VMware® NSX™ for vSphere® (NSX-V) and F5® BIG-IP®

Design Guide

vmware®

Table of Contents

Intended Audience

The intended audience for this document includes virtualization and network architects seeking to deploy VMware® NSX™ for vSphere® in combination with F5® BIG-IP® Local Traffic Manager™ devices.

Note: A solid understanding based on hands-on experience with both NSX-v and F5 BIG-IP LTM is a pre-requisite to successfully understanding this design guide.

Overview

The Software Defined Data Center is defined by server virtualization, storage virtualization and network virtualization and server virtualization has already proved the value of SDDC architectures in reducing costs and complexity of compute infrastructure. VMware NSX network virtualization provides the third critical pillar of the SDDC and extends the same benefits to the data center network to accelerate network service provisioning, simplify network operations and improve network economics.

VMware NSX-v is the leading network virtualization solution in the market today and is being deployed across all vertical markets and market segments. NSX reproduces L2-L7 networking and security including L2 Switching, L3 Routing, Firewalling, Load Balancing, and IPSEC/VPN secure access. services completely in software and allows programmatic provisioning and management of these services. More information about these functions is available in the NSX Design Guide here.

F5 BIG-IP is the leading application delivery controller in the market today. The BIG-IP product family provides Software-Defined Application Services™ (SDAS) designed to improve the performance, reliability and security of mission-critical applications. BIG-IP is available in a variety of form factors, ranging from ASIC-based physical appliances to vSphere-based virtual appliances. NSX deployments can be coupled with F5 BIG-IP appliances or Virtual Edition form factors. Learn more about F5 SDAS here.

Furthermore, F5 offers a centralized management and orchestration platform called BIG-IQ.

By deploying BIG-IP and NSX together, organizations are able to achieve service provisioning automation and agility enabled by the SDDC combined with the richness of the F5 application delivery services they have come to expect.

This design guide provides recommended practices and topologies to optimize interoperability between the NSX platform and F5 BIG-IP physical and virtual appliances. This interoperability design guide is intended for those customers who would like to adopt the SDDC while ensuring compatibility and minimal disruption to their existing BIG-IP environment. The Recommended practice guide will provide step-by-step guidance to implement the topologies outlined in this document.

NSX/F5 Topology Options

"BIG-IP Form Factor" / "NSX overlay or not" / "BIG-IP placement" Relationships

There are about 20 possible topologies that can be used when connecting BIG-IP to an NSX environment but this Design Guide will focus on the three that best represent the form factor, connection method, and logical topology combinations. In addition, the Design Guide will highlight the Pros and Cons of each of the three topologies.

The following figure describes the relationship of:

- BIG-IP form factor:
  - BIG-IP Virtual Edition ("VE")
  - BIG-IP physical appliance
- With NSX overlay/Without NSX overlay:
  - VXLAN
  - non-VXLAN (VLAN tagged on untagged)
- BIG-IP placement:
  - BIG-IP parallel to NSX Edge
  - BIG-IP parallel to DLR
  - BIG-IP One-Arm connected to server network(s)
  - BIG-IP on top of NSX Edge
  - BIG-IP on top of NSX DLR

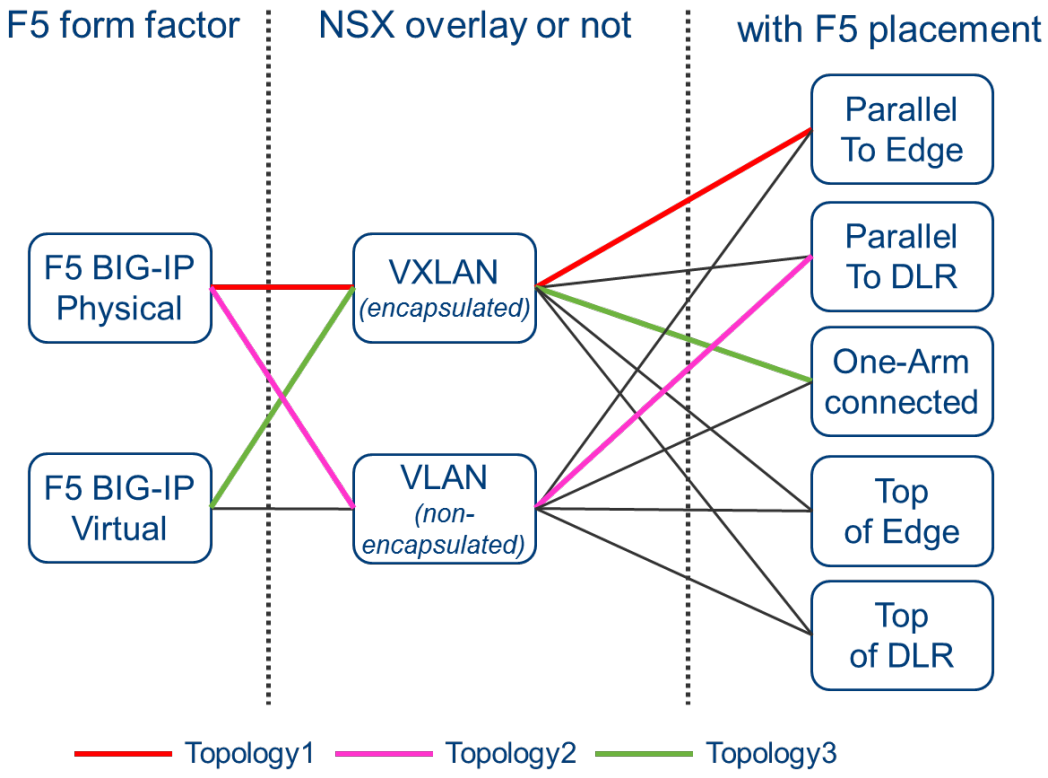

Figure 1 – "BIG-IP Form Factor" / "NSX overlay or not" / "BIG-IP placement" Relationships

Topology #1: Parallel to NSX Edge Using VXLAN Overlays with BIG-IP Physical Appliances

In this topology, the BIG-IP devices are placed parallel to the Edge. BIG-IP appliances are used in this description but BIG-IP VE could be used the same way.

This topology is popular on Layer 3 physical fabrics, such as Leaf/Spine, but also works on Layer 2 physical fabrics.

This topology option is represented by the red lines in Figure.

Logical and Physical view

The BIG-IP appliances are logically installed parallel to the NSX Edges.

Those Edges can be installed in active/standby mode or active/active mode (active/active requires dynamic routing configuration between the BIG-IP appliances and Edges).

Below the Edges, a Distributed Logical Router (DLR) provides connectivity to the different applications tiers: Web, App, and DB. Those internal networks can be VLAN or VXLAN (for flexibility VXLAN is recommended and represented in the diagram below).

The default gateway of the different servers (Web, App, and DC) is the DLR.

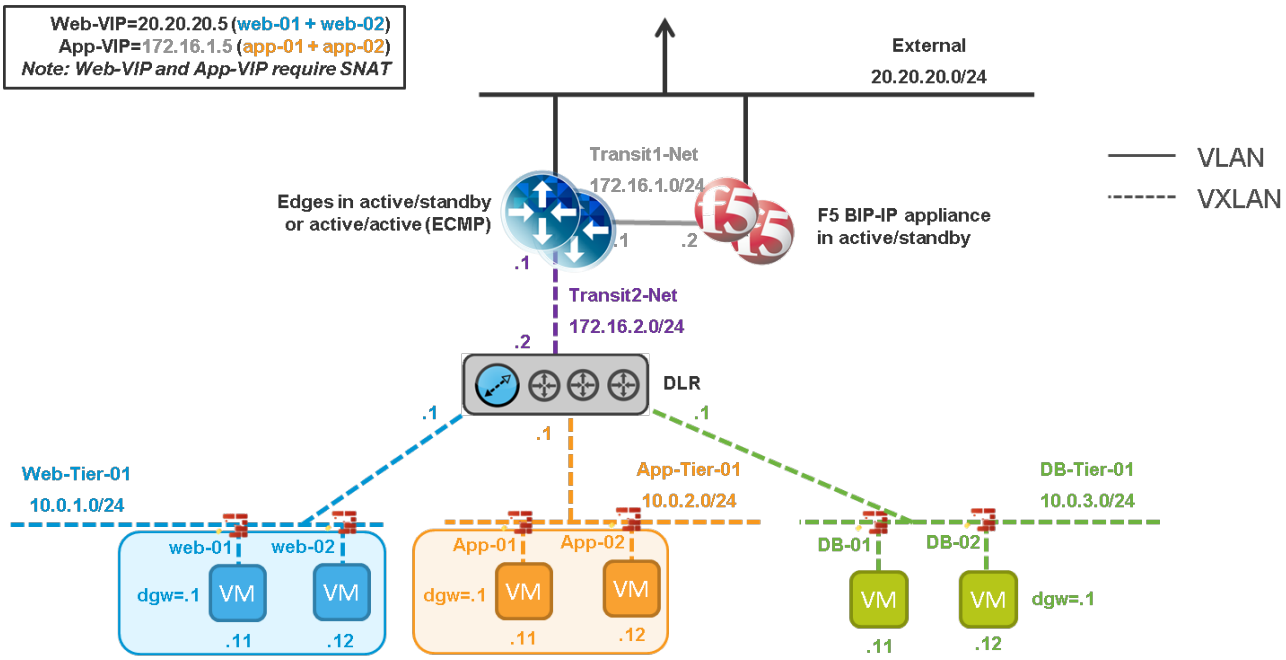

Figure 2– Logical View "Parallel to Edge" with BIG-IP appliances

Following the recommended design (see NSX-v Design Guide here), the recommendation is to physically install the BIG-IP appliances in the Edge Racks where external network is available.

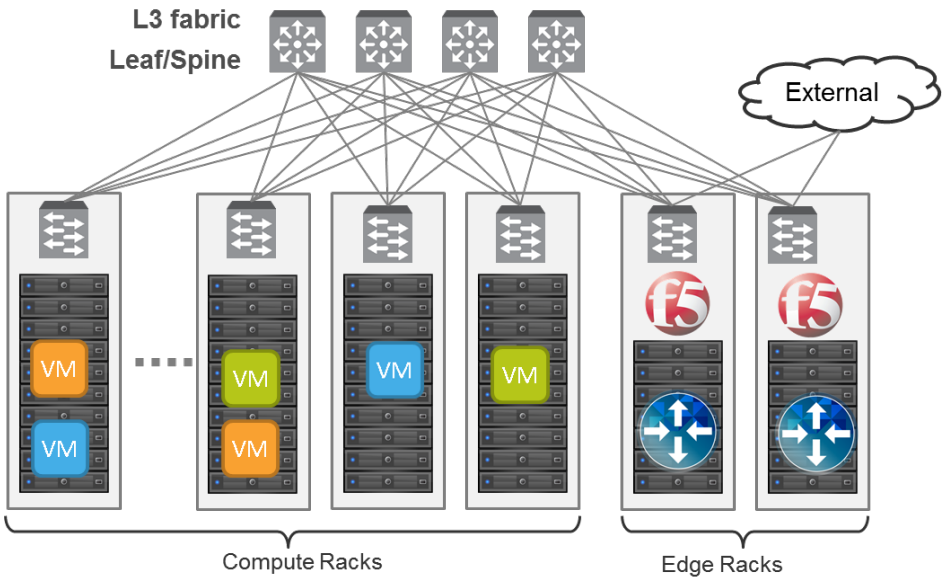

Figure 3– Physical View "Parallel to Edge" with BIG-IP appliances

Traffic Flows

- **North-South Traffic**

  Logical Traffic flows are as follows:
  - From External to BIG-IP appliance.
  - From BIG-IP appliance to Edge to DLR to Web
  - From Web to DLR to Edge to BIG-IP appliance
  - From BIG-IP appliance to Edge to DLR to App
  - From App to DLR to DB

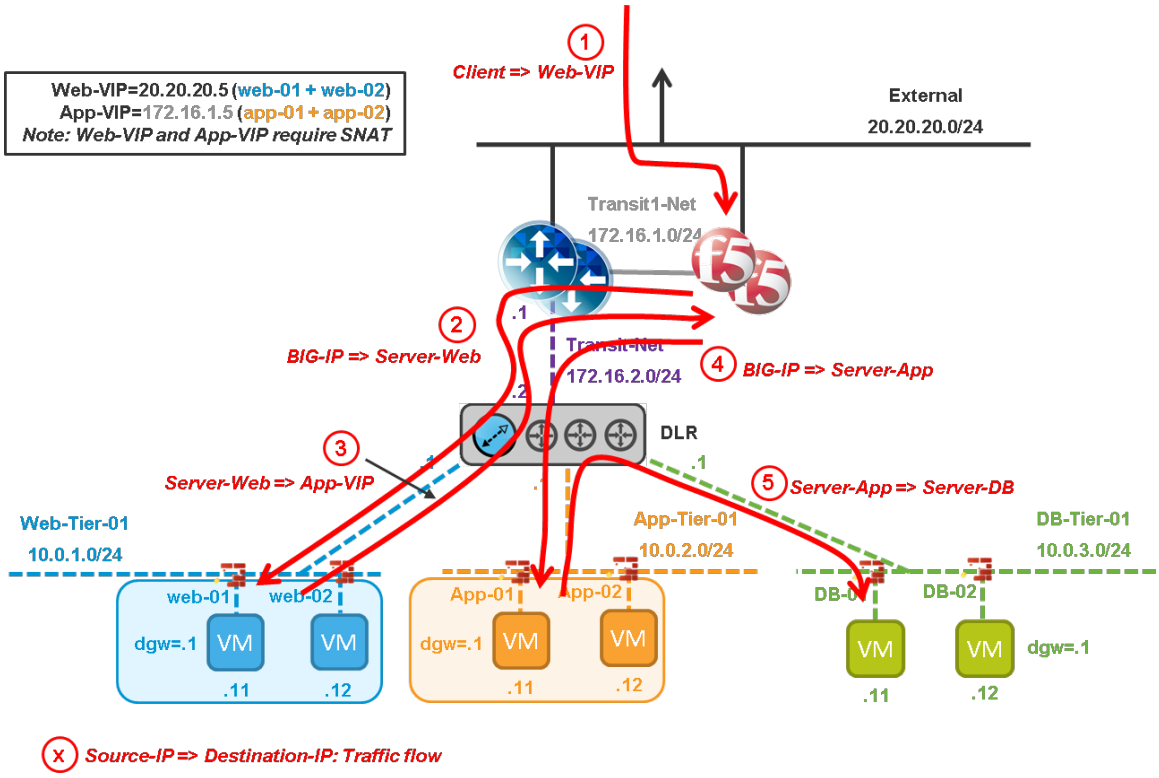

Figure 4 – North-South logical traffic flow "Parallel to Edge" with BIG-IP appliances

Physical Traffic flows as follows:

- From External to BIG-IP appliance
- From BIG-IP appliance to ESXi-hosting-Edge to ESXi-hosting-Web (DLR is not a step since distributed and done here in ESXi-hosting-Edge kernel)
- From ESXi-hosting-Web to ESXi-hosting-Edge to BIG-IP appliance (DLR is not a step since distributed and done here in ESXi-hosting-Web kernel)
- From BIG-IP appliance to ESXi-hosting-Edge to ESXi-hosting-App (DLR is not a step since distributed and done here in ESXi-hosting-Edge kernel)
- From ESXi-hosting-App to ESXi-hosting-DB (DLR is not a step since distributed and done here in ESXi-hosting-App kernel)

Note: BIG-IP and DLR can be configured with dynamic routing (OSPF or BGP) + ECMP. In that case both Edges process traffic. This is represented in figures 6 and 7.

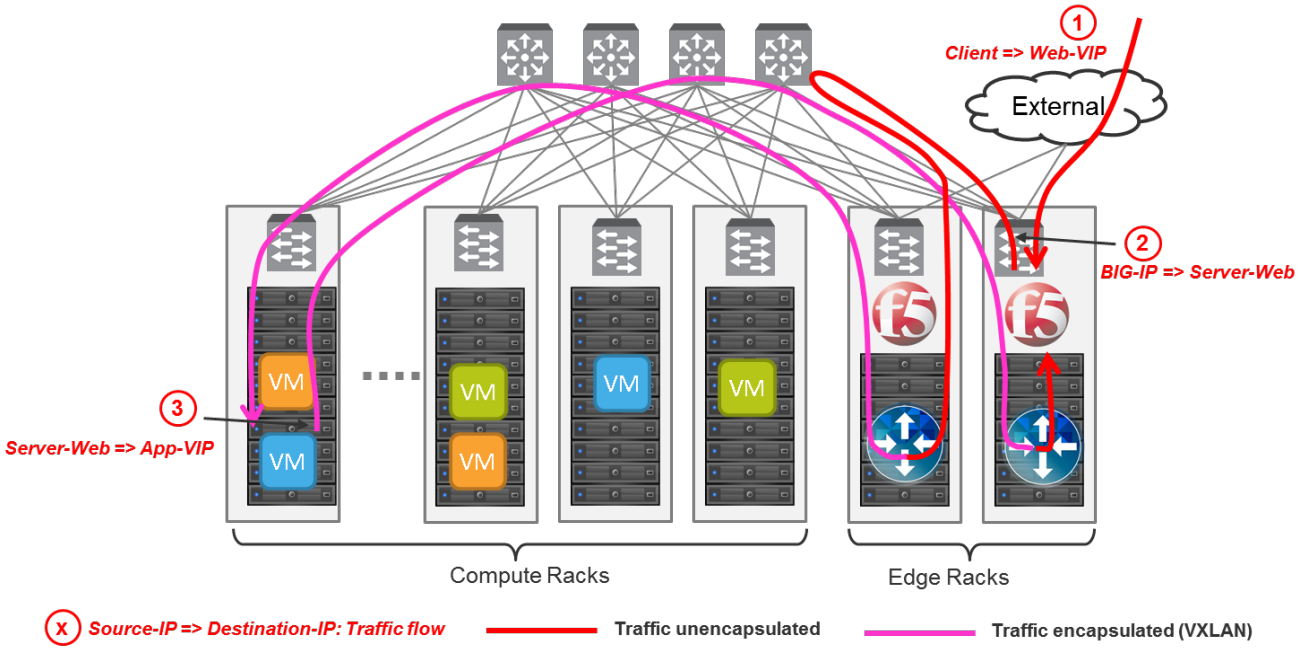

Figure 5 – North-South physical traffic flow 1 "Parallel to Edge" with BIG-IP appliances

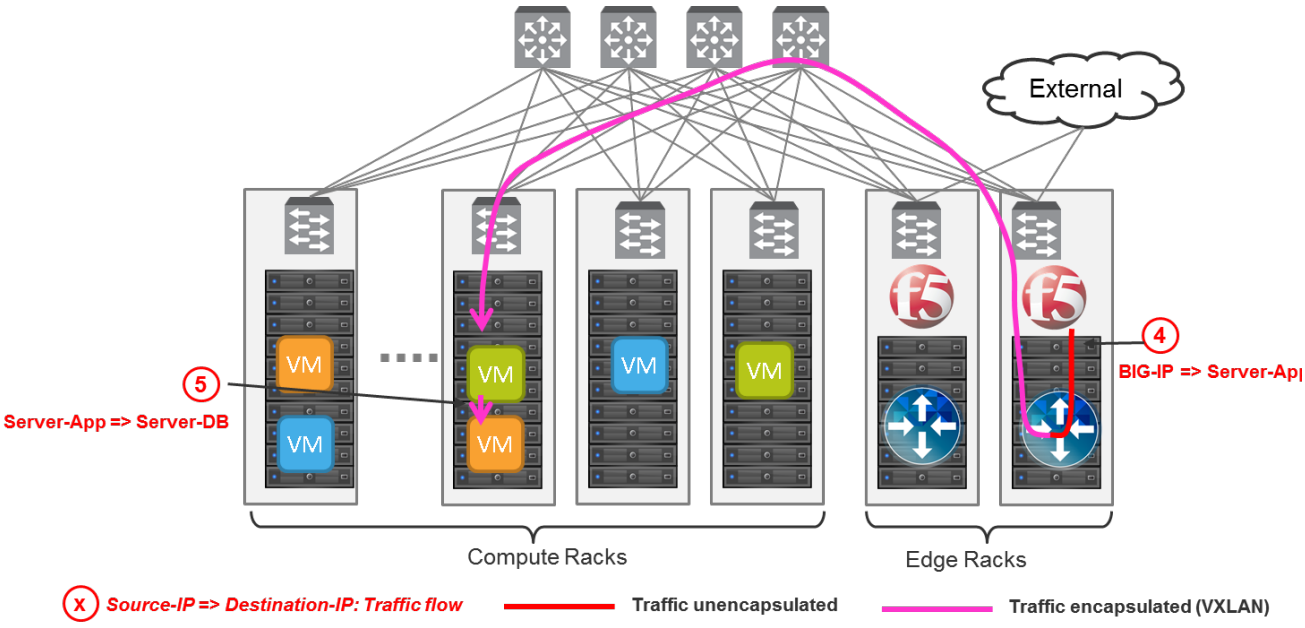

Figure 6 – North-South physical traffic flow 2 "Parallel to Edge" with BIG-IP appliances

- **South-North traffic**
  In this example, a server initiates a security update.

  Logical traffic flows as follows:
  - From Web/App/DB to DLR to Edge to External

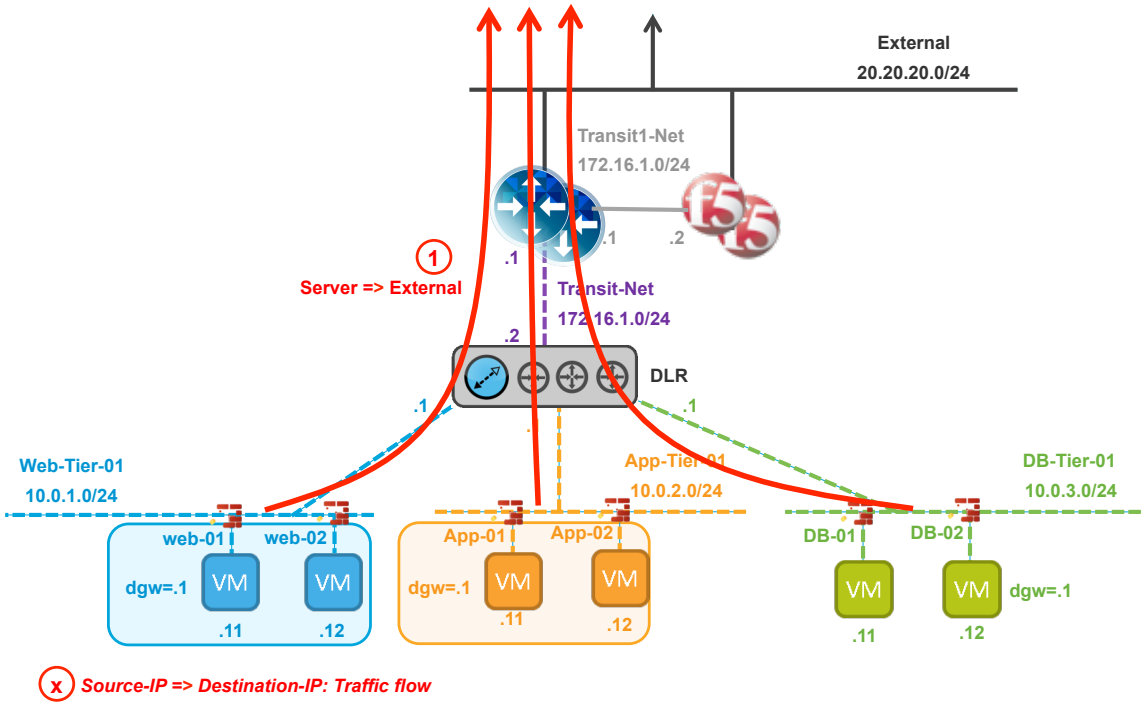

Figure 7 – South-North logical traffic flow "Parallel to Edge" with BIG-IP appliances

Physical traffic flows as follows:
- From ESXi-hosting-Web/App/DB to ESXi-hosting-Edge to External (DLR is not a step since distributed and done here in ESXi-hosting-Web/app/DB kernel)

Note: DLR can be configured with dynamic routing + ECMP. In that case both Edges process traffic. This is represented in figure 8.

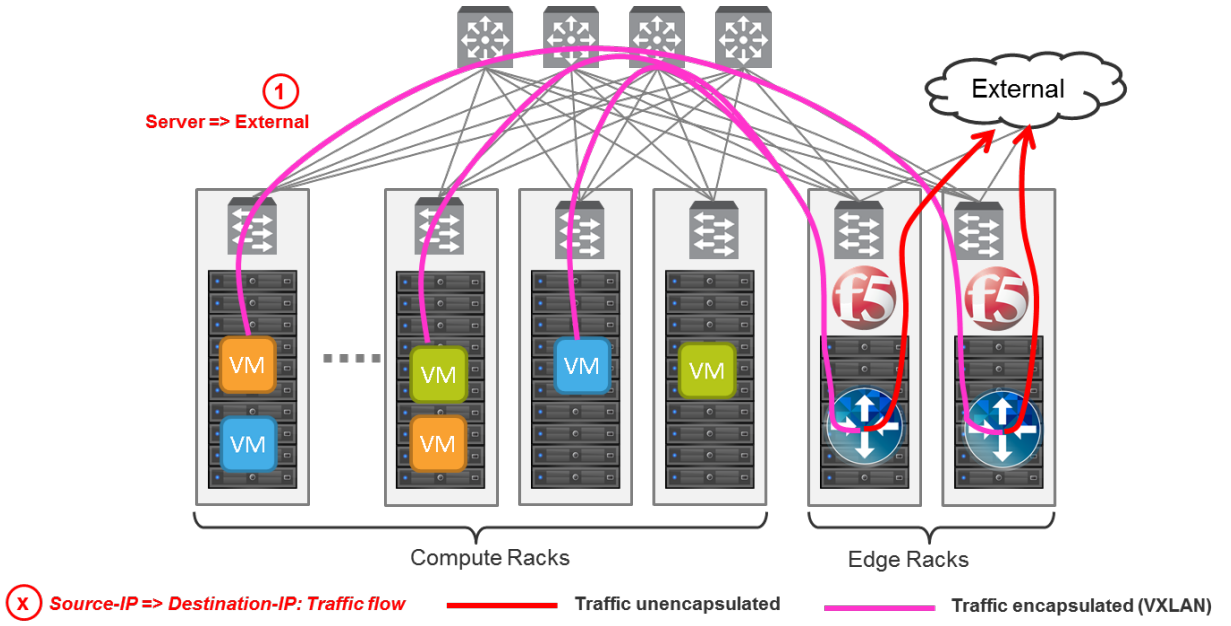

Figure 8 – South-North physical traffic flow "Parallel to Edge" with BIG-IP appliances

VIP requirements

- VIP-Web requires SNAT
- VIP App requires SNAT

Topology #2: Parallel to DLR using VLANs with BIG-IP Appliances

In this topology, the BIG-IP appliances are placed parallel to the DLR. In this topology, the BIG-IP appliances are used but BIG-IP VE could also be used the same way.

This topology requires a Layer 2 physical fabric (external network is available in all racks).

This topology is represented by the magenta lines in figure 1.

Logical and Physical view

BIG-IP appliances are logically installed parallel to the DLR.

The DLR provides connectivity to the different applications tiers: Web, App, and DB. Those internal networks could be VLAN or VXLAN, but in this topology we are describing a non-VXLAN overlay scenario. Standard 802.1q tagging is used to connect the VLAN-backed virtual networking segments to the physical networking segments to which the BIG-IPs are connected

The default gateway of the different servers (Web, App, and DC) is the DLR.

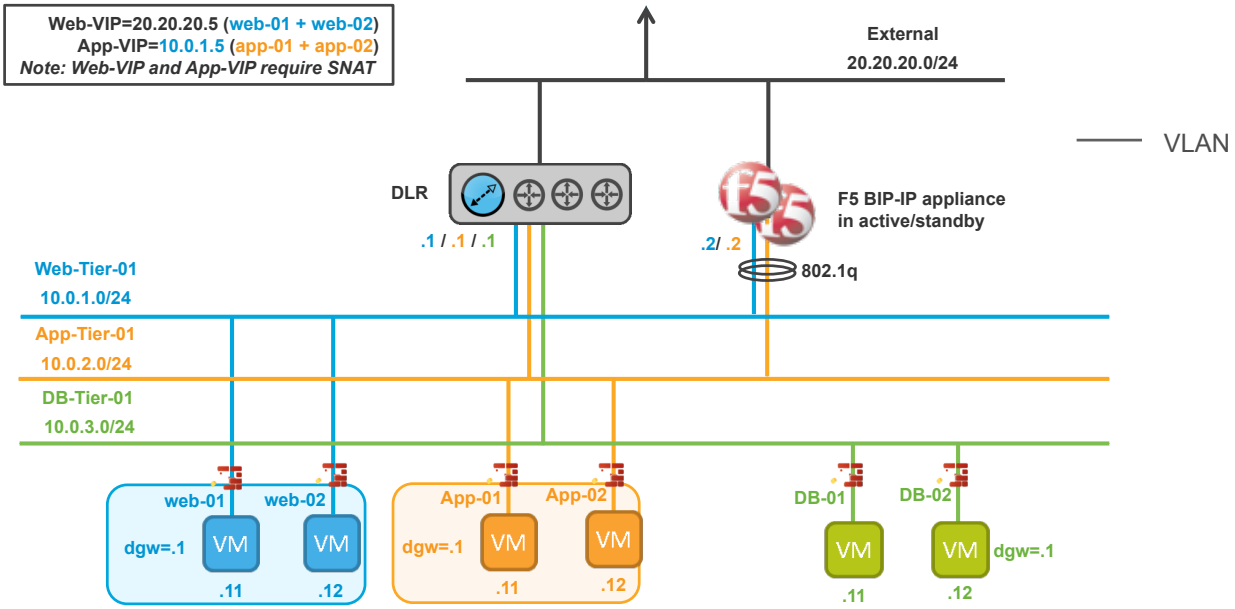

Figure 9 – Logical View "Parallel to DLR" with BIG-IP appliances

The recommended topology has the BIG-IP appliances physically connected to the top distribution physical routers.

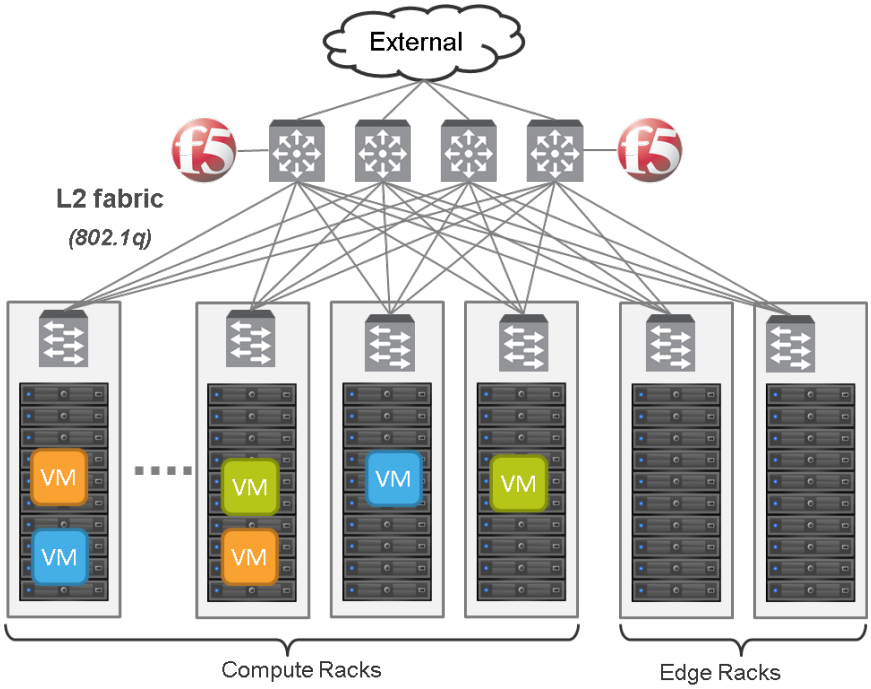

Figure 10 – Physical View "Parallel to DLR" with BIG-IP appliances

Traffic flows

- **North-South traffic**

Logical traffic flows as follows:

- From External to BIG-IP appliance.
- From BIG-IP appliance to Web
- From Web to BIG-IP appliance
- From BIG-IP appliance to App
- From App to DLR to DB

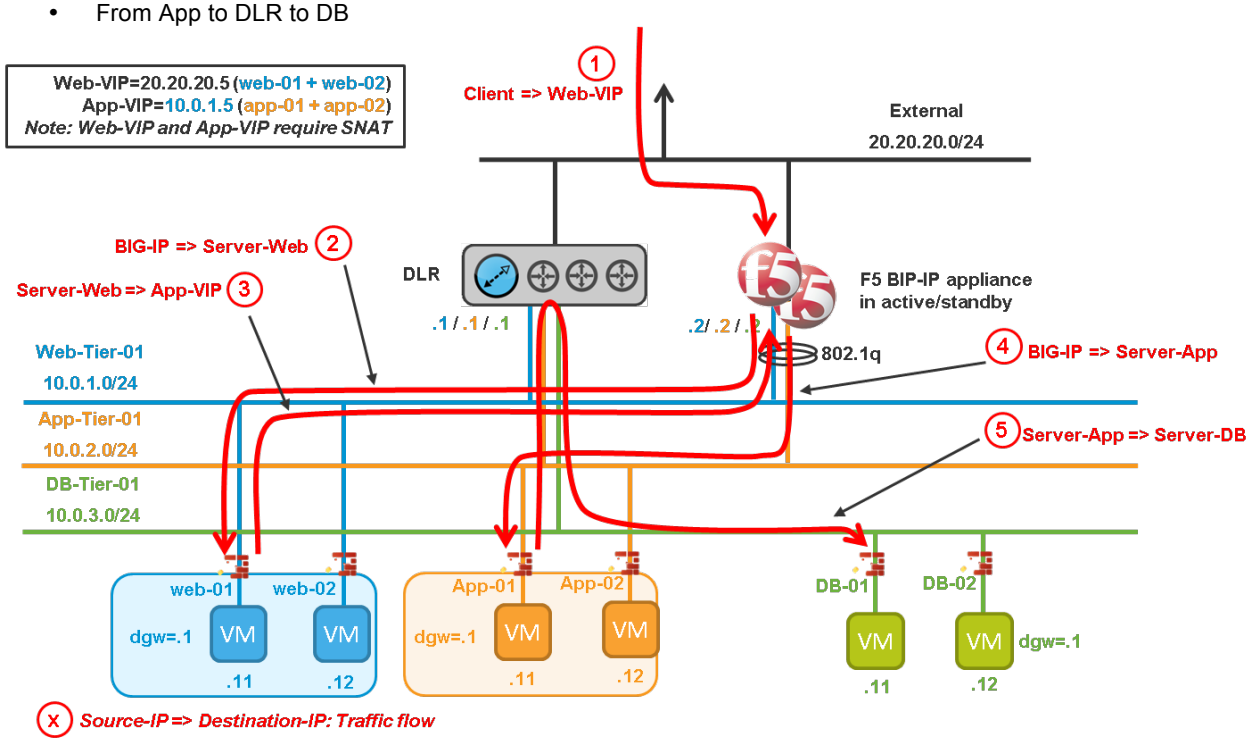

Figure 11 – North-South logical traffic flow "Parallel to DLR" with BIG-IP appliances

Physical traffic flows as follows:

- From External to BIG-IP appliance
- From BIG-IP appliance to ESXi-hosting-Web
- From ESXi-hosting-Web to BIG-IP appliance
- From BIG-IP appliance to ESXi-hosting-App
- From ESXi-hosting-App to ESXi-hosting-DB (DLR is not a step since distributed and done here in ESXi-hosting-App kernel)

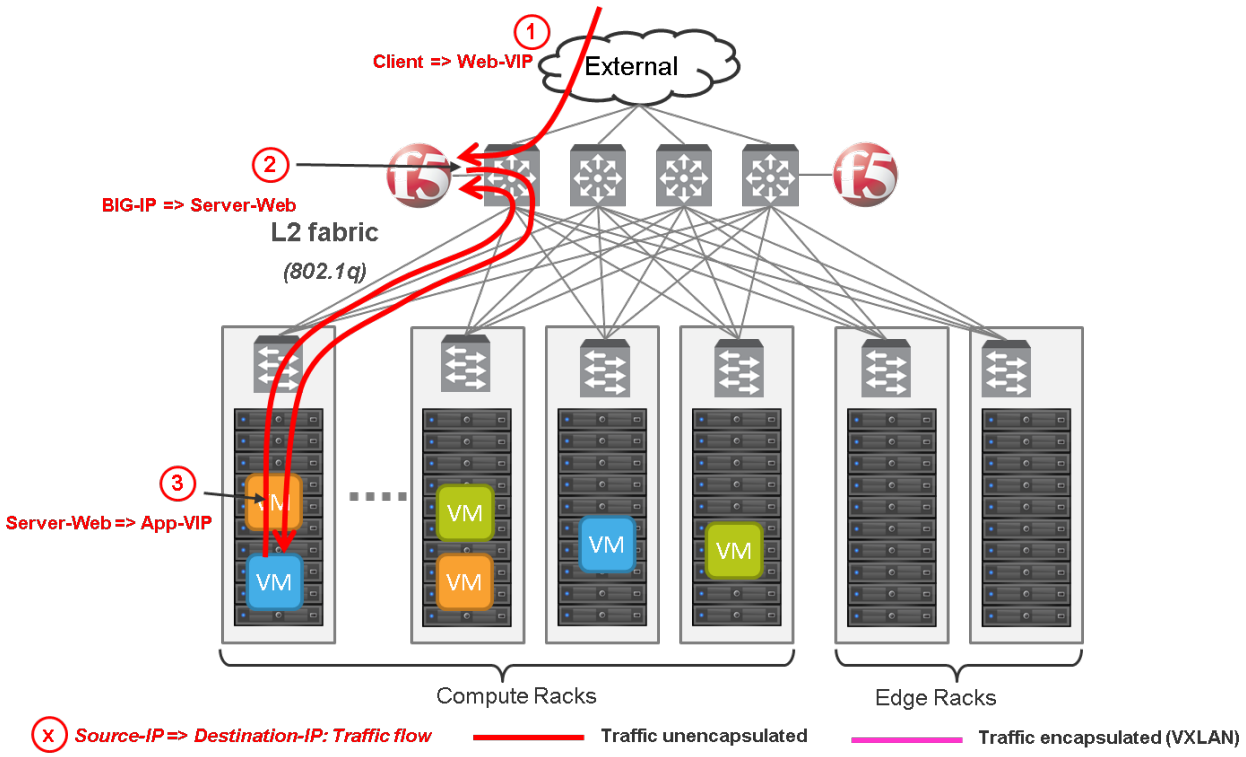

Figure 12 – North-South physical traffic flow1 "Parallel to DLR" with BIG-IP appliances

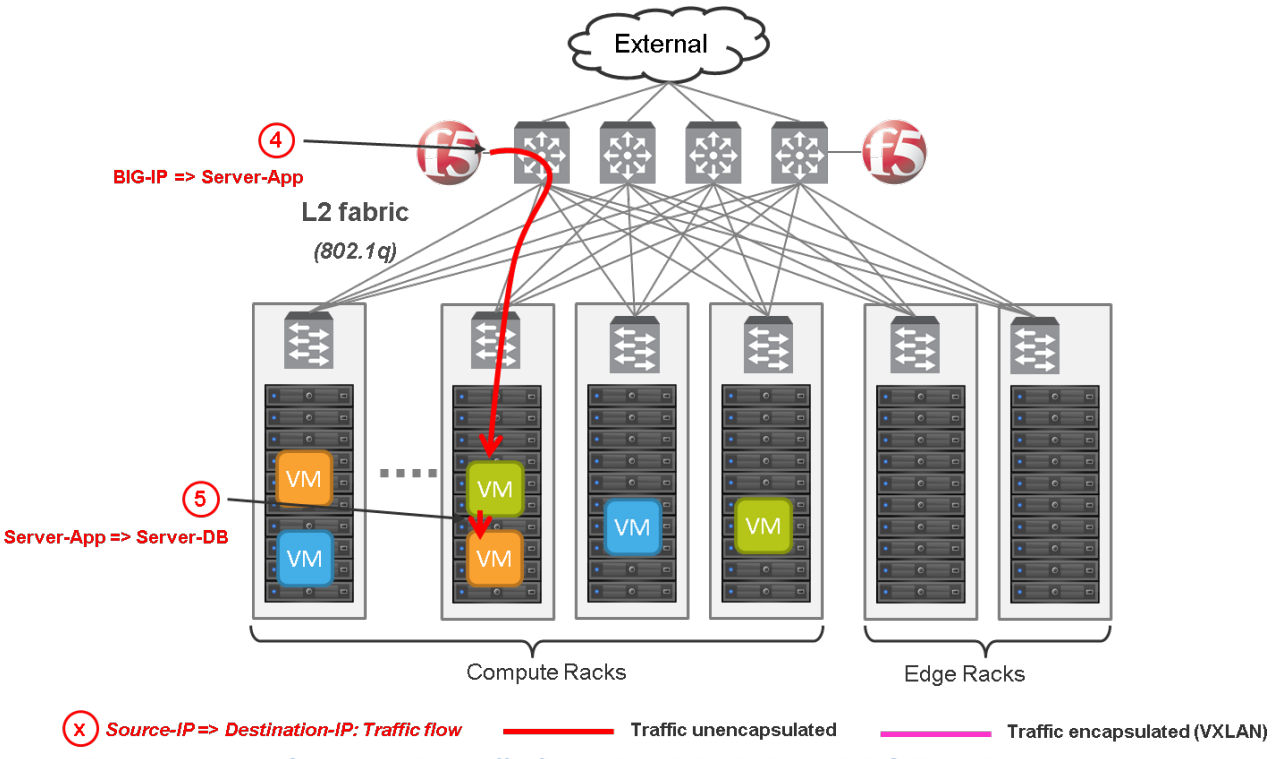

Figure 13 – North-South physical traffic flow2 "Parallel to DLR" with BIG-IP appliances

- **South-North traffic**

  In this example, the server initiates a security update.

  Logical traffic flows as follows:
  - From Web/App/DB to DLR to External

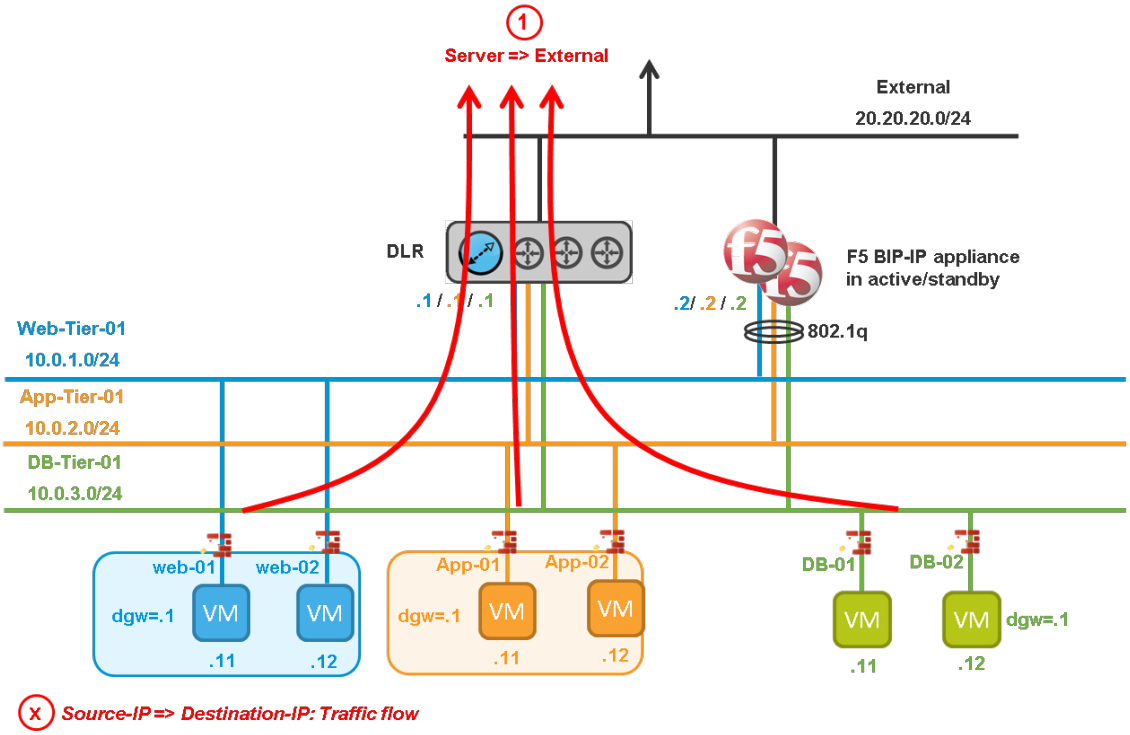

Figure 14 – South-North logical traffic flow "Parallel to DLR" with BIG-IP appliances

Physical traffic flows as follows:

- From ESXi-hosting-Web/App/DB to External (DLR is not a step since distributed and done here in ESXi-hosting-Web/app/DB kernel)

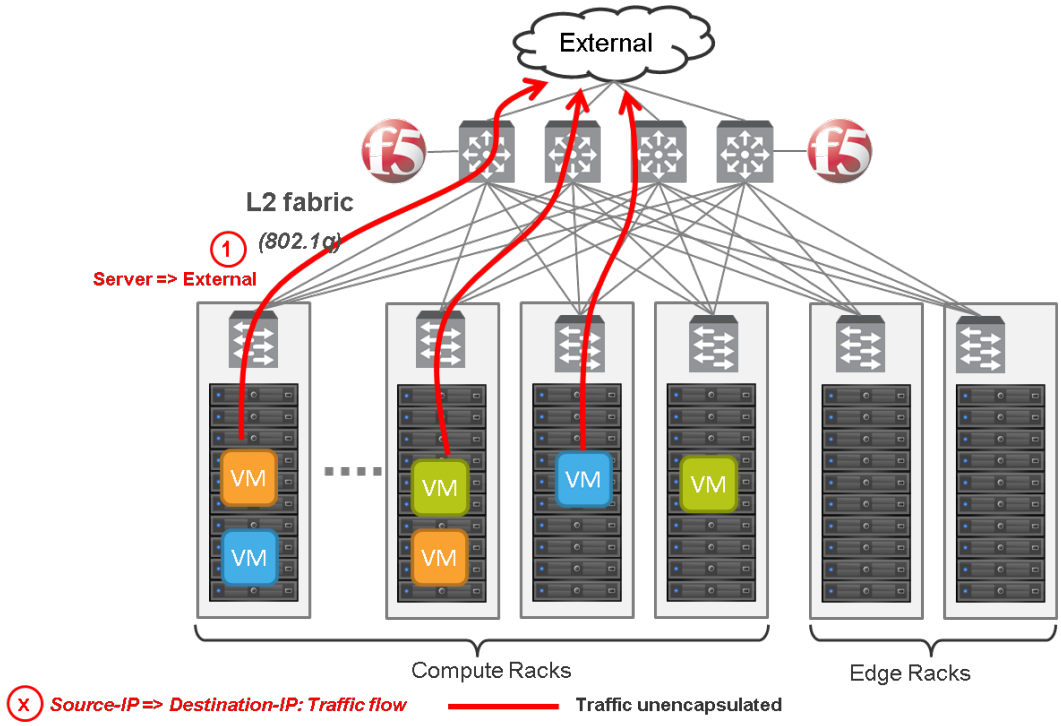

Figure 15 – South-North physical traffic flow "Parallel to DLR" with BIG-IP appliances

VIP requirements

- VIP-Web requires SNAT
- VIP App requires SNAT

Topology #3: One-Arm connected using VXLAN overlays with BIG-IP VE

In this topology, the BIG-IP VEs are placed in one-arm mode in the load-balanced server networks.

This topology is popular on Layer 3 physical fabrics, such as Leaf/Spine but also works on Layer 2 physical fabrics.

Logical and Physical view

At the top, Edges can be installed in active/standby mode or active/active mode (active/active requires dynamic routing configuration between the DLR and Edges).

Below the Edges, a Distributed Logical Router (DLR) provides connectivity to the different applications tiers: Web, App, and DB. Those internal networks can be VLAN or VXLAN (for flexibility VXLAN is recommended and represented in the diagram below).

BIG-IP VEs are logically installed in one-arm mode on the load-balanced tiers' logical switches.

The default gateway of the different servers (Web, App, and DC) is the DLR.

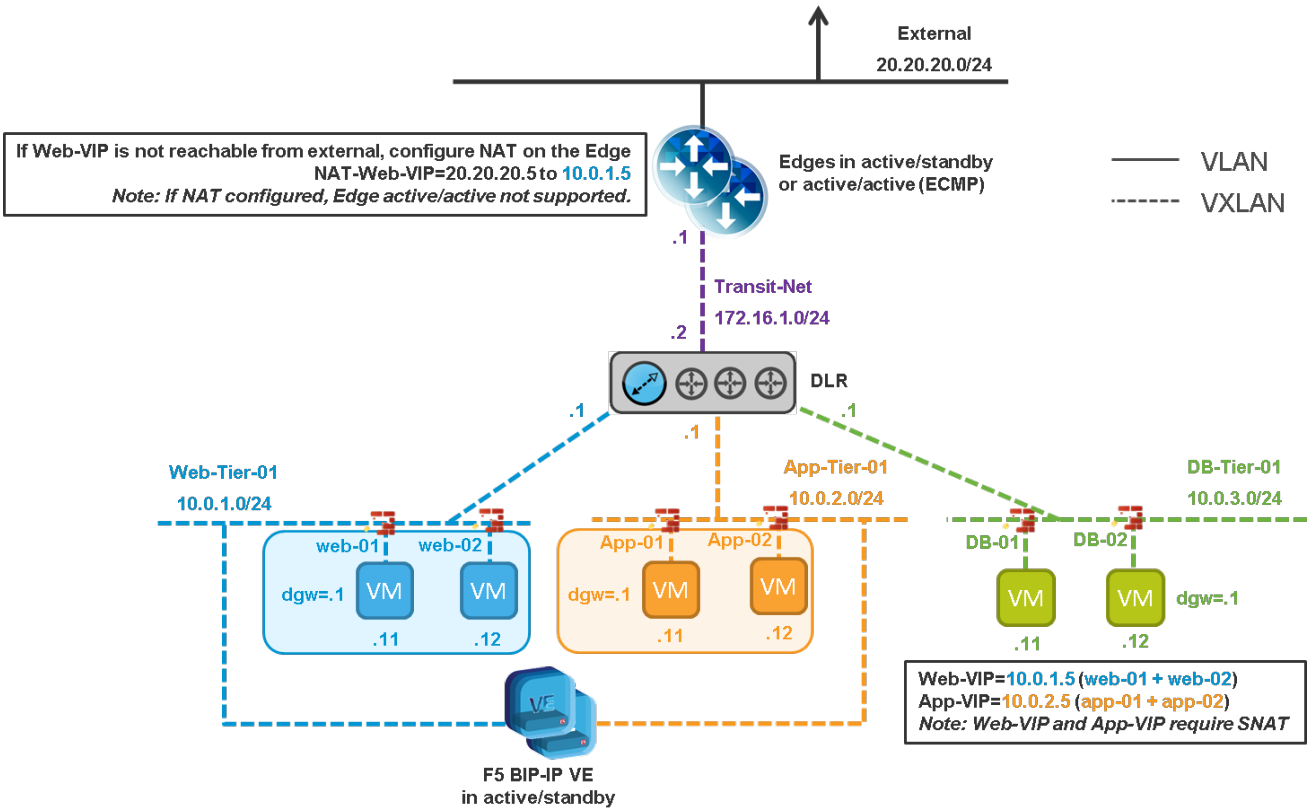

Figure 16 – Logical View "One-Arm connected" with BIG-IP VE

The recommended topology has the BIG-IP VEs deployed in the compute rack close to the load-balanced servers.

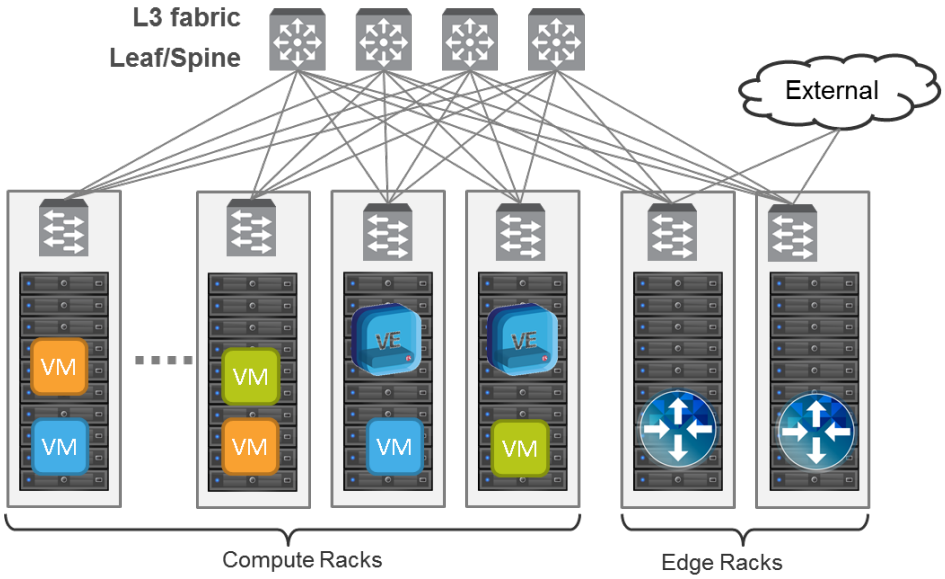

Figure 17– Physical View "One-Arm connected" with BIG-IP VE

Traffic flows

- **North-South traffic**

  Logical traffic flows as follows:
  - From External to Edge to DLR to BIG-IP VE
  - From BIG-IP VE to Web
  - From Web to BIG-IP VE
  - From BIG-IP VE to App
  - From App to DLR to DB

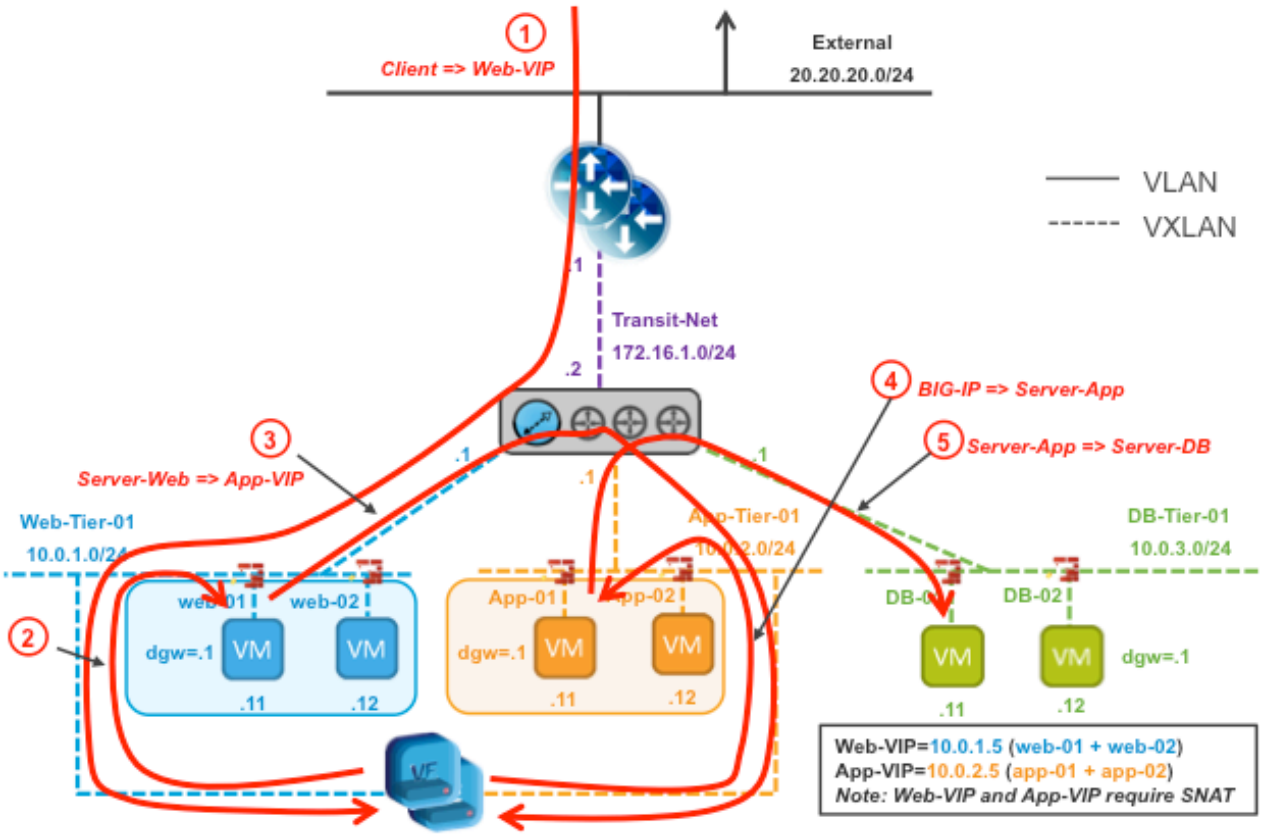

Figure 18 – North-South logical traffic flow "One-Arm connected" with BIG-IP VE

Physical traffic flows as follows:

- From External to ESXi-hosting-Edge
- From ESXi-hosting-Edge to ESXi-hosting-BIG-IP-VE (DLR is not a step since distributed and done here in ESXi-hosting-Edge kernel)
- From ESXi-hosting-BIG-IP-VE to ESXi-hosting-Web
- From ESXi-hosting-Web to ESXi-hosting-BIG-IP-VE (DLR is not a step since distributed and done here in ESXi-hosting-Web kernel)
- From ESXi-hosting-BIG-IP-VE to ESXi-hosting-App
- From ESXi-hosting-App to ESXi-hosting-DB (DLR is not a step since distributed and done here in ESXi-hosting-App kernel)

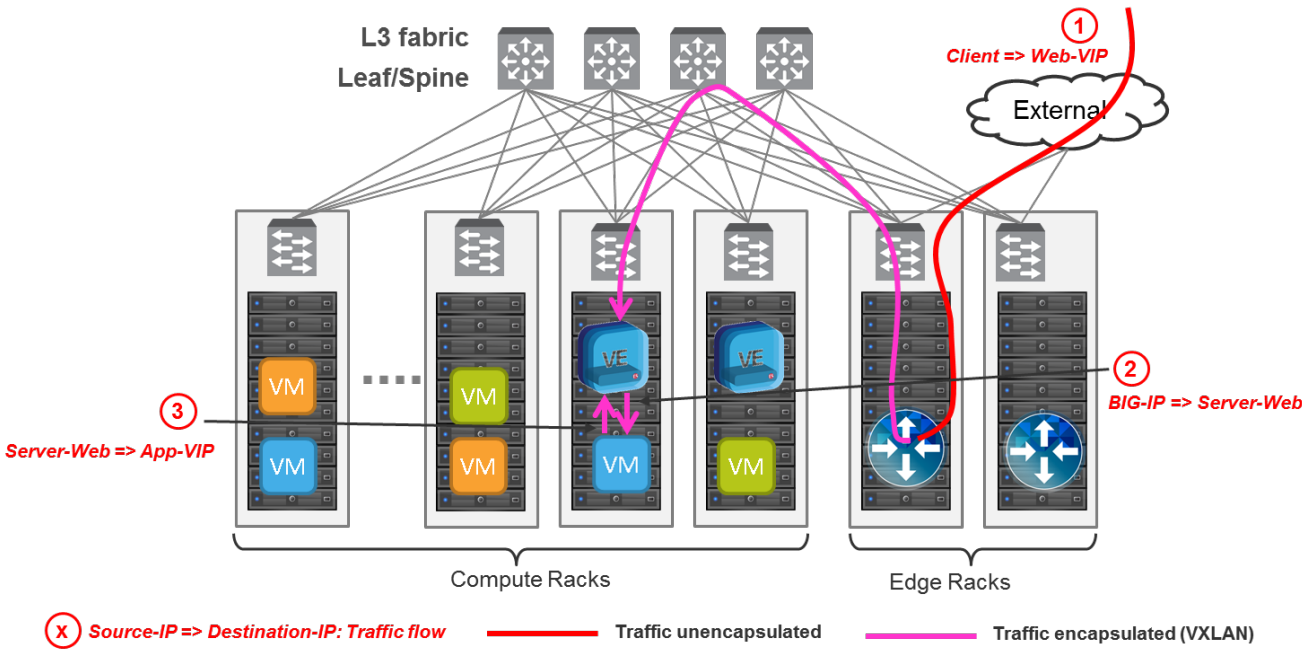

Figure 19 – North-South physical traffic flow1 "One-Arm connected" with BIG-IP VE

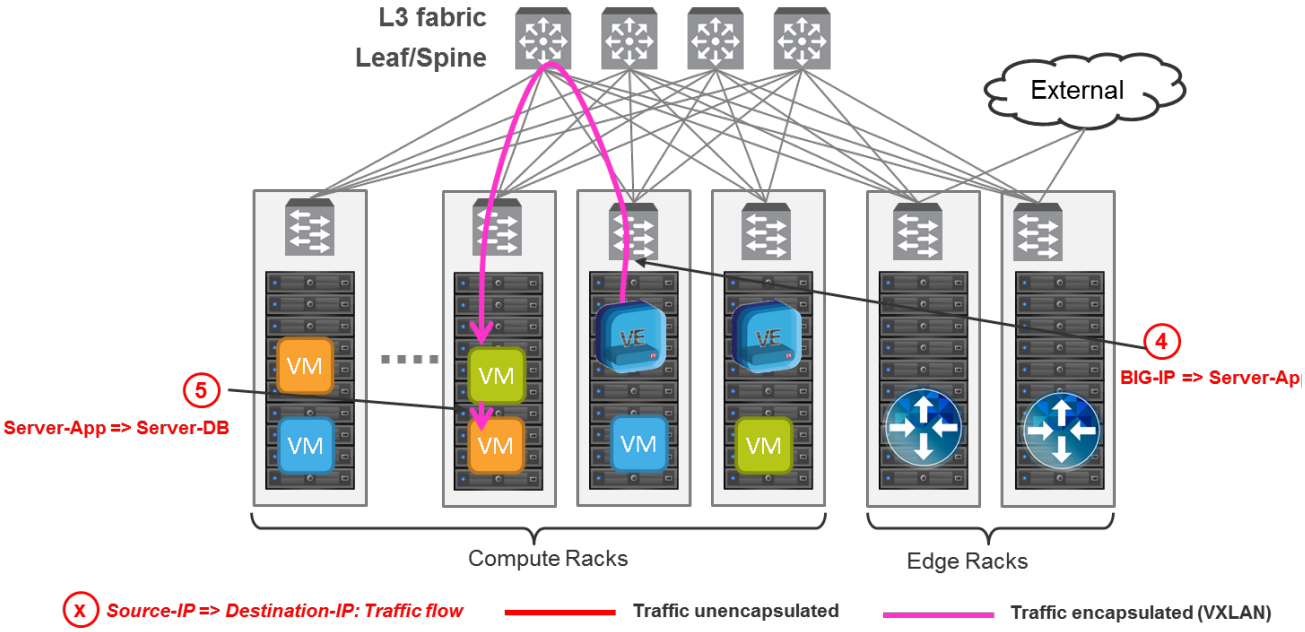

Figure 20 – North-South physical traffic flow2 "One-Arm connected" with BIG-IP VE

- **South-North traffic**

  In this example, a server initiates a security update.

  Logical traffic flows as follows:

  - From Web/App/DB to DLR to Edge to External

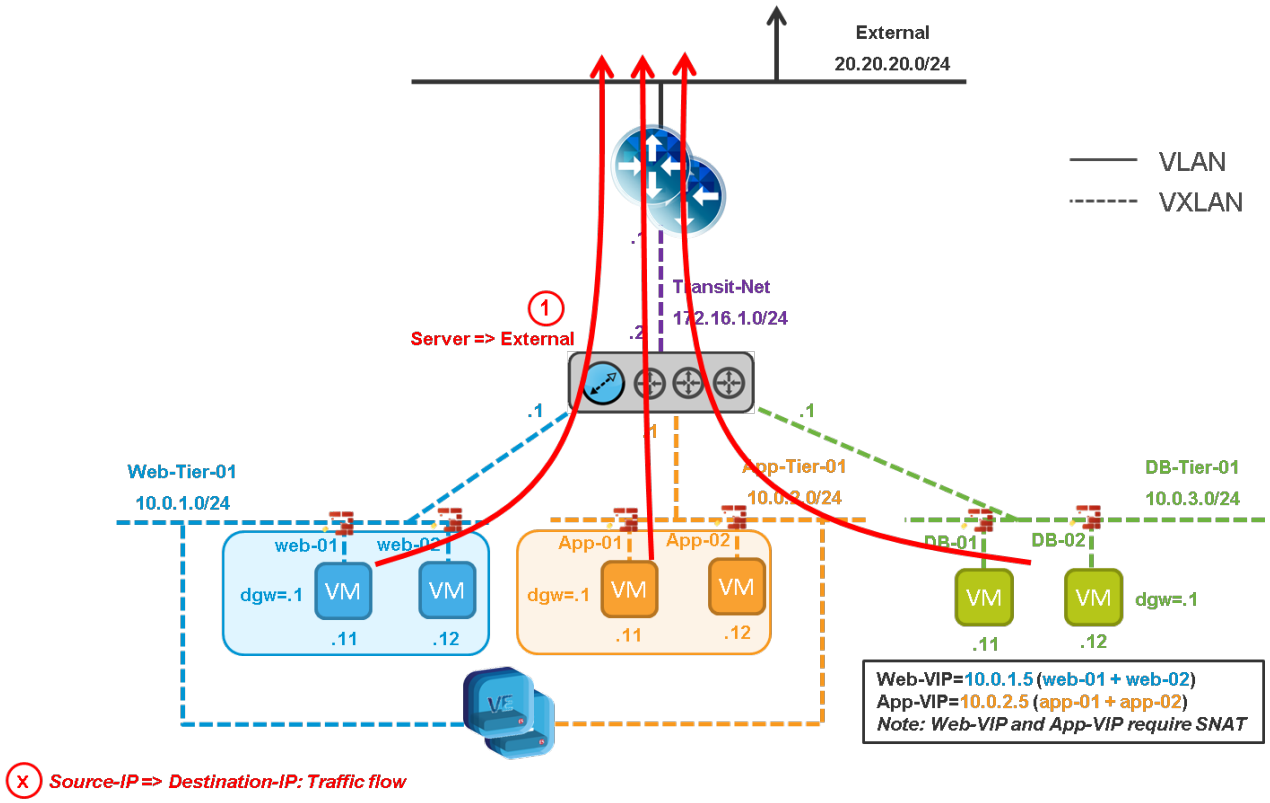

Figure 21 – South-North logical traffic flow "One-Arm connected" with BIG-IP VE

Physical traffic flows as follows:

- From ESXi-hosting-Web/App/DB to ESXi-hosting-Edge to External (DLR is not a step since distributed and done here in ESXi-hosting-Web/app/DB kernel)

Note: DLR can be configured with dynamic routing + ECMP. In that case both Edges process traffic. This is represented in figure 23.

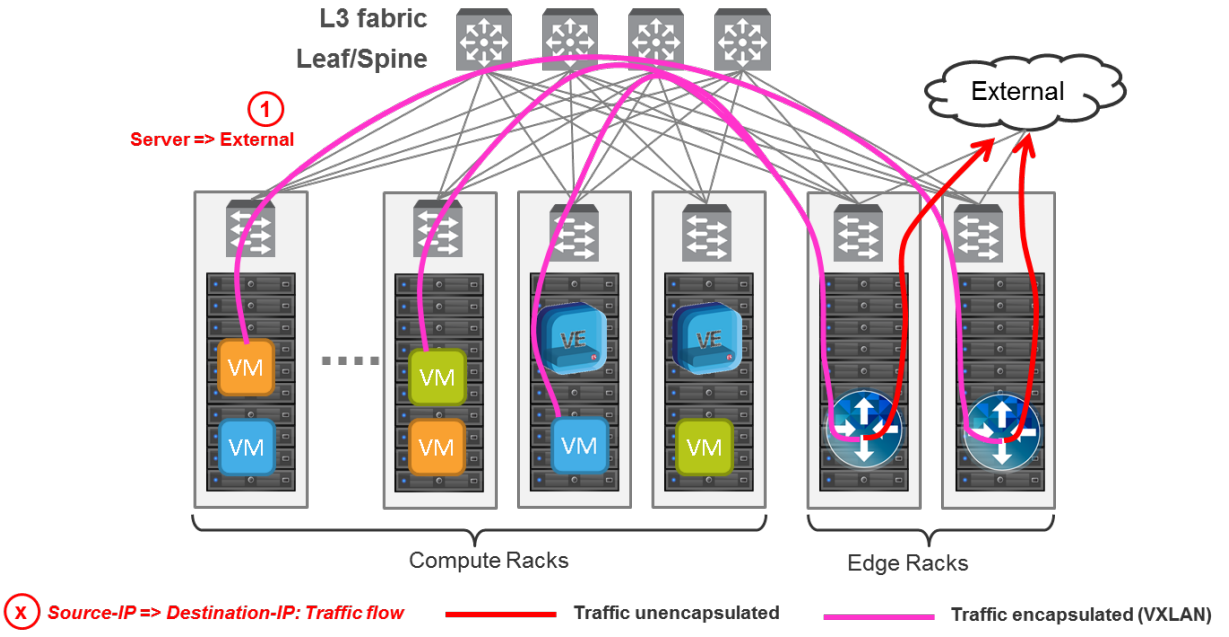

Figure 22 – South-North physical traffic flow "One-Arm connected" with BIG-IP VE

VIP requirements

- VIP-Web requires SNAT
- VIP App requires SNAT

Pros/Cons of the Alternative Topologies

Parallel to Edge	Pros	Cons
BIG-IP appliances or VE	• Easy to implement • Supports very large throughput (multiple 10Gbps) with BIG-IP appliances • East-West non load-balanced traffic is DLR-optimized • Offers DFW for security even within the same subnet • Works on Layer 2 or Layer 3 fabrics	• East-West load-balanced traffic is not DLR-optimized (goes through Edge) • VIP-Web and VIP-App require SNAT

Parallel to DLR	Pros	Cons
BIG-IP appliances or VE	• Supports very large throughput (multiple 10Gbps) with BIG-IP appliances • East-West non load-balanced traffic is DLR-optimized • Offers DFW for security even within the same subnet	• VIP-Web and VIP-App require SNAT • Requires Layer 2 fabric (requires all VLANs on all ESXi hosts)

One-Arm Connected	Pros	Cons
BIG-IP VE only	• Easy to implement • East-West traffic between non load-balanced network is DLR optimized • Offers DFW for security even within same subnet • Works on Layer 2 or Layer 3 fabric	• VIP-Web and VIP-App require SNAT

Note: Topologies that do not require SNAT on the VIP are detailed in the section under “Alternative Topologies”.

Alternative Topologies

Alternative Topology A: “On Top of Edge with BIG-IP appliances / VE”

This topology is similar to Topology #1, but with the BIG-IP appliances or VE on top of the Edge.

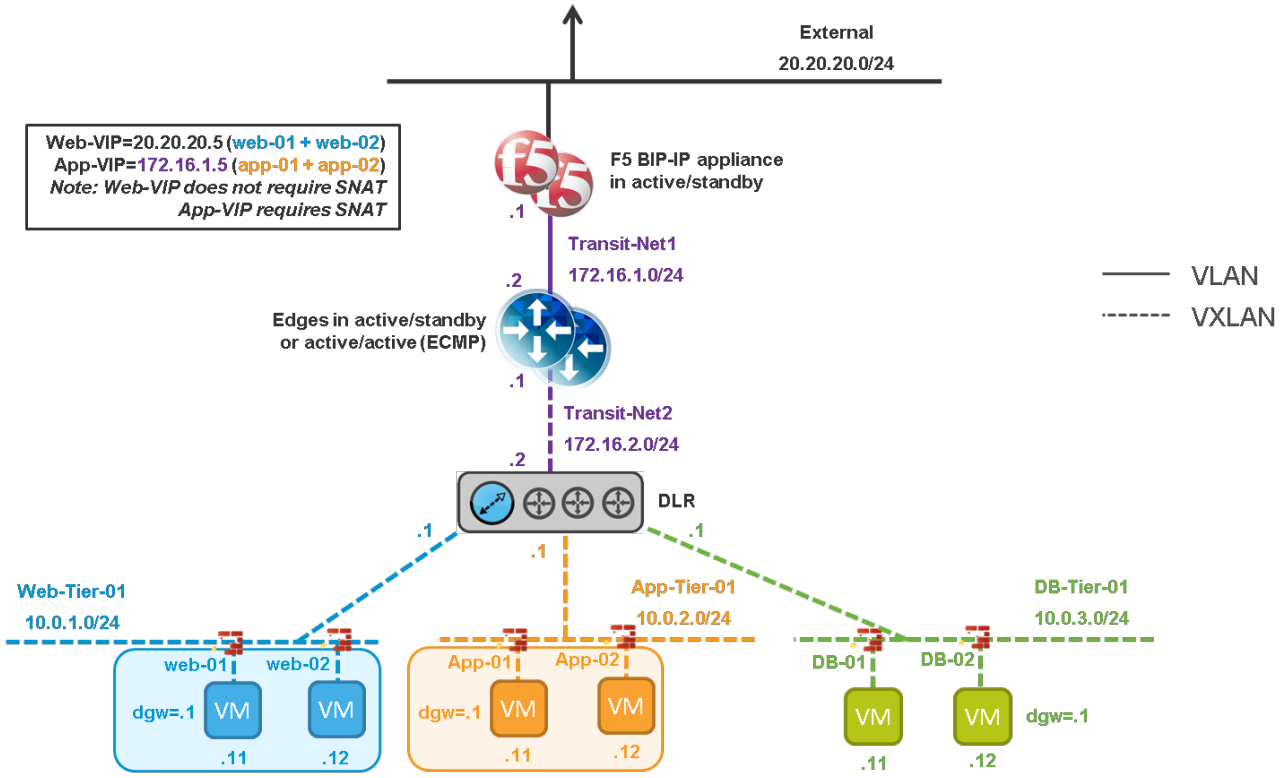

Figure 23 – Logical View “On Top of Edge” with BIG-IP appliances

This offers the benefit of not requiring SNAT for the Web-VIP.

All North/South traffic (even non load-balanced traffic) crosses the BIG-IP.

Alternative Topology B:
"On Top of DLR with BIG-IP appliances / VE"

This topology is similar to the Topology #2, but with the BIG-IP appliances or VE on top of the DLR.

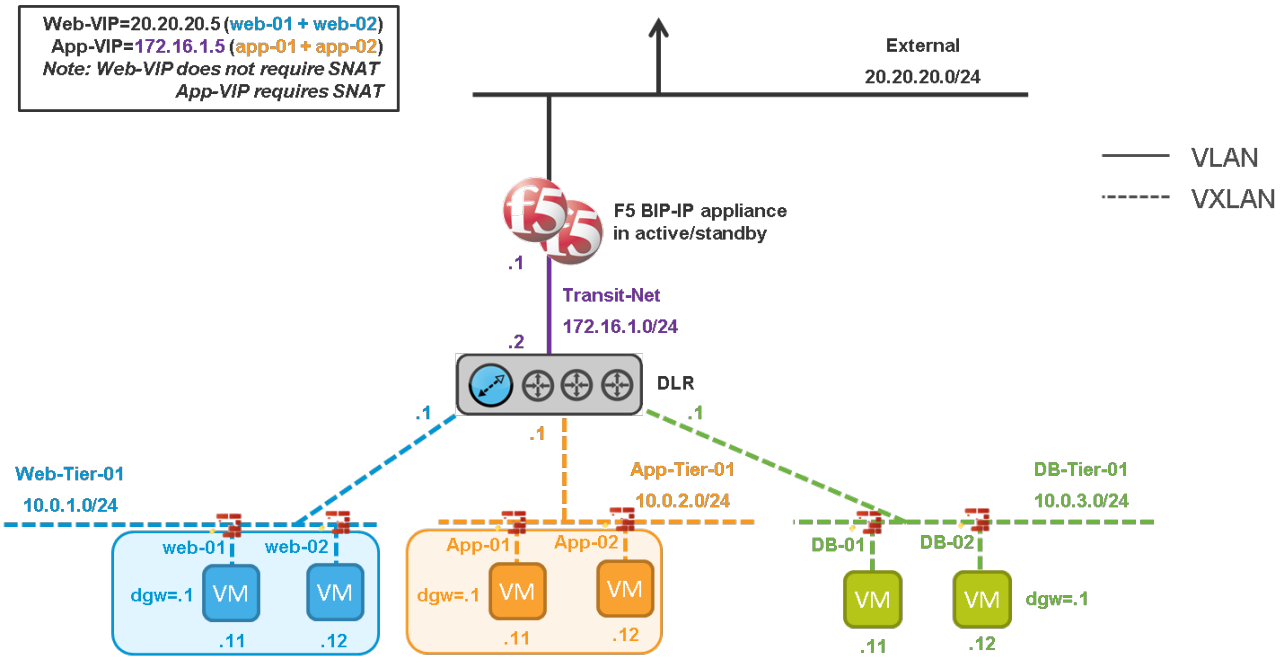

Figure 24 – Logical View "On Top of DLR" with BIG-IP appliances

This offers the benefit of not requiring SNAT for the Web-VIP.

All North/South traffic (even non load-balanced traffic) crosses the BIG-IP.

Conclusion

This document showcases several possible NSX and BIG-IP interoperability scenarios and the network topologies to accomplish those scenarios.

F5 and VMware are working on a jointly developed API integration between NSX and BIG-IQ. This will enable IT organizations to fully leverage the combined strengths of NSX virtualization and automation with richer application delivery services enabled by F5 BIG-IP. This planned NSX/F5 integration will allow users to configure BIG-IP settings (e.g. pools, VIPs, iApps) from NSX. The integration will also allow for automated BIG-IP VE deployment, licensing, and configuration.

Many of the scenarios depicted in this document will be deployable using this upcoming integration.

For more information these solutions, please contact your local F5 or VMware representative.