Hewlett Packard Enterprise

HPE VXLAN configuration between 12900 and 5400R

Technical configuration guide
Ver. 1 March 2015

Technical white paper

Contents

Introduction 2
Logical Network Diagram 2
Working mechanisms 3
Hardware and Software Requirements 3
Configuration 4
  HPE FlexFabric 12900 Configuration 4
  HPE FlexFabric 5400R Configuration 5
Verification 6
Appendix A Switches configuration 9
  12900 switch configuration 9
  5400R VSF-stack configuration 12
Additional links 14

Introduction

Virtual eXtensible LAN (VXLAN) is a MAC-in-UDP technology that provides Layer 2 connectivity between distant network sites across an IP network. VXLAN is typically used in data centers for multitenant services.

VXLAN provides the following benefits:

- Support for more virtual switched domains than VLANs. Each VXLAN is uniquely identified by a 24-bit VXLAN ID. The total number of VXLANs can reach 16777216 (2^{24}). This specification makes VXLAN a better choice than 802.1Q VLAN to isolate traffic for VMs.
- Easy deployment and maintenance. VXLAN requires deployment only on the edge devices of the transport network. Devices in the transport network perform typical Layer 3 forwarding.

The VXLAN tunnel endpoints (VTEP) performs Layer 2 or Layer 3 forwarding for VXLANs depending on your configuration:

- In Layer 3 forwarding mode, the VTEP uses the ARP table to forward traffic for VXLANs. Use Layer 3 forwarding mode if you want to use the VTEP as a VXLAN IP gateway.
- In Layer 2 forwarding mode, the VTEP uses the MAC address table to forward traffic for VXLANs.

This Technical Configuration Guide (TCG) describes the Layer 2 forwarding processes. It describes how to configure a VXLAN tunnel between an HPE FlexFabric 12900 and an HPE 5406R that support VXLAN feature. Please note that Only FC, FE, and FX cards support VXLANs.

The intended audience for this TCG is HP Networking Solution Architects, HP Networking Technical Consultants, and HP Networking technical pre-sales staff who are familiar with Comware and ProVision switches.

Note: This example of configuration should be limited to very small deployments, considering the large amount of configuration required to get these static tunnels

---

Logical Network Diagram

As shown in the figure below, the end devices are HPE 5700 and HPE 2920 switches. The VTEPs are the VXLAN tunnel endpoints. For this configuration guide we use the HPE 12900 and the HPE 5400R switches as VTEPs.

An HPE VTEP uses VSIs and VXLAN tunnels to provide VXLAN services.

- VSI. A virtual switching instance is a virtual Layer 2 switched domain. Each VSI provides switching services only for one VXLAN. VSIs learn MAC addresses and forward frames independently of one another. VMs in different sites have Layer 2 connectivity if they are in the same VXLAN.
- VXLAN tunnel. Logical point-to-point tunnels between VTEPs over the transport network. Each VXLAN tunnel can trunk multiple VXLANs. VTEPs encapsulate VXLAN traffic in the VXLAN, outer UDP, and outer IP headers. The devices in the transport network forward VXLAN traffic only based on the outer IP header.

VTEPs encapsulate VXLAN traffic in the VXLAN, outer UDP, and outer IP headers. The devices in the transport network forward VXLAN traffic only based on the outer IP header

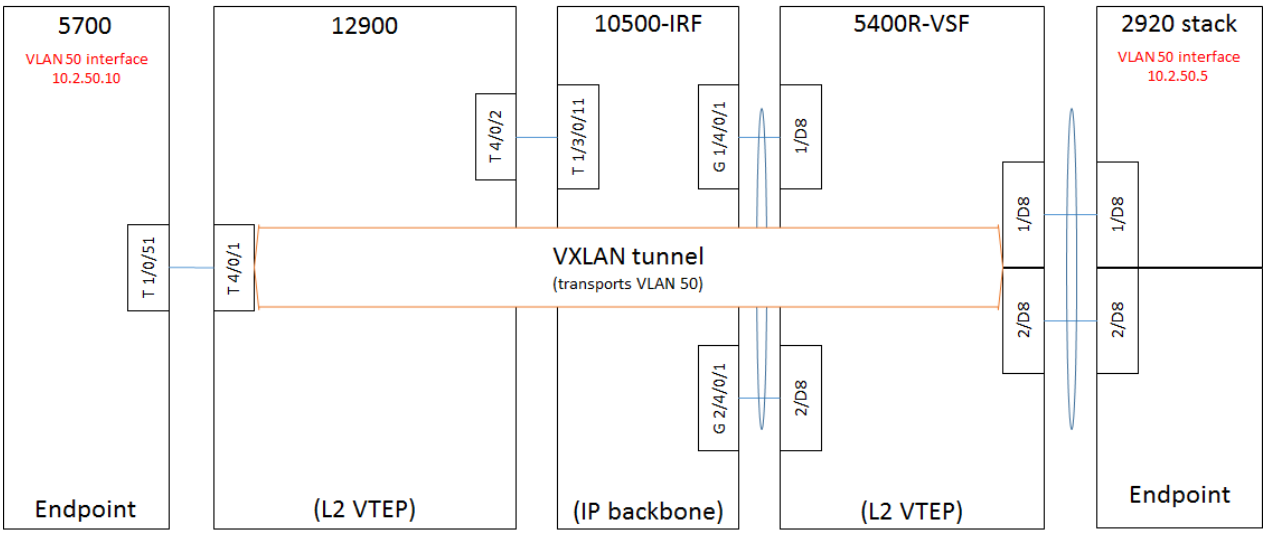

Figure 1 Logical Network Diagram

Working mechanisms

The VTEP uses the following process to forward an inter-site frame:

- Assigns the frame to its matching VXLAN if the frame is sent between sites.
- Performs MAC learning on the VXLAN's VSI.
- Forwards the frame.

VXLAN tunnel establishment and assignment

To provide Layer 2 connectivity for a VXLAN between two sites, you must create a VXLAN tunnel between the sites and assign the tunnel to the VXLAN.

Traffic from the local site to a remote site

The VTEP uses an Ethernet service instance to match a list of VLANs on a site-facing interface. The VTEP assigns customer traffic from the VLANs to a VXLAN by mapping the Ethernet service instance to a VSI.

Traffic from a remote site to the local site

When a frame arrives at a VXLAN tunnel, the VTEP uses the VXLAN ID in the frame to identify its VXLAN.

MAC learning

The VTEP performs source MAC learning on the VSI as a Layer 2 switch.

- For traffic from the local site to the remote site, the VTEP learns the source MAC address before VXLAN encapsulation.
- For traffic from the remote site to the local site, the VTEP learns the source MAC address after removing the VXLAN header.

Traffic forwarding

The VTEP performs Layer 2 or Layer 3 forwarding for VXLANs depending on your configuration:

- In Layer 3 forwarding mode, the VTEP uses the ARP table to forward traffic for VXLANs. Use Layer 3 forwarding mode if you want to use the VTEP as a VXLAN IP gateway.
- In Layer 2 forwarding mode, the VTEP uses the MAC address table to forward traffic for VXLANs.

Hardware and Software Requirements

For this configuration guide we used the following hardware and software versions in order to implement this specific VXLAN solution:

- HPE 5400R zl2 series switches and V3 modules (J9995A): KB.16.01.0004
- HPE 12900 and FX module (LSXM1TGS48FX1): version 7.1.045, Release 1138P01

Configuration

The following steps are to configure VXLAN on the HPE FlexFabric 12900 and on the HPE 5406R. The prerequisites are to configure IP addresses and unicast routing settings on all transport switches:

- Assign IP addresses to the interfaces
- Configure OSPF on all transport network switches
- Configure OSPF to advertise routes

HPE FlexFabric 12900 Configuration

Create a VLAN for the end-point devices

```
vlan 50
interface Vlan-interface50
 ip address 10.2.50.10 255.255.255.0
```

Reserve VXLAN globally

```
reserve-vlan-interface 200 global
```

Enable L2VPN

```
system-view
l2vpn enable
```

Enable Layer 2 forwarding for VXLANs.

```
undo vxlan ip-forwarding
```

Create the VSI vni200 and VXLAN 200.

```
vsi vni200
 vxlan 200
 quit
 quit
```

Assign an IP address to Loopback 0. The IP address will be used as the source IP address of the VXLAN tunnels to Switch B and Switch C.

```
interface loopback0
 ip address 1.220.0.2 255.255.255.255
 quit
```

Create a VXLAN tunnel to the 5400R. The tunnel interface name is Tunnel 1.3

```
interface tunnel 13 mode vxlan
 source 10.220.0.2
 destination 10.2.10.2
 quit
```

Assign Tunnel 13 to VXLAN 200.

```
vsi vni200
 vxlan 200
 tunnel 13
 quit
 quit
```

On the uplink to the 5400R create Ethernet service instance 2.

```
interface Ten-GigabitEthernet4/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan all
 service-instance 2
  encapsulation s-vid 50
```

Map Ethernet service instance 2 to the VSI vni200.

```
interface Ten-GigabitEthernet4/0/1
service-instance 2
xconnect vsi vni200
quit
quit
```

Configure the uplink to the end-point device

```
interface Ten-GigabitEthernet4/0/2
 port link-mode route
 ip address 10.220.0.2 255.255.255.252
```

HPE FlexFabric 5400R Configuration

Enable VXLAN

```
vxlan enable
```

Create Virtual Network and associate Virtual Network ID to VLAN

```
virtual-network 200 50 "vni200"
```

Create VxLAN tunnel between the two switches

```
interface tunnel 13
   tunnel name "VXLAN_Tunnel02"
   tunnel mode vxlan
   tunnel source 10.2.10.2
   tunnel destination 10.220.0.2
   exit
```

Map Overlay-VLAN (VLAN 50) to Tunnel 13

```
vxlan tunnel 13 overlay-vlan 50
```

Verification

As mentioned in the introduction, the VXLAN tunnel endpoints are HPE 5700 and HPE 2920 switches. Below are the details of the endpoints:

Table 1. VXLAN tunnel endpoints

MODEL	IP ADDRESS	MAC ADDRESS
HPE 2920	10.2.50.3	7446-a0ff-48a3
HPE 5700	10.2.50.100	784859-ed0ad0

- The commands to determine the above information are:

```
HP-Stack-2920# show arp

 IP ARP table

  IP Address       MAC Address       Type    Port
  ---------------  ----------------- ------- ----
  10.2.20.5        288023-977b3f     dynamic Trk1
  10.2.50.100      784859-ed0ad0     dynamic Trk1

[5700-GigabitEthernet1/0/48]display arp
  Type: S-Static   D-Dynamic   O-Openflow   R-Rule   M-Multiport  I-Invalid
IP address      MAC address    VLAN     Interface                Aging Type
…
…
10.2.50.3       7446-a0ff-48a3 50       XGE1/0/51                6     D
```

- To confirm that the above data are correct, issue the following commands on both endpoitns to find out the MAC Address for each interface:

```
[5700]display int vlan 50
Vlan-interface50
Current state: UP
Line protocol state: UP
Description: Vlan-interface50 Interface
Bandwidth: 10000000 kbps
Maximum transmission unit: 1500
Internet address: 10.2.50.100/24 (primary)
IP packet frame type: Ethernet II, hardware address: 7848-59ed-0ad0
IPv6 packet frame type: Ethernet II, hardware address: 7848-59ed-0ad0
Last clearing of counters: Never
Last 300 seconds input rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Last 300 seconds output rate: 0 bytes/sec, 0 bits/sec, 0 packets/sec
Input: 0 packets, 0 bytes, 0 drops
Output: 0 packets, 0 bytes, 0 drops

HP-Stack-2920# show system info

 Status and Counters - General System Information

  System Name        : HP-Stack-2920
  System Contact     :
  System Location    : Campus-access
  MAC Age Time (sec) : 300
  Time Zone          : 0
  Daylight Time Rule : None

  Software revision  : WB.16.01.0004
  Base MAC Addr      : 7446a0-ff48a3
```

- To verify the correct neighbor information, issue the LLDP commands:

```
HP-VSF-Switch# show lldp info remo

 LLDP Remote Devices Information

  LocalPort | ChassisId                 PortId PortDescr SysName
  --------- + ------------------------- ------ --------- --------------------
  1/D4      | 70 10 6f 8c 3f 45         2      1/2       Aruba-Stack-3810M
  1/D5      | 74 46 a0 ff 48 a3         54     2/2       HP-Stack-2920
  1/D8      | 78 48 59 e1 32 00         Gig... Gigabi... 10500-IRF-Core
  2/D4      | 70 10 6f 8c 3f 45         98     2/2       Aruba-Stack-3810M
  2/D5      | 74 46 a0 ff 48 a3         2      1/2       HP-Stack-2920
  2/D8      | 78 48 59 e1 32 00         Gig... Gigabi... 10500-IRF-Core

[12900]display lldp neigh list
Chassis ID : * -- -- Nearest nontpmr bridge neighbor
             # -- -- Nearest customer bridge neighbor
             Default -- -- Nearest bridge neighbor
System Name          Local Interface Chassis ID      Port ID
5700                 XGE4/0/1        7848-59ed-0ab5  Ten-GigabitEthernet1/0/51
10500-IRF-Core       XGE4/0/2        7848-59e1-3200  Ten-GigabitEthernet1/3/0/11
OOB                  M-GE0/0/0       b8af-67da-1135  GigabitEthernet2/0/61
```

- Now we can start passing traffic between the two endpoints:

```
[5700-GigabitEthernet1/0/48]ping 10.2.50.3
Ping 10.2.50.3 (10.2.50.3): 56 data bytes, press CTRL_C to break
56 bytes from 10.2.50.3: icmp_seq=0 ttl=255 time=1.825 ms
56 bytes from 10.2.50.3: icmp_seq=1 ttl=255 time=1.460 ms
56 bytes from 10.2.50.3: icmp_seq=2 ttl=255 time=1.653 ms
56 bytes from 10.2.50.3: icmp_seq=3 ttl=255 time=1.486 ms
56 bytes from 10.2.50.3: icmp_seq=4 ttl=255 time=1.465 ms


HP-Stack-2920#  ping 10.2.50.100
10.2.50.100 is alive, time = 3 ms
```

- Besides pinging the two end-point devices, in order to verify the correct functionality, the status of the tunnels and MAC learning, issue the following commands:

On the 12900

- Verify that the VXLAN tunnel interfaces on the VTEP are up

```
[12900]display interface tunnel 13
Tunnel13
Current state: UP
Line protocol state: UP
Description: Tunnel13 Interface
Bandwidth: 64 kbps
Maximum transmission unit: 64000
Internet protocol processing: Disabled
Last clearing of counters: Never
Tunnel source 10.220.0.2, destination 10.2.10.2
Tunnel protocol/transport UDP_VXLAN/IP
```

- Verify that the VXLAN tunnels have been assigned

```
[12900]display l2vpn vsi verbose
VSI Name: vni200
  VSI Index               : 0
  VSI State               : Up
```

```
  MTU                     : 1500
  Bandwidth               : -
  Broadcast Restrain      : -
  Multicast Restrain      : -
  Unknown Unicast Restrain: -
  MAC Learning            : Enabled
  MAC Table Limit         : -
  Drop Unknown            : -
  Flooding                : Enabled
  Statistics              : Disabled
  VXLAN ID                : 200
  Tunnels:
    Tunnel Name          Link ID    State    Type        Flooding proxy
    Tunnel13             0x500000d  UP       Manual      Disabled
  ACs:
    AC                              Link ID    State
    XGE4/0/1 srv2                   0          Up
```

- Verify that the VTEP has learned the MAC addresses of remote

```
[12900]display l2vpn mac-address
MAC Address      State    VSI Name                        Link ID/Name  Aging
7446-a0ff-48a3   Dynamic  vni200                          Tunnel13      Aging
7848-59ed-0ad0   Dynamic  vni200                          0             Aging
```

On the 5400R

- Verify the interface tunnel

```
HP-VSF-Switch(config)# show interface tunnel

 Tunnel Configuration :

  Tunnel               : 251659491
  Tunnel Name          : VXLAN_Tunnel02
  Tunnel Status        : Enabled
  Source Address       : 10.2.10.2
  Destination Address  : 10.220.0.2
  Mode                 : VXLAN Tunnel
  TOS                  : -1
  TTL                  : 64
  IPv6                 : n/a
  MTU                  : 1460
```

- Verify that the VXLAN tunnels have been assigned

```
HP-VSF-Switch(config)# show virtual-network
  Max. Supported Virtual Networks    : 64
  Virtual Networks Configured        : 1

  VN-ID VN-Name                        VLAN-ID VLAN-Name
  ----- ------------------------------ ------- --------------------------------
  200   vni200                         50      VLAN50
```

Appendix A Switches configuration

12900 switch configuration

```
#
 version 7.1.045, Release 1138P01
#
mdc Admin id 1
#
mdc Production-MDC id 2
 mdc start
#
 sysname 12900
#
 telnet server enable
#
 undo vxlan ip-forwarding
#
ospf 1
 area 0.0.0.0
  network 1.220.0.2 0.0.0.0
  network 10.220.0.0 0.0.0.3
#
 lldp global enable
#
 mvrp global enable
#
 reserve-vlan-interface 3000 to 3100
 reserve-vlan-interface 200 global
#
 system-working-mode standard
 password-recovery enable
 lpu-type f-series
#
vlan 1
#
vlan 50
#
vlan 129
#
 stp global enable
#
 l2vpn enable
#
vsi vni200
 vxlan 200
  tunnel 13
#
interface NULL0
#
#
interface LoopBack0
 ip address 1.220.0.2 255.255.255.255
#
interface Vlan-interface1
 mtu 9008
#
interface Vlan-interface50
 ip address 10.2.50.10 255.255.255.0
#
interface FortyGigE5/0/1
 port link-mode route
#
```

```
interface FortyGigE5/0/13
#
# ... other interfaces ...
#
interface M-GigabitEthernet0/0/0
 ip address 10.10.10.44 255.255.255.0
#
interface Ten-GigabitEthernet4/0/2
 port link-mode route
 ip address 10.220.0.2 255.255.255.252
#
interface Ten-GigabitEthernet4/0/1
 port link-mode bridge
 port link-type trunk
 port trunk permit vlan all
 service-instance 2
  encapsulation s-vid 50
  xconnect vsi vni200
#
interface Ten-GigabitEthernet4/0/3
 port link-mode bridge
#
interface Ten-GigabitEthernet4/0/4
 port link-mode bridge
#
interface Ten-GigabitEthernet4/0/5
 port link-mode bridge
#
interface Ten-GigabitEthernet4/0/6
 port link-mode bridge
#
interface Ten-GigabitEthernet4/0/7
 port link-mode bridge
#
interface Ten-GigabitEthernet4/0/8
 port link-mode bridge
#
# ... other interfaces ...
#
#
interface Tunnel13 mode vxlan
 source 10.220.0.2
 destination 10.2.10.2
#
interface Blade-Aggregation1
#
 scheduler logfile size 16
#
line class aux
 user-role network-admin
#
line class vty
 user-role network-operator
#
line aux 0 1
 user-role network-admin
#
#
line vty 0 15
 authentication-mode scheme
 user-role network-operator
#
line vty 16 63
```

```
 user-role network-operator
#
 ip route-static 0.0.0.0 0 10.10.10.254
#
 snmp-agent
 snmp-agent local-engineid 800063A28080F62E82C30700000001
 snmp-agent community write private
 snmp-agent community read public
 snmp-agent sys-info version all
 snmp-agent target-host trap address udp-domain 10.10.10.10 params securityname public v2c
 snmp-agent target-host trap address udp-domain 10.3.10.220 params securityname public v2c
#
acl number 2000
#
acl number 3000
#
acl number 4000
domain system
#
 domain default enable system
#
role name level-0
 description Predefined level-0 role
#
role name level-1
 description Predefined level-1 role
#
role name level-2
 description Predefined level-2 role
#
role name level-3
 description Predefined level-3 role
#
role name level-4
 description Predefined level-4 role
#
role name level-5
 description Predefined level-5 role
#
role name level-6
 description Predefined level-6 role
#
role name level-7
 description Predefined level-7 role
#
role name level-8
 description Predefined level-8 role
#
role name level-9
 description Predefined level-9 role
#
role name level-10
 description Predefined level-10 role
#
role name level-11
 description Predefined level-11 role
#
role name level-12
 description Predefined level-12 role
#
role name level-13
 description Predefined level-13 role
#
```

```
role name level-14
 description Predefined level-14 role
#
user-group system
#
local-user admin class manage
 password hash
$h$6$ucKzWby6Pa3zRhCP$uCDJbw5pvcGP9gIFXP0I4++QDxc9sXvPK8WrwhpwbgK976oHF5r06yLmvdzUcJwOxz6PxwgKu/MRa
pealGtgsA==
 service-type telnet
 authorization-attribute user-role network-admin
 authorization-attribute user-role network-operator
#
return
```

5400R VSF-stack configuration

```
; J9850A Configuration Editor; Created on release #KB.16.01.0004
; Ver #0c:01.7c.59.f4.7b.ff.ff.fc.ff.ff.3f.ef:
hostname "HP-VSF-Switch"
module 1/D type j9995a
module 1/F type j9991a
module 2/D type j9995a
vsf
   enable domain 1
   member 1
      type "J9850A" mac-address 288023-976b00
      priority 128
      link 1 1/D1
      link 1 name "I-Link1_1"
      exit
   member 2
      type "J9850A" mac-address 3ca82a-3b6e00
      priority 128
      link 1 2/D1
      link 1 name "I-Link2_1"
      exit
   exit
no rest-interface
vxlan enable
vxlan tunnel 13 overlay-vlan 50
trunk 1/D8,2/D8 trk1 lacp
trunk 1/D5,2/D5 trk2 lacp
max-vlans 4000
ip routing
interface loopback 0
   ip address 1.2.10.2
   ip ospf 1.2.10.2 area backbone
   exit
interface tunnel 13
   tunnel name "VXLAN_Tunnel02"
   tunnel mode vxlan
   tunnel source 10.2.10.2
   tunnel destination 10.220.0.2
   exit
snmp-server community "public" operator unrestricted
snmp-server community "private"
oobm
   ip address dhcp-bootp
   vsf member 1
      ip address dhcp-bootp
      exit
```

```
   vsf member 2
      ip address dhcp-bootp
      exit
   exit
router ospf
   area backbone
   enable
   exit
vlan 1
   name "DEFAULT_VLAN"
   no untagged 1/D4,2/D4,Trk1-Trk2
   untagged 1/D2-1/D3,1/D6-1/D7,1/F1-1/F24,2/D2-2/D3,2/D6-2/D7
   ip address dhcp-bootp
   exit
vlan 10
   name "VLAN10"
   untagged Trk1
   ip address 10.2.10.2 255.255.255.0
   ip ospf 10.2.10.2 area backbone
   exit
vlan 20
   name "VLAN20"
   untagged 1/D4,2/D4,Trk2
   ip address 10.2.20.5 255.255.255.0
   ip ospf 10.2.20.5 area backbone
   exit
vlan 30
   name "VLAN30"
   tagged 2/D4,Trk2
   ip address 10.2.30.5 255.255.255.0
   ip ospf 10.2.30.5 area backbone
   exit
vlan 40
   name "VLAN40"
   ip address 10.2.40.5 255.255.255.0
   ip ospf 10.2.40.5 area backbone
   exit
vlan 50
   name "VLAN50"
   tagged Trk2
   ip address 10.2.50.5 255.255.255.0
   exit
vlan 100
   name "VLAN100"
   ip address 10.100.0.5 255.255.255.0
   exit
spanning-tree Trk1 priority 4
spanning-tree Trk2 priority 4
no allow-v2-modules
virtual-network 200 50 "vni200"
password manager
```

Additional links

R213x-HPE FlexFabric 12900 VXLAN Configuration Guide

Learn more at

https://www.hpe.com/us/en/networking.html

Share with colleagues

Rate this document

Version 1.1, March 2015

Hewlett Packard Enterprise