F5 BIG-IP and Cisco Nexus 9000 Series Switches: OpenStack Deployment Guide

Guide

Contents

Purpose of This Document 3
What Is OpenStack? 3
Cisco Nexus 9000 Series Servers 3
Cisco Application Centric Infrastructure 3
F5 BIG-IP 3
Installation and Provisioning 4
Provision the Cisco Nexus 9000 Series in NX-OS Mode 4
Provision F5 BIG-IP 4
Install the F5 Plug-in on OpenStack 5
Configure and Load-Balance an Application Service 6
Cisco ACI and OpenStack Integration with F5 13
Conclusion 14
For More Information 14

Purpose of This Document

This document explains how to provision F5 BIG-IP application delivery controllers (ADCs) and Cisco Nexus® 9000 Series Switches in an OpenStack environment. The document begins with an overview and then provides a detailed process for creating the solution.

What Is OpenStack?

OpenStack is an open-source cloud platform that provides a collection of tools for creating and managing computing clouds for infrastructure as a service (IaaS). These tools can be used for both public and private clouds and for small, medium-sized, and large organizations. The main capabilities of OpenStack discussed in this document are the computing and network functions for application services. The specific tool used is the OpenStack Horizon dashboard.

With F5 BIG-IP, applications become highly available, more reliable, and easy to scale. You can perform routine maintenance without application downtime; because multiple servers share the application load, any single server can be taken offline without affecting availability. Because multiple server instances are also responding to requests sent to the virtual application server, as workloads increase, the capability to quickly instantiate additional servers in OpenStack becomes critical; after those servers are up and running, they can quickly be added to the virtual application server pool, allowing application bandwidth to expand and contract on demand.

Cisco Nexus 9000 Series Servers

The Cisco Nexus 9000 Series offers high-performance data center switches that include both fixed-configuration and modular models. It has a variety of physical interfaces that include 1-, 10-, 40-, and, 100-Gbps connectivity to all types of devices in a typical data center or cloud environment. The Cisco Nexus 9000 Series has two modes of operation: one with Cisco® NX-OS Software and the other with Cisco Application Centric Infrastructure (Cisco ACI™). This document focuses in detail on OpenStack integration with NX-OS. This document also presents the principles of Cisco ACI integration.

Cisco Application Centric Infrastructure

Cisco ACI is a policy-based software-defined networking (SDN) architecture with outstanding scalability and security and the capability to quickly and repeatedly deploy applications and policies. Cisco ACI lets you use southbound APIs for integration with F5 and service appliances for service chaining, and northbound APIs for integration with OpenStack. This capability allows users to choose open-source tools such as Horizon and API calls for provisioning and orchestration, or any other orchestrator tool using OpenStack.

F5 BIG-IP

F5 BIG-IP offers high-performance ADC platforms that provide application services. The BIG-IP platforms are delivered as both physical and virtual devices. The physical devices are available as fixed-configuration or modular chassis (F5 VIPRION). The overall features are uniform across the platforms because they share a single software base, but differences in scalability, performance, and interface flexibility make each product platform unique.

Installation and Provisioning

Provision the Cisco Nexus 9000 Series in NX-OS Mode

The Cisco Nexus 9000 Series Switches provide the physical connectivity for the network-connected devices. Only basic configuration of the switches is necessary, with the management IP address and Secure Shell (SSH) access. Communication between OpenStack Neutron and the Cisco Nexus switch is through standard NETCONF commands.

See the following links for detailed configuration details and more information about the Cisco Nexus 9000 Series and OpenStack:

- Neutron modular Layer 2 (ML2) driver for Cisco Nexus devices: https://wiki.openstack.org/wiki/Neutron/ML2/MechCiscoNexus
- Cisco NX-OS release notes for Cisco Nexus 9000 Series Switches: http://www.cisco.com/c/en/us/support/switches/nexus-9000-series-switches/products-release-notes-list.html

Note: Refer to the Cisco documentation for the latest supported tenant network types. At the writing of the document, only network type VLAN (IEEE 802.1Q) is supported.

Provision F5 BIG-IP

The F5 BIG-IP Local Traffic Manager (LTM) provides the load-balancing functions for OpenStack. You can use either BIG-IP physical devices or BIG-IP Virtual Edition (VE) because the feature sets are similar. You can follow the normal BIG-IP provisioning steps to bring the system online. You need to configure only the management IP address and licenses before you integrate BIG-IP into an OpenStack environment.

If you use physical BIG-IP, then you set the management IP address and licensing in the normal way.

If you use BIG-IP VE, you can choose among three platform images. Table 1 lists the differences in the platforms.

Table 1: F5 BIG-IP VE Images

Image Name	Features Available	Minimum Required Disk Size[1]
BIGIP-11.6.0.0.0.401.LTM_1SLOT.qcow2.zip	Single-slot LTM module	7 GB
BIGIP-11.6.0.0.0.401.LTM.qcow2.zip	Two LTM modules plus hot-fix space allocated	31 GB
BIGIP-11.6.0.0.0.401.ALL.qcow2.zip	All modules and combinations plus hot-fix space allocated	124 GB

[1] Refer to the latest documentation for the exact volume size requirements, virtual CPU (vCPU) requirements, and memory requirements.

For additional information about BIG-IP VE, also see the following links:

- BIG-IP getting-started guide: https://support.f5.com/kb/en-us/products/big-ip_ltm/manuals/product/bigip-ve-setup-amazon-ec2-11-6-0/1.html
- BIG-IP VE 11.6.0 release notes: https://support.f5.com/kb/en-us/products/big-ip_ltm/releasenotes/product/relnote-ve-11-6-0.html
- Procedure for extending disk space: https://support.f5.com/kb/en-us/solutions/public/14000/900/sol14952.html

After the BIG_IP virtual machine has been created, you need to update the platform file so that the BIG-IP software is aware of the infrastructure on which it is running. During the initial power up, log in to the system console and update the /PLATFORM file with the following entries:

```
vi /PLATFORM

Platform = Z100
Family = 0xC0000000
Host = Z100
Systype = 0x71 Platform = Z100
Family = 0xC0000000
Host = Z100
Systype = 0x71
```

After this task has been completed, save the file and reboot BIG-IP VE. Continue with the initial license and management setup as needed.

Install the F5 Plug-in on OpenStack

F5 BIG-IP communicates with the Neutron server. To accomplish this, a plug-in is provided for Neutron. This plug-in, detailed installation information, and most current versions are available through DevCentral.F5.com, at https://devcentral.f5.com/d/openstack-neutron-lbaas-driver-and-agent.

Download the plug-in driver to the host running your Neutron API server process. The F5 load-balancing-as-a-service (LBaaS) plug-in driver and the LBaas Agent with iControl driver support the OpenStack Neutron Havana and Icehouse releases. For deployment modes that include tenant tunnels, Neutron must use the ML2 core plug-in.

To install the driver, use the following command:

```
sudo dpkg –I f5-lbass-driver_1.0-1_all.deb
```

After the plug-in has been successfully installed, update the neutron.conf file with the following entries:

```
service_plugins = neutron.services.loadbalaner.plugin.LoadBalancerPlugin

service_provider=LOADBALANCER:F5:neutron.services.loadbalancer.drivers.f5.plugin_driver.F5PluginDriver:default
```

With the driver installed, you now need to install the LBaaS agent. Install the agent on the host that runs agent processes. Download the agent f5-bigip-lbaas-agent_1.0-1_all.deb and run the following command to install the agent:

```
sudo dpkg -i f5-bigip-lbaas-agent_1.0-1_all.deb
```

With the agent installation complete, you need to stop the agent service and configure the F5 LBaaS agent. This configuration file specifies how BIG-IP is accessed, what ports are connected to the network, and various other critical parameters. Use the following command to stop the agent:

```
sudo service f5-bigip-lbaas-agent stop
```

After the agent has been stopped, you need to update the configuration file (f5-bigip-lbaas-agent.ini). The default configuration file is located at:

```
/etc/neutron/f5-bigip-lbaas-agent.ini
```

Details about each value setting are provided in the f5lbaas-readme.html file included with the driver file download from DevCentral.f5.com. The file also contains detailed steps for the installation process. At minimum, you should update or customize the following items for your specific installation:

- f5_external_physical_mappings
- f5_ha_type
- icontrol_hostname
- icontrol_username
- icontrol_password

Configure and Load-Balance an Application Service

At this point, both the Cisco Nexus 9000 Series Switch and BIG-IP are configured with OpenStack. You now need to configure and load-balance your application service. This section provides a step-by-step example of this process using the Horizon dashboard. The process involves these main steps:

1. Create a network for your servers.
2. Create a virtual application server pool.
3. Assign members to the pool.
4. Create the virtual IP address.
5. Create and assign a monitor service.

The first step is to log into the Horizon OpenStack dashboard (Figure 1). This dashboard is where you create the virtual machines and networks and set up load balancing.

Figure 1: Horizon Login Screen

After you have logged in to the dashboard, you create a network for the application servers (Figure 2). First enter a name for the network.

Figure 2: Create a Network Name in OpenStack

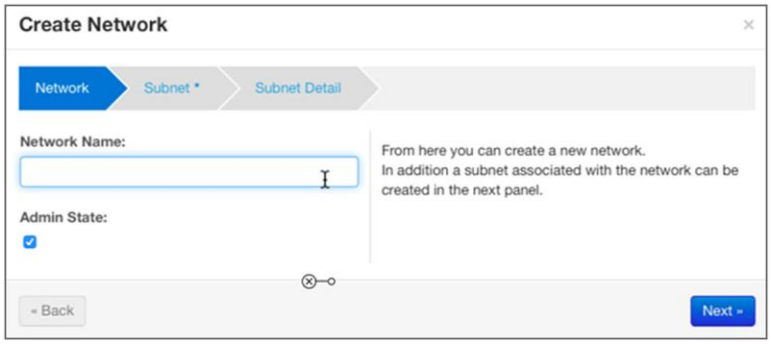

Next you need to provide a name for the subnet. This name is equivalent to the VLAN description on a networking switch. You also need to provide the network address, allowed IP version, and address of the default gateway (Figure 3).

Figure 3: Network Details

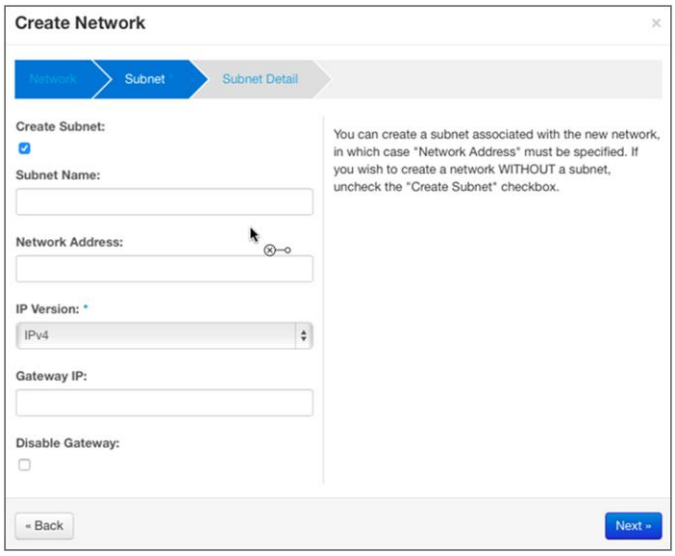

After you have provided the higher-level network parameters, you have the option of providing Domain Host Configuration Protocol (DHCP) services for any machine that joins the subnet. Here you enter the start and end of the DHCP address range, the Domain Name System (DNS) servers, and any host routes that should be provided by DHCP (Figure 4).

Figure 4: Configure DHCP Options for the Subnet

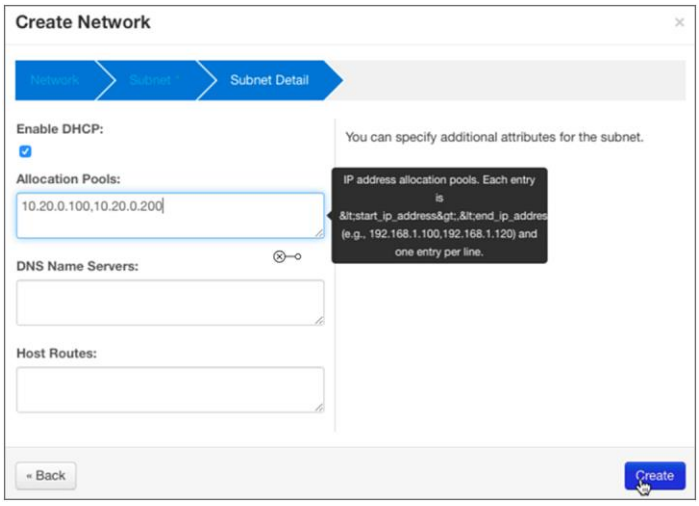

After you have specified the network, review the summary (Figure 5). Make note of the network ID; it will be used to automatically name the BIG-IP partition in future steps.

Figure 5: Network Configuration Summary

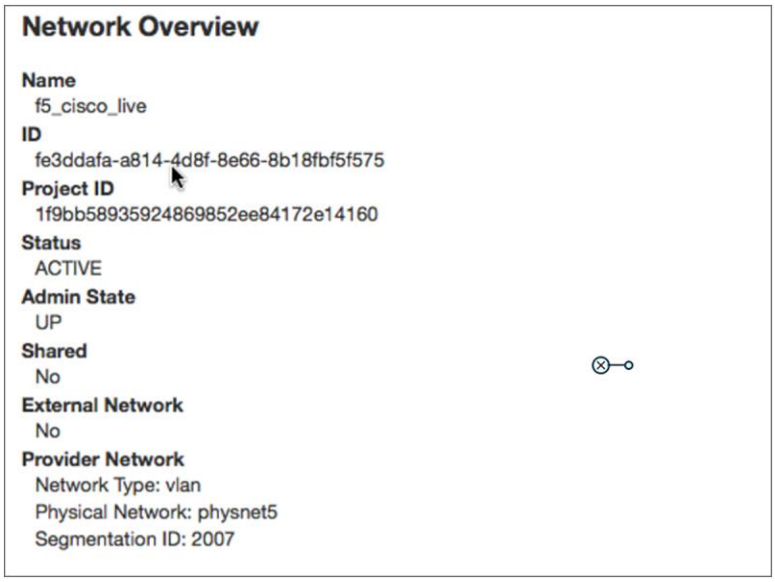

You can now create the virtual machines to be used as application servers. After they have been created and provisioned for the network that was created, the network topology will reflect their existence in the system. Figure 6 shows that two virtual machines are attached to the network for use.

Figure 6: Network Topology Showing Two Application Servers Attached

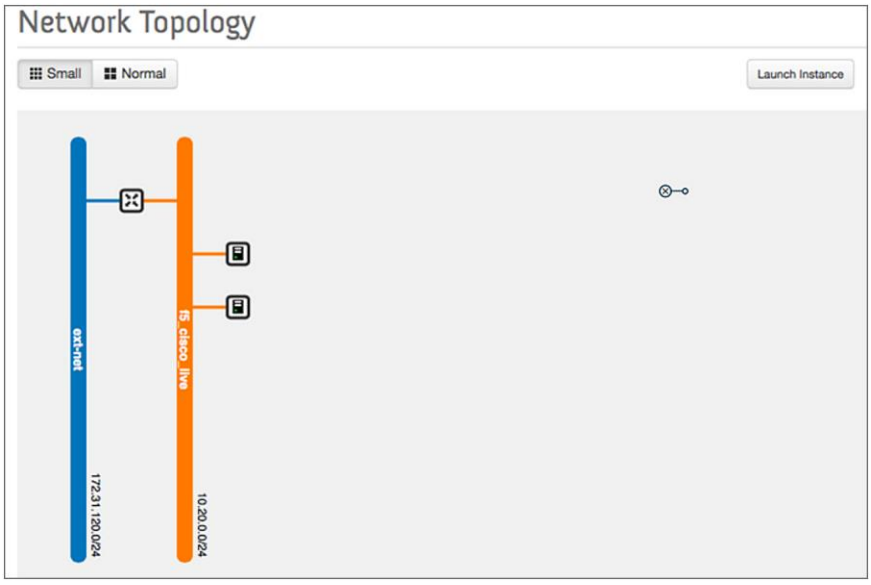

With the application servers attached to the network, you can add the load-balancing function for the servers (Figure 7). In OpenStack, this process is referred to as making an application highly available.

Figure 7: Create a Load-Balancing Pool

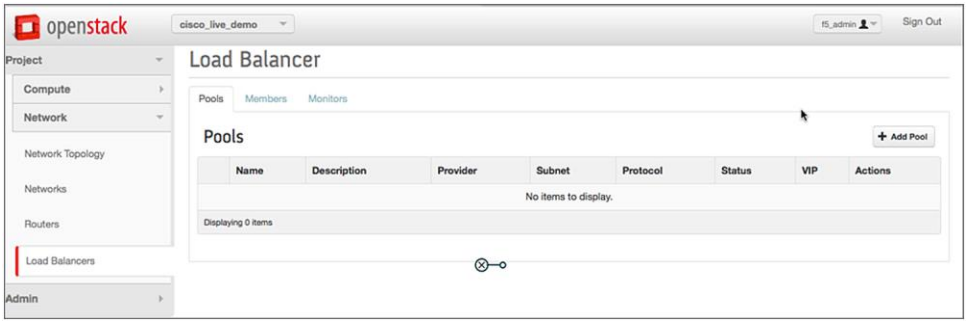

Use the Add Pool option to create the pool object in OpenStack. Then specify the necessary entries (Figure 8). The provider is the BIG-IP device that was previously provisioned. Be sure to also include the subnet, protocol, and load-balancing method.

Figure 8: Pool Parameters

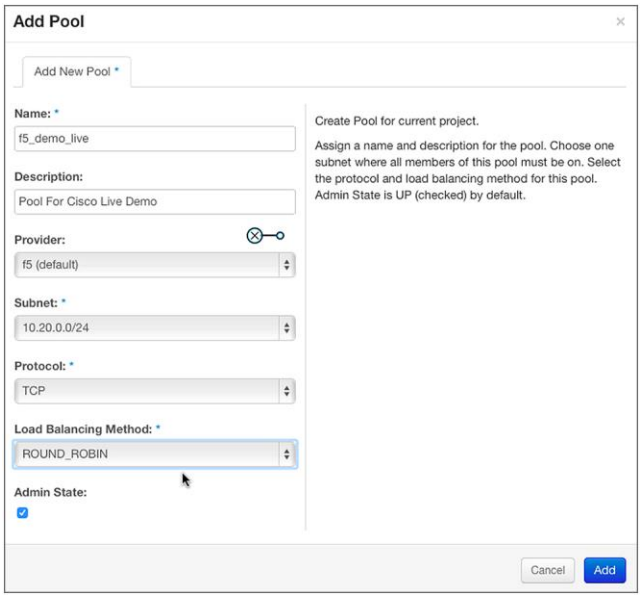

With the pool created, you now need to add servers to the pool (Figure 9). These application servers will be used to service requests. Available servers are shown under members. By having multiple servers responding to incoming requests, you enhance application bandwidth, availability, and ease of maintenance.

Figure 9: Add Pool Members

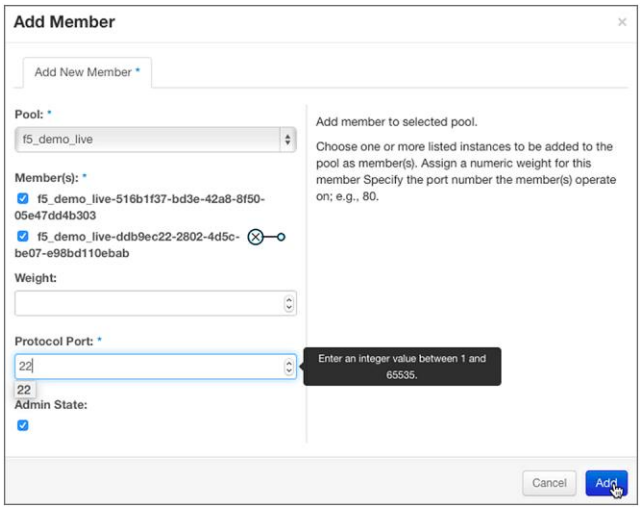

With the pool members added to the pool, you now need to create a virtual IP address for the pool (Figure 10). This is the address that clients will access to get to the load-balanced application.

Figure 10: Add Virtual IP Address

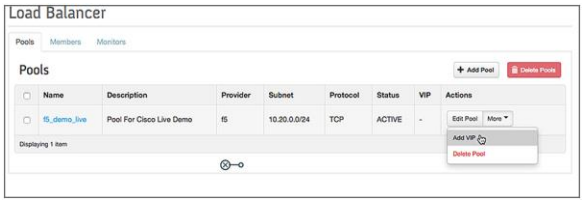

Click Add VIP and then enter the IP address to use (Figure 11). Make sure that the virtual IP address is unique and is not in any defined DHCP pool ranges, to prevent conflicts. Also configure Protocol Port, Protocol, Session Persistence, and Connection Limit as needed.

Figure 11: Virtual IP Address Configuration Details

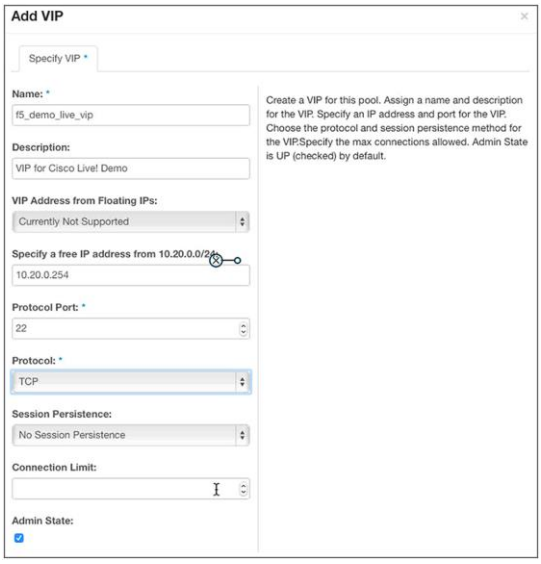

The final step is to create and configure a pool monitor (Figure 12). This monitor is used to check the health of servers in the pool and also to help ensure that incoming requests are sent only to healthy servers. The pool monitor type, delay between probes, timeout values, and number of retries all need to be configured properly to determine when to declare a server as healthy or down.

Figure 12: Create a Pool Monitor

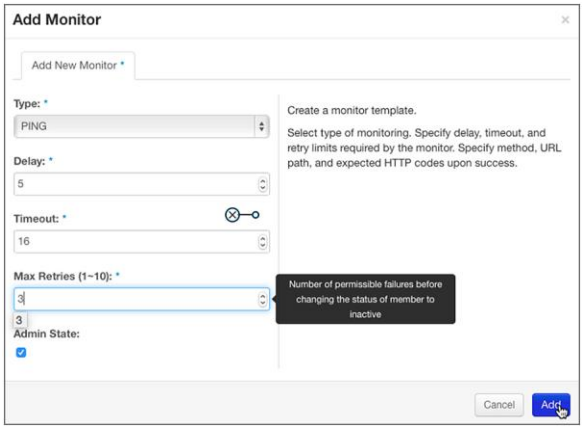

After the pool monitor has been created, you need to associate it with the application server pool (Figure 13).

Figure 13: Monitor Association

Now the configuration of the application server is complete, and the application server is ready for use.

Log in to the BIG-IP to confirm that the same configurations have been automatically applied by the Horizon dashboard (Figure 14, 15, and 16).

Figure 14: BIG-IP Virtual Server Verified as Properly Configured

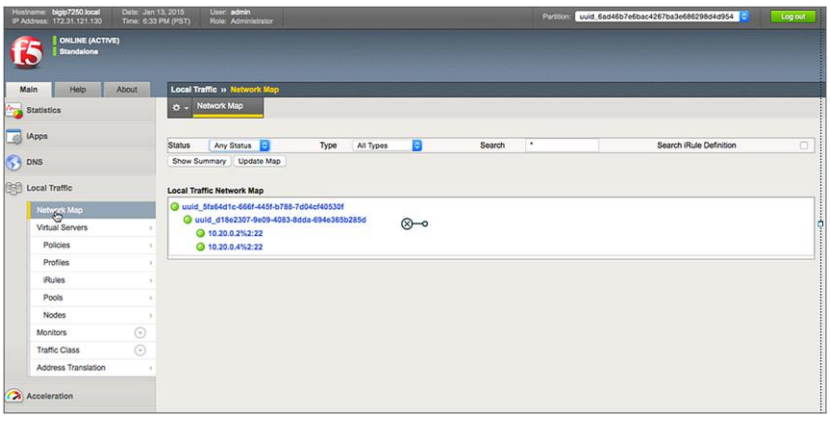

Figure 15: BIG-IP Self-IP Verified as Properly Configured

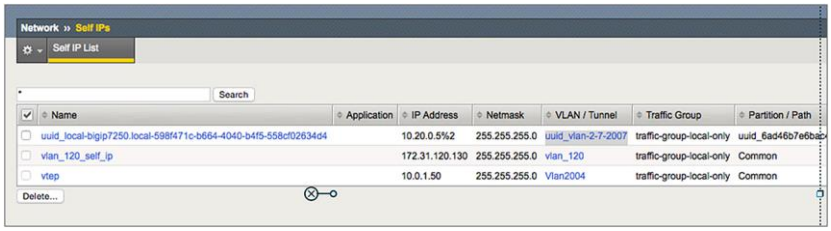

Figure 16: VLANs Verified as Properly Configured

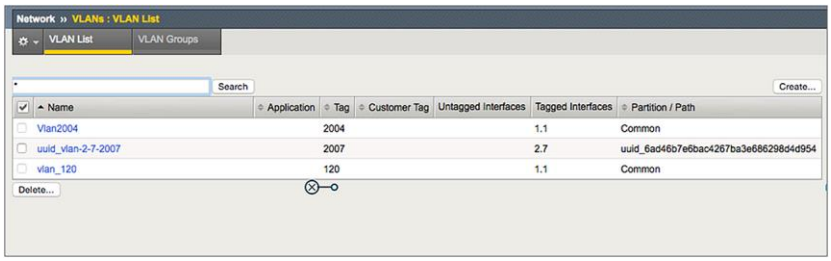

Cisco ACI and OpenStack Integration with F5

The procedure for integrating the Cisco Application Policy Infrastructure Controller (APIC) and OpenStack is similar to that for integrating Cisco Nexus 9000 Series Switches. The integration provides the capability to provision the fabric for end-to-end connectivity and policy deployment. Through the use of the APIC plug-in for OpenStack Neutron, tenants can be transparently configured and managed. The plug-in allows the Cisco ACI fabric to automatically configure items such as VLANs, subnets, and routers to enable end-to-end connectivity and communication through OpenStack tools.

For more information, see:

- Cisco Nexus 9000 Series and OpenStack installation video: https://www.youtube.com/watch?v=pWMXTb237Vk
- Cisco Nexus 9000 Series and OpenStack overview video: http://youtu.be/pQXysWvCPRQ

With this Cisco ACI integration, BIG-IP physical and virtual devices can be used as endpoints in the fabric. Three methods are available to use BIG-IP in the fabric:

- F5 OpenStack LBaaS plug-in for BIG-IP: Use the plug-in to enable the use of a single pane for configuration. This approach allows both the fabric and load balancing to be configured and deployed at the same time. This mode uses an OpenStack orchestrator as the single management pane.
- Standalone F5: With this method, you keep the management and configuration of BIG-IP separate from that for OpenStack. Use this method in cases in which the desired capabilities are beyond the capabilities of the LBaaS plug-in: for example, when you want to use features specific to BIG-IP such as a health monitor specific to F5, or when you want to use layered F5 functions such as firewalls, web application firewalls (WAFs), and LTM in parallel.

- The third method of integration of BIG-IP with ACI and OpenStack involves the use of the Application Policy Infrastructure Controller (APIC) specific feature known as Service Insertion. With service insertion, Cisco APIC is responsible for the creation, management, stitching and traffic redirection into the BIG-IP devices for L4-L7 services in the ACI fabric. This tight integration with APIC allows for the creation of L4-L7 policies in to be applied to endpoints needing communication. Since the policies are defined in APIC they can be easily reused and redeployed for other endpoints requiring the same APIC policy. APIC service insertion is accomplished through the F5 BIG-IP device package for APIC. The device package pushes configuration information from APIC into BIG-IP for L4-L7 services. The below high level procedure would be used.
  - Create the BIG-IP device service node in APIC.
  - Create and configure the necessary service policies and graphs in APIC for the applications that utilize the BIG-IP device cluster.
  - Using Openstack and the ML2 plugin for APIC create the necessary endpoints, networking configuration, and virtual machines.
  - While in Openstack specify the networking policy/contract that should be used for communication between the endpoints
  - Once the contract has been created my Openstack, an update in APIC is necessary to create the necessary service graph for the L4-L7 services.

Additional general information on device package integration between F5 BIG-IP and Cisco ACI can be found at the following links http://www.f5.com/cisco.

Conclusion

With OpenStack growing in popularity, you must be sure that application deployments are configured correctly to use it. To help ensure that applications remain highly available, you can use load balancers, using F5 BIG-IP LTM. The low cost and high performance of the Cisco Nexus 9000 Series coupled with the dependability and stability of BIG-IP provide and excellent OpenStack solution for the modern data center and both public and private clouds.

For More Information

http://www.cisco.com/go/nexus9000

http://www.cisco.com/go/aci

http://www.cisco.com/go/openstack

http://www.f5.com

http://www.openstack.org

F5 (NASDAQ: FFIV) provides solutions for an application world. F5 helps organizations seamlessly scale cloud, data center, and software defined networking (SDN) deployments to successfully deliver applications to anyone, anywhere, at any time. F5 solutions broaden the reach of IT through an open, extensible framework and a rich partner ecosystem of leading technology and data center orchestration vendors. This approach lets customers pursue the infrastructure model that best fits their needs over time. The world's largest businesses, service providers, government entities, and consumer brands rely on F5 to stay ahead of cloud, security, and mobility trends. For more information, go to f5.com.

C11-735572-01 11/15