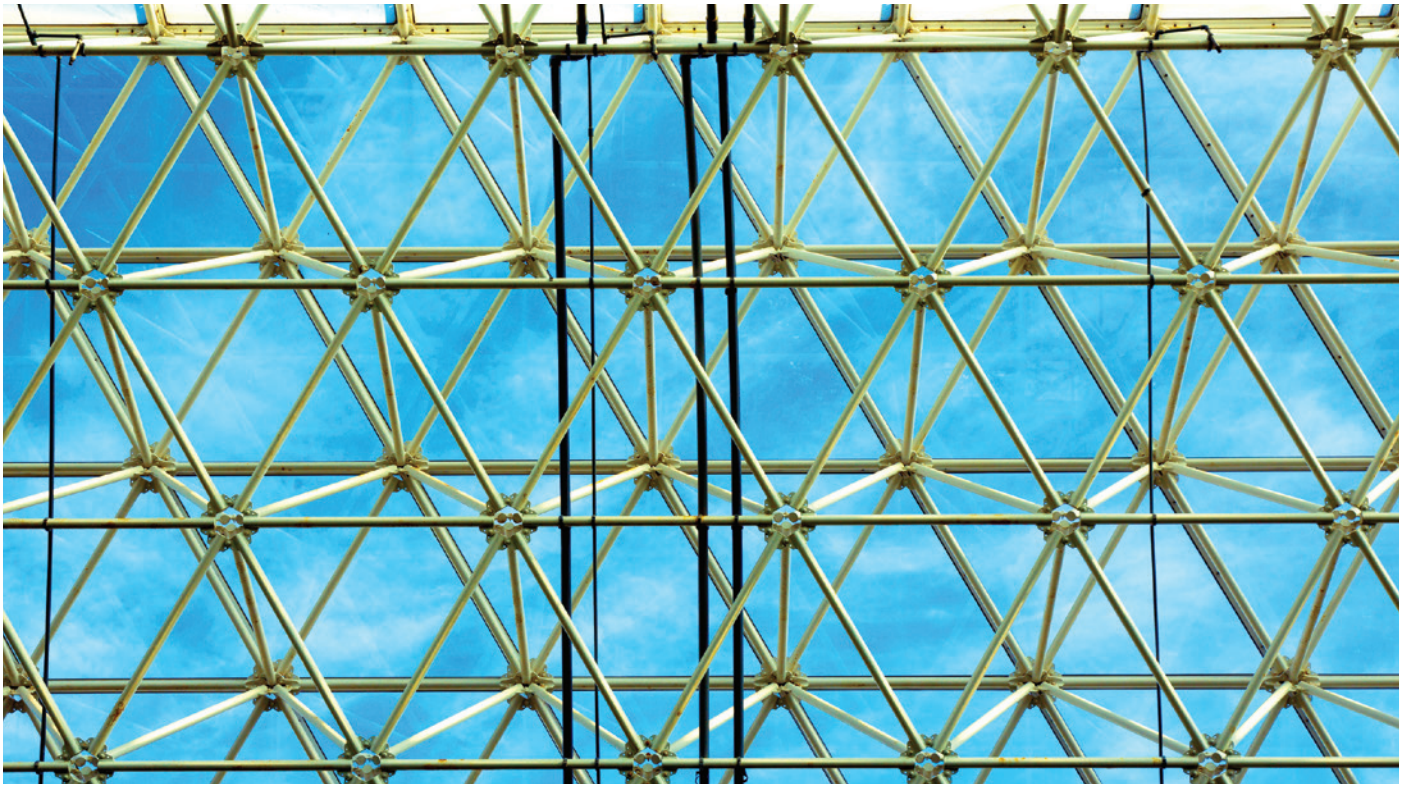

VXLAN GATEWAY USE CASE:

CONNECTING SERVERS WITH THE OPEN VIRTUAL SWITCH AND OMNISWITCH VXLAN GATEWAY

APPLICATION NOTE

Alcatel·Lucent Enterprise

EXECUTIVE SUMMARY

Open Virtual Switch (OVS) is a popular software virtual switch that runs on servers along with a hypervisor, offering a flexible way to connect virtual machines (VMs) to various networks. OVS can connect VMs together with Virtual Extensible LAN (VXLAN) tunnels to build private virtual networks. Because of this, many cloud management systems such as OpenStack® support OVS. The Alcatel-Lucent OmniSwitch OS6900-Q32 VXLAN gateway functionality can easily connect virtual machines that are using VXLAN to external networks.

This paper covers the configuration of a VXLAN network with servers running OVS and a OS6900-Q32 as the gateway. While there are many software suites (such as OpenStack) that can automate the steps described in this document, these are the basic building blocks of the VXLAN network with which users should be familiar.

EXAMPLE NETWORK

The following diagram shows the components of the example VXLAN network.

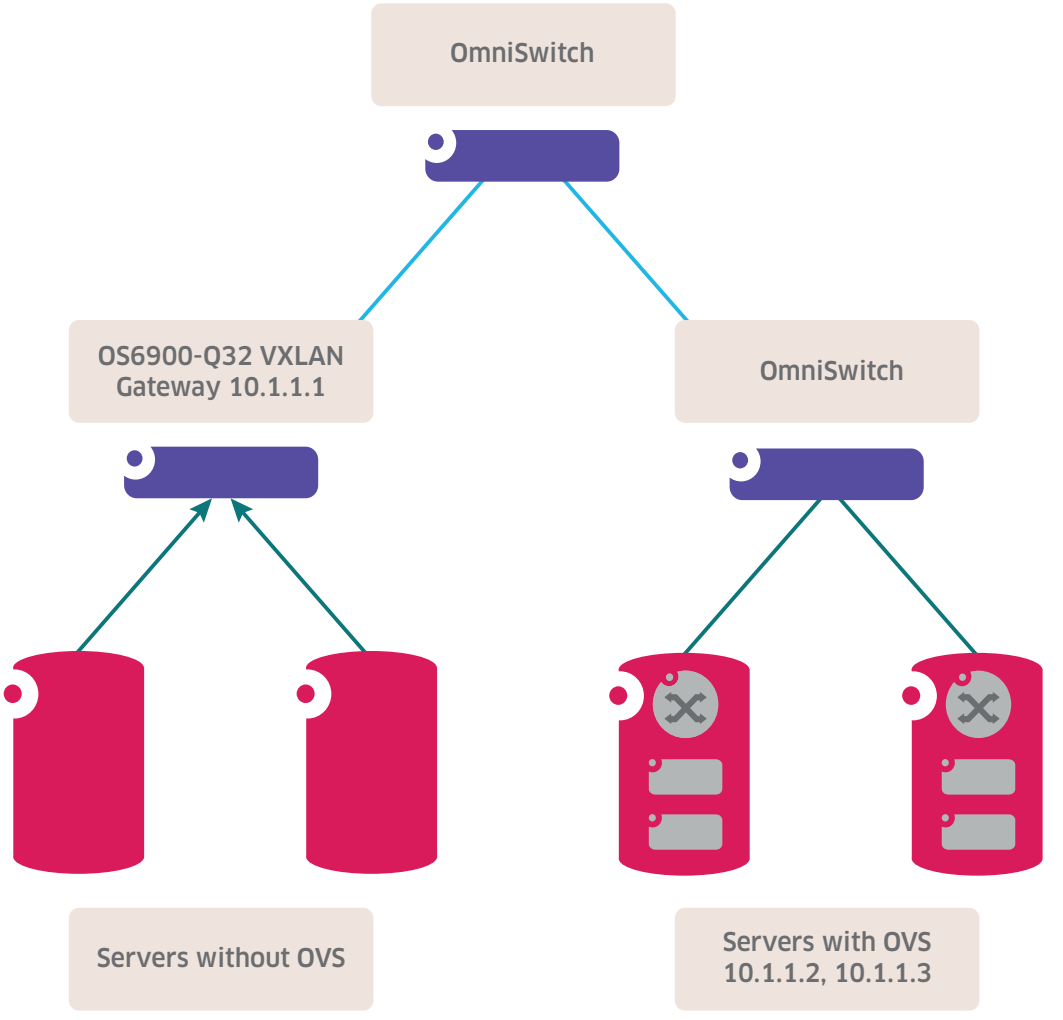

In the above diagram there are three OmniSwitch devices that are configured as routers, with one of them enabled as a VXLAN gateway. There are several servers that have virtual machines that use OVS with VXLAN tunnels to communicate. There are also several servers that do not have VXLAN tunnels that use the OS6900-Q32 as a gateway to communicate with the VMs within the VXLAN domain.

CONFIGURING THE BASICS OF THE OMNISWITCH VXLAN GATEWAY

The Alcatel-Lucent Operating System (AOS) VXLAN gateway is a combination of a router and a VXLAN tunnel endpoint (VTEP). Unlike the VXLAN tunnel endpoint (VTEP) running on a server, which acts like a normal IP host, a VXLAN gateway must interact with the IP network.

The routing configuration of the VXLAN gateway is straightforward and uses standard routing protocols and commands that are well understood. The user configures the gateway with the required IP interfaces and protocols that are in use in the routed network which carries the VXLAN traffic. This configuration can include Open Shortest Path First (OSPF), Intermediate System to Intermediate System (IS-IS), or even static routes. The main requirement is that the IP addresses of the remote VTEPs of the servers or other gateways must be reachable.

Here are the minimum routing configuration requirements for the VXLAN gateway to operate:

- A **loopback0** IP interface must be created. This is used by the gateway to identify VXLAN frames destined for it.
  - ¬ **ip interface "Loopback0" address 10.1.1.1**
- If a routing protocol is being used in the network, it must be properly configured according the AOS User Guide for Release 7.3.4.R01.
- If static routes are being used, a route must be entered for every remote VTEP that the gateway will communicate with.
  - ¬ **ip static-route 10.1.1.2/32 gateway 10.1.1.5**
  - ¬ **ip static-route 10.1.1.3/32 gateway 10.1.1.5**

The VXLAN Network Identifier (VNI) related configuration on the OmniSwitch gateway is based on the Service Manager framework which abstracts each VNI as an individual service with a set of access points and distribution points. The Service Access Points (SAPs) define the port and traffic parameters that are used to place traffic into the VNI. The Service Distribution Points (SDPs) define the remote devices (in this case VTEPs) that are also part of the same VNI.

Here are the basic steps and commands for configuring the VXLAN-specific part of the VXLAN gateway:

- Create a service for each VNI.
  - ¬ **service 1 vxlan vnid 1000 stats enable description "VxLAN Service for VNID 1000"**
- Identify the ports on which the devices or networks that need to talk to the VXLAN VNI are located, and make them access ports. Add a SAP to place the correct traffic into the VNI. In this example, the two servers on the left will be sending untagged frames and thus require a 'null'-encapsulated SAP.
  - ¬ **service access port 1/1/3**
  - ¬ **service access port 1/1/4**
  - ¬ **service 1 sap port 1/1/3:0 stats enable**

- ¬ **service 2 sap port 1/1/4:0 stats enable**

- Create an SDP for each of the VTEPs that are to talk to this gateway.
  - ¬ **service sdp 10 vxlan far-end 10.1.1.2 description "To Server A"**
  - ¬ **service sdp 20 vxlan far-end 10.1.1.3 description "To Server B"**
- Bind the VTEPs through their SDP designation to each service that represents the VNIs in which they participate.
  - ¬ **service 1 bind-sdp 10**
  - ¬ **service 1 bind-sdp 20**

CONFIGURING THE OVS ON THE SERVERS

Configuring VXLAN on a server running OVS is a straightforward matter. Each tenant will have its own virtual bridge to which all of the VMs of the other tenants on that server will be attached. Then, a VXLAN tunnel must be created to connect the local virtual bridge to the remote VTEPs that are all participating in the service. There are many possible options for configuring OVS, this document only covers one example.

- Create a bridge for each tenant network, on each server.
  - ¬ **ovs-vsctl add-br my_subnet**
- Create an interface on each bridge that represents a VXLAN tunnel to each remote VTEP.
  - ¬ **ovs-vsctl add-port my_subnet vx101 -- set interface vx101 type=vxlan options:remote_ip=100.1.1.1.1 options:key=1000 options:dst_port=4789**
  - ¬ **ovs-vsctl add-port my_subnet vx102 -- set interface vx102 type=vxlan options:remote_ip=100.1.1.1.2 options:key=1000 options:dst_port=4789**

CONCLUSION

The combination of OVS and the OmniSwitch VXLAN gateway simplifies the transformation of data centers greatly, using legacy and new server VM technologies. In addition, the OmniSwitch 6900 provides unparalleled high performance compared to server- and software-based gateways. Finally, the OmniSwitch VXLAN gateway can be managed as an entity of the fully managed switching solution by the OmniVista® 2500 Network Management System.

Alcatel·Lucent Enterprise