if you have a production or Eval license, choose **Automatic** and click **Next**. If you chose **Manual**, do the following:

a. In the **Step 1: Dossier** field, copy all text, and then **Click here to access F5 Licensing Server**.

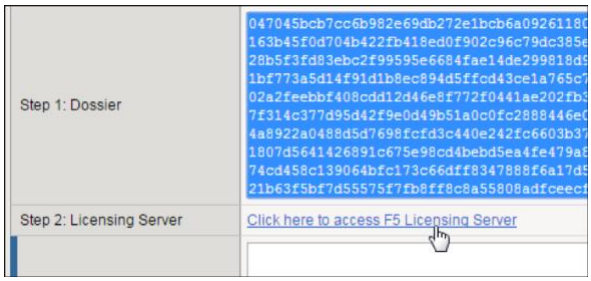

Figure 6. F5 BIG-IP instance licensing step 1

A separate web page opens.

b. In the **Enter your dossier** field, paste the text, and then click **Next**.

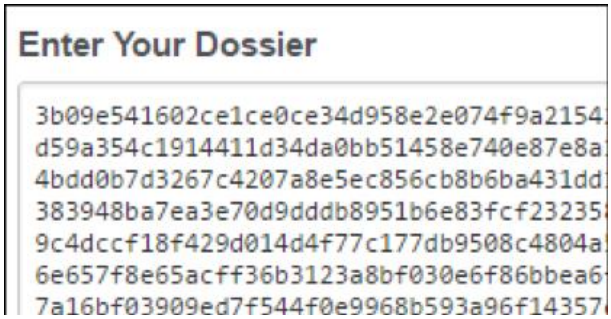

Figure 7. F5 BIG-IP instance licensing step 2

4. Accept the agreement and click **Continue**.
5. On the **Activate F5 Product** page, copy the license text in the box. Return to the BIG-IP Configuration utility and paste the text into the **Step 3: License** field.

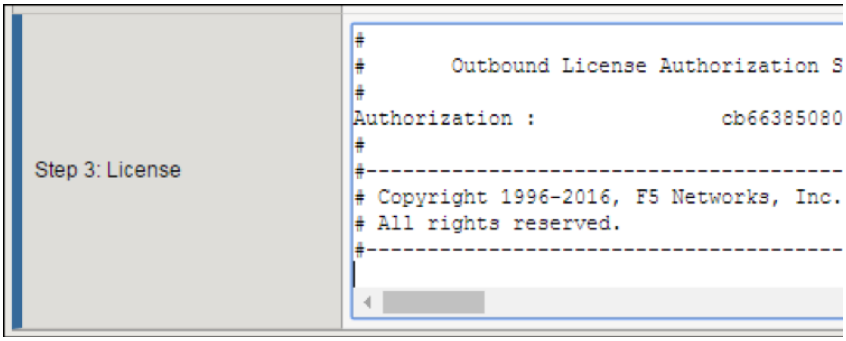

Figure 8. F5 BIG-IP instance licensing step 3

Provision BIG-IP VE

You must confirm the modules you want to run, before you can begin to work in the BIG-IP Configuration utility.

1. Open a Web browser and log in to the BIG-IP Configuration utility.
2. On the **Resource Provisioning** screen, change settings if necessary, and then click **Next**.
3. On the **Device Certificates** screen, click **Next**.
4. On the **Platform** screen, in the **Admin Account** text box, re-enter the password for the root account, and then click **Next**.

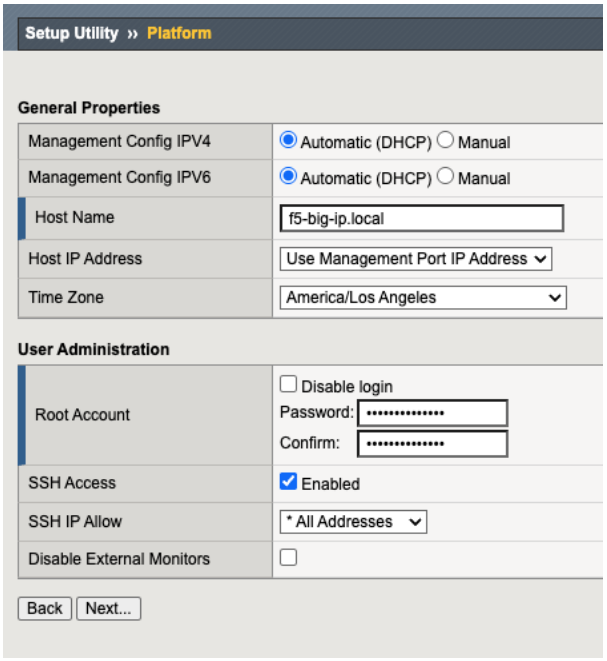

Figure 9. F5 BIG-IP instance platform configuration step 1

5. On the **Advanced Network Configuration** section, click **Finished**.

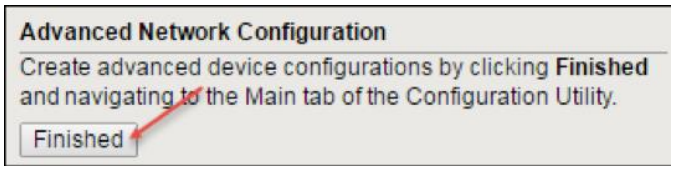

Figure 10. F5 BIG-IP instance platform configuration step 2

Configuring VLANs on F5 BIG-IP

In BIG-IP VE management interface, we need to create an external and internal VLAN that corresponds to the subnets previously created.

> **NOTE:** To ensure correct VLAN-to-interface mapping on F5 BIG-IP, verify that the VNIC MAC addresses in Compute Cloud@Customer or Private Cloud Appliance match the MAC addresses displayed under **Network > Interfaces** in the F5 BIG-IP management interface.

1. On the **Main tab**, click **Network -> VLANs**.
2. Click **Create** and complete the following information for the external VLAN:
   - **Name**: External
   - **Interface**: 1.1
   - **Tagging**: Untagged
3. Click **Finished**.

4. Click **Create** and complete the following information for the HA VLAN.
    - **Name**: Internal
    - **Interface**: 1.3
    - **Tagging**: Untagged
5. Click **Finished**.

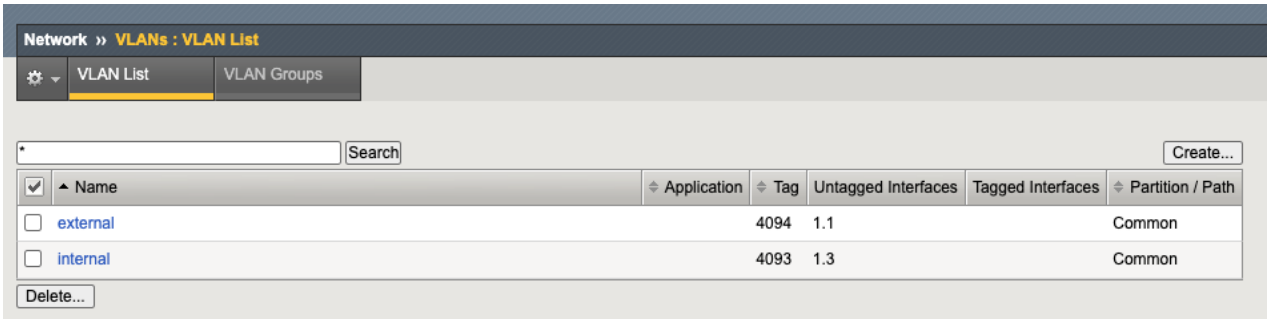

Figure 11. F5 BIG-IP VLANs configuration

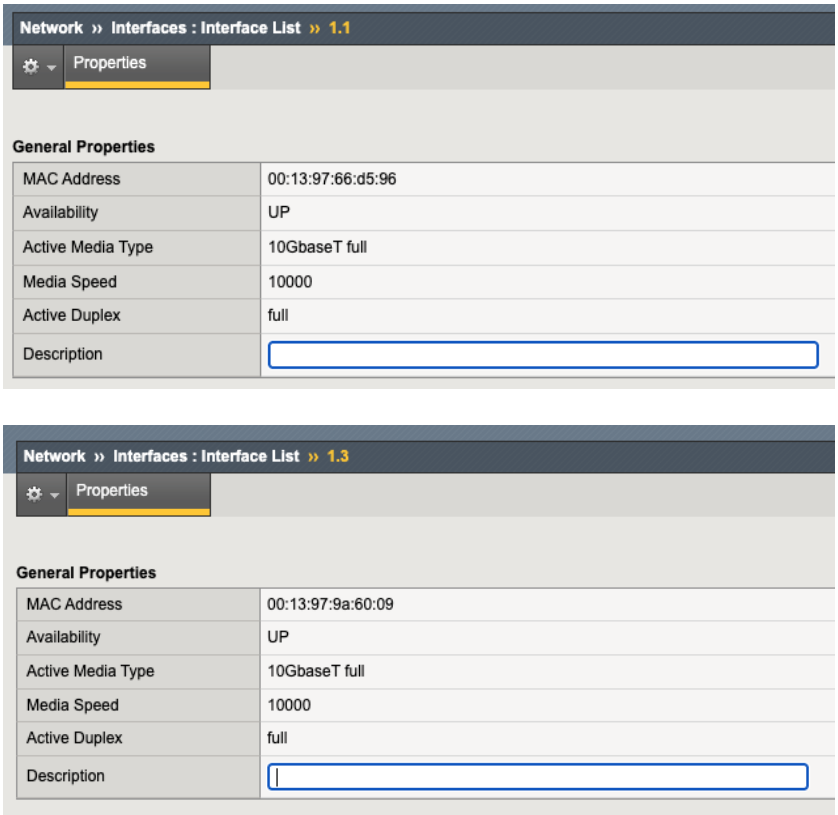

Figure 12. F5 BIG-IP Interface list

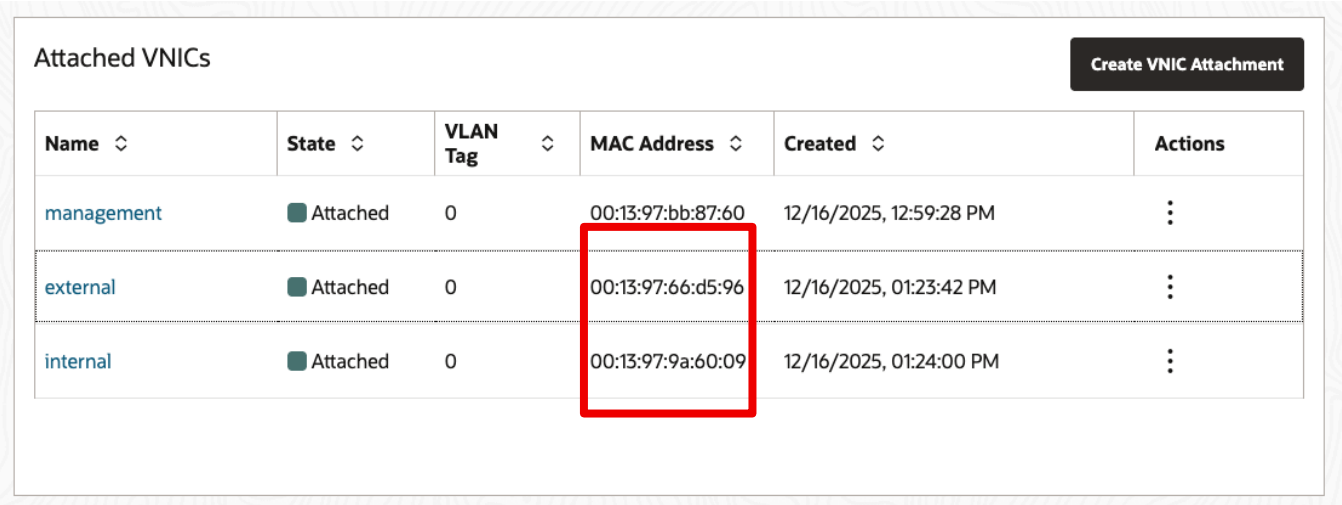

Figure 13. F5 BIG IP VNICs list on Compute Cloud@Customer or Private Cloud Appliance Virtual Network list – MAC Address

Configuring self-IPs on F5 BIG IP

Before using the BIG-IP VE Configuration utility tool, in the Compute Cloud@Customer or Private Cloud Appliance, create one Self IP on F5 BIG-IP VE instance with the following configuration:

management interface, copy the **primary private IP addresses** for the following interfaces: **external and internal.**

Additionally, copy the **secondary IP addresses** for the following network interfaces: **external and internal network**

In BIG-IP VE Configuration utility tool, to create a self IP address, based on these private IP addresses.

1. Click **Repeat** and complete the following information for the internal self IP address.
    - **Name**: InternalSelfIP
    - **IP Address**: 10.0.2.2
    - **Netmask**: 255.255.255.0
    - **VLAN/Tunnel**: internal
    - **Port Lockdown**: Allow None
2. Click **Finished**.

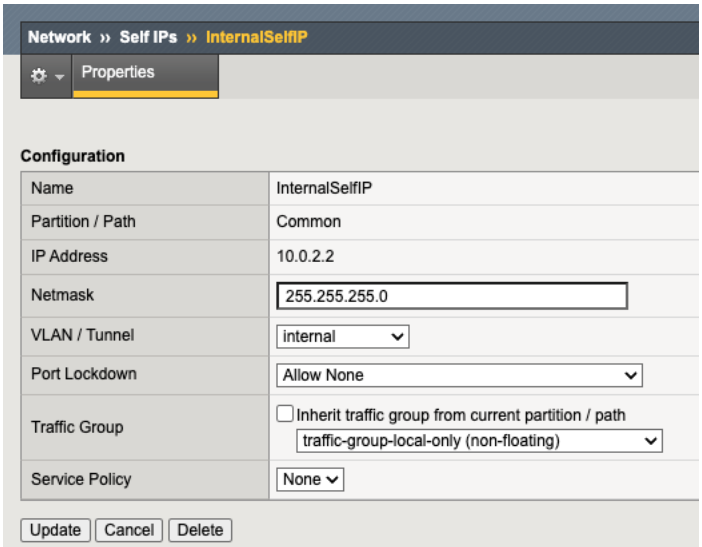

Figure 14. F5 BIG IP Self IPs configuration.

The screen refreshes, and the new self IP addresses appear in the list.

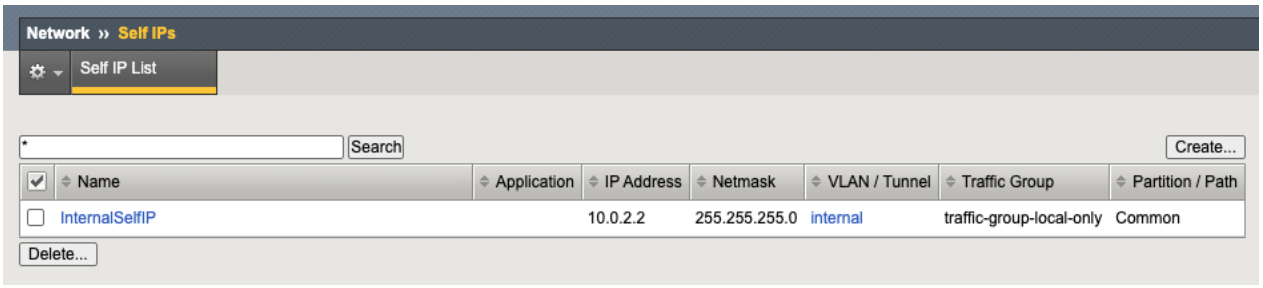

Figure 15. F5 BIG IP Self IPs configuration.

Create a Pool of Servers

A **pool** (also called a "server pool" or "pool member") on F5 BIG-IP is a logical grouping of backend servers that provide the same application or service. It's a fundamental load balancing component that defines which servers will receive traffic distributed by the BIG-IP.

Key Concepts

Pool Components:

- **Pool Members:** Individual backend servers (identified by IP address and port)
- **Health Monitors:** Checks that determine if pool members are available
- **Load Balancing Method:** Algorithm used to distribute traffic (round-robin, least connections, etc.)

How It Works:

1. A virtual server receives incoming client requests

2. The virtual server is associated with a pool
3. BIG-IP distributes requests across available pool members based on the load balancing algorithm
4. Only healthy pool members (passing health checks) receive traffic

In our architecture, we have a basic pool of web servers (10.0.2.3, 10.0.2.4, and 10.0.2.5) all serving on the HTTP/HTTPS which are behind of the Virtual Server IP 10.0.1.2 with an external/public IP 10.122.56.208. When a user connects to https://my-site.domain.com (10.122.56.208), the F5 virtual server receives the request and forwards it to one of the pool members based on the configured load balancing algorithm.

- **Virtual Server:** 10.0.1.2 on the External subnet
- **Backend Server: Web Pool** - This would be a pool containing your three web servers:
  - Pool Member 1: 10.0.2.3:80
  - Pool Member 2: 10.0.2.4:80
  - Pool Member 3: 10.0.2.5:80

Common Pool Configuration

Load Balancing Methods:

- **Round Robin:** Distributes requests sequentially to each server
- **Least Connections:** Sends traffic to the server with fewest active connections
- **Ratio:** Distributes based on weighted ratios assigned to each member

Health Monitoring:

- HTTP/HTTPS checks (verify web server responds)
- TCP checks (verify port is open)
- Custom health checks for specific applications

To configure the pool of server, follow the steps listed below:

1. On the **Main** tab, click **Local Traffic, Pools**, then click **Create**.
2. In the **Name** text box, type webservers. Names must begin with a letter, be fewer than 63 characters, and can contain only letters, numbers, and the underscore (_) character.
3. For **Health Monitors**, move **https** from the **Available** to the **Active** list.
4. Choose the **load balancing method** or retain the default setting.
5. In the **New Members** section, in the **Address** text box, type the IP address of the web servers.
6. In the **Service Port** text box, type a service port; for example, **443**.
7. Click **Add**. The member appears in the list.
8. Add more pool members as needed, and then click **Finished**.

Create a virtual server

A **virtual server** on F5 BIG-IP is the front-end listener that accepts incoming client traffic and acts as the entry point for applications. It's essentially the public-facing or client-facing endpoint that clients connect to, which then routes traffic to backend servers.

Key Concepts

- Listens on a specific IP address and port for incoming connections

- Acts as a reverse proxy between clients and backend servers
- Applies security policies, SSL/TLS termination, and traffic management rules
- Routes traffic to pools of backend servers based on configuration

Core Components:

- **Virtual Server IP Address:** The IP that clients connect to (e.g., 10.0.1.2)
- **Port/Protocol:** What port and protocol it listens on (HTTP/80, HTTPS/443, custom ports)
- **Pool Association:** Which pool of backend servers receives the traffic
- **Profiles:** SSL, HTTP, TCP configurations applied to connections
- **iRules/Policies:** Advanced traffic routing and manipulation logic

How It Works

Traffic Flow:

1. **Client connects** to the virtual server IP (e.g., https://my-site.domain.com → 10.122.56.208)
2. **Public IP maps** to the virtual server's private IP (10.122.56.208 → 10.0.1.2)
3. **Virtual server processes** the request (SSL termination, security checks, etc.)
4. **Virtual server forwards** the request to a backend server from the associated pool
5. **Backend server responds** through the virtual server back to the client

A virtual server needs to be created for the private IP address that's associated with the external network interface, on this case, 10.0.1.2. To create a new Virtual Server on F5 BIG-IP VE, follow the steps listed below:

1. In the BIG-IP Configuration utility, on the **Main** tab, click **Local Traffic -> Virtual Servers**.
2. Click **Create** and complete the following information:
    - **Name**: A unique name.
    - **Destination Address/Mask**: 10.0.1.2 (The private IP address on the external NIC).
    - **Service Port**: A port number or a service name from the Service Port list.
    - **HTTP Profile**: http.
    - **Source Address Translation**: Auto Map.
    - **Default Pool**: webpool.
3. Configure any other settings as needed, and then click **Finished**.

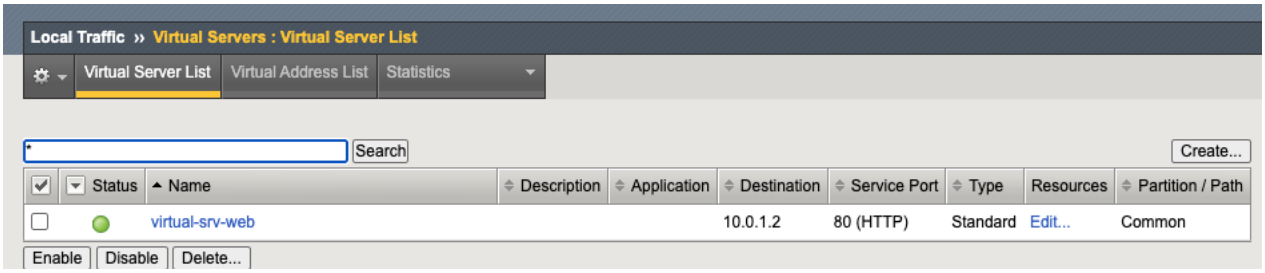

Figure 16. F5 BIG IP Virtual Servers configuration.

Traffic to the virtual server IP address will now go to the pool members. Testing the access to the https://my-site.domain.com → 10.122.56.208, you will see the webpage of your backend web servers, loadbalance between the three servers, web1, web2, and web3.

Figure 17. Web servers deployed with F5 BIG IP Virtual Servers.

ORACLE

Connect with us

Call +**1.800.ORACLE1** or visit **oracle.com**. Outside North America, find your local office at: **oracle.com/contact**.

blogs.oracle.com facebook.com/oracle twitter.com/oracle

United States only or elsewhere. Not all products and features discussed are currently offered for sale in the United States or elsewhere. Products and features of the actual offering may differ from those discussed in this document and may vary from country to country. Any timelines contained in this document are indicative only. Timelines and product features may depend on regulatory approvals or certification for individual products or features in the applicable country or region.

Author: Anderson Souza