

OVSDB VTEP with Midonet Configuration

OVSDB VTEP with Midonet Configuration

Requirements

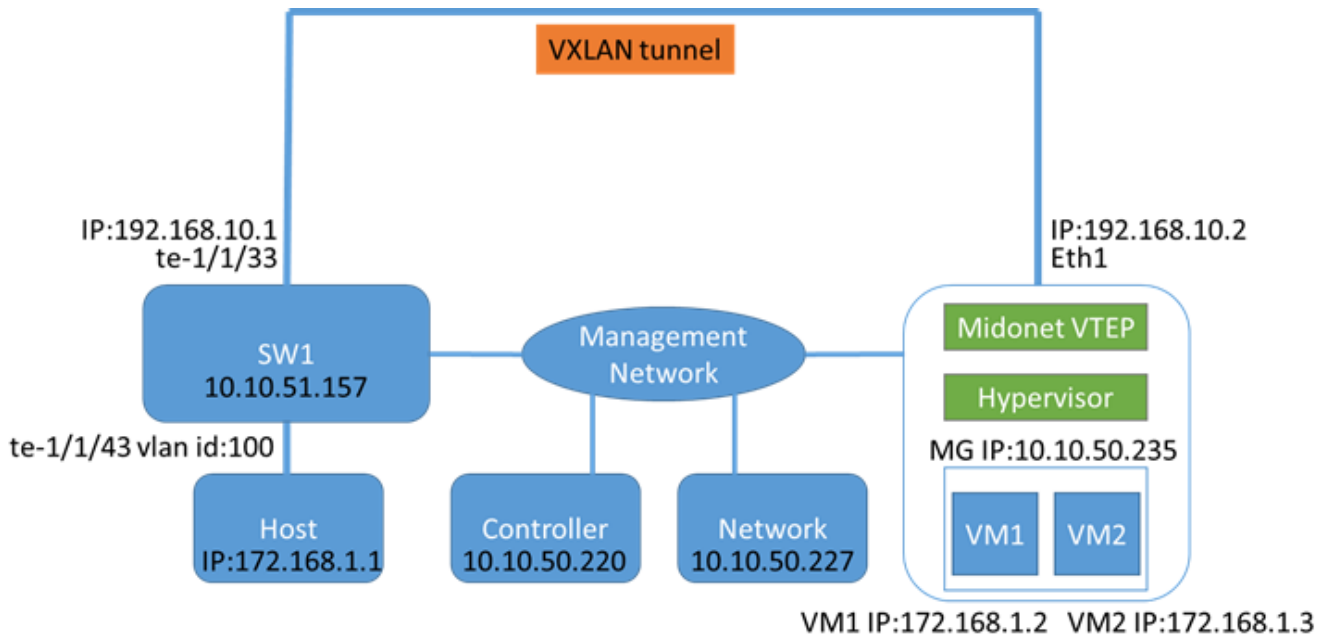
This example uses the following hardware and software components:

- Trident-II Series switch
- PicOS OS Release 2.6 and higher
- Midonet controller

Overview

In this example, OVSDB VTEP is configured to run on a VXLAN domain. VTEP interface sources are configured to the loopback address. Interfaces are configured for VLAN tagging and encapsulation. Static route protocols are configured to facilitate unicast routing.

Topology



The VXLAN networks Host and VM IP address and Mac address information:

Host IP Address: 172.168.1.1, Mac address: 00:07:43:05:45:11

VM1 IP Address: 172.168.1.2, Mac address: fa:16:3e:00:0c:f3

VM2 IP Address: 172.168.1.3, Mac address: fa:16:3e:28:aa:cd

Configuring OVSDB on Trident-II Series Switches

CLI Quick Configuration

To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and copy and paste the commands into the CLI.

SW1 Configure:

```
set interface gigabit-ethernet te-1/1/33 speed "1000"
set interface gigabit-ethernet te-1/1/33 family ethernet-switching native-vlan-id 1000
set interface gigabit-ethernet te-1/1/43 family ethernet-switching port-mode "trunk"
set interface gigabit-ethernet te-1/1/43 family ethernet-switching vlan members 100
set protocols ovssdb management-ip 10.10.51.157
set protocols ovssdb controller c1 protocol "ptcp"
set protocols ovssdb interface te-1/1/43
set l3-interface vlan-interface 1000 address 192.168.10.1 prefix-length 24
set vlans vlan-id 100
set vlans vlan-id 1000 l3-interface "1000"
set vxlans source-interface 1000 address 192.168.10.1
set vxlans ovssdb-managed true
```

Configuring OVSSDB Step-by-Step Procedure

The following example shows how to set up a basic VTEP OVSSDB configuration with VXLAN domain. To configure VXLAN on Trident-II Series switches, follow these steps:

Configuring OVSSDB Step-by-Step for SW1

1. Configure the VLAN ID to 10 for vxlan domain.

```
set vlans vlan-id 100
```

2. Configure the VLAN ID to 1000 for IP routing.

```
set vlans vlan-id 1000 l3-interface vlan-1000
```

3. Configure the te-1/1/1 interface VLAN ID to 1000.

```
set interface gigabit-ethernet te-1/1/33 family ethernet-switching native-vlan-id 1000
```

4. Configure VLAN trunk for te-1/1/43.

```
set interface gigabit-ethernet te-1/1/43 family ethernet-switching port-mode "trunk"
```

5. Configure the te-1/1/43 interface VLAN ID to 100.

```
set interface gigabit-ethernet te-1/1/43 family ethernet-switching vlan members 100
```

6. Configure IP address for the loopback interface.

```
set l3-interface loopback lo address 10.10.10.1 prefix-length 32
```

7. Configure IP address for the vlan-interface vlan-1000.

```
set l3-interface vlan-interface vlan-1000 address 192.168.10.1 prefix-length 24
```

8. Configure VTEP interface sources IP address.

```
set vxlans source-interface vlan-1000 address 192.168.10.1
```

9. Enable VXLAN managed by ovssdb

```
set vxlan ovssdb-managed true
```

10. Configure ovssdb management interface IP address

```
set protocols ovssdb management-ip 10.10.51.157
```

11. Configure ovssdb controller protocol

```
set protocols ovssdb controller c1 protocol ptcp
```

12. Configure ovssdb controller port

```
set protocols ovssdb controller ovssdb port 6632
```

13. Configure ovssdb the interface on the VTEP

```
set protocols ovssdb interface te-1/1/43
```

Manually Configuring the Switch as a VTEP Gateway by Midonet cli

Configuring the ovssdb by midonet cli for SW1

1. Add a virtualization image on OpenStack Dashboard

Images

Images Project (1) Shared with Me (0) Public (0) + Create Image ✕ Delete Images

Image Name	Type	Status	Public	Protected	Format	Size	Actions
test	Image	Active	No	No	QCOW2	12.6 MB	Launch

Displaying 1 item

2. Create a network on OpenStack Dashboard

Networks

Networks + Create Network ✕ Delete Networks

Project	Network Name	Subnets Associated	DHCP Agents	Shared	Status	Admin State	Actions
admin	vxlan	m1 172.168.1.0/24	1	No	ACTIVE	UP	Edit Network

Displaying 1 item

3. Create two virtual hosts, and add them to network on OpenStack Dashboard

Instances

Instances Instance Name Filter Filter Launch instance Soft Reboot Instances Terminate Instances

Instance Name	Image Name	IP Address	Size	Key Pair	Status	Availability Zone	Task	Power State	Time since created	Actions
vm2	test	172.168.1.3	m1.tiny	-	Active	nova	None	Running	3 hours, 19 minutes	Create Snapshot
vm1	test	172.168.1.2	m1.tiny	-	Active	nova	None	Running	3 hours, 21 minutes	Create Snapshot

Displaying 2 items

4. Create a tunnel zone of type 'vtep'

```
midonet> tunnel-zone create name vtep_zone1 type vtep
tzone0
```

5. Add a VTEP to MidoNet, and assign it to the 'vtep' tunnel zone that you created.

```
midonet> vtep add management-ip 10.10.51.157 management-port 6632 tunnel-zone tzone0
name br0 description OVS VTEP Emulator management-ip 10.10.51.157 management-port 6632 tunnel-zone
tzone0 connection-state CONNECTED
midonet> list vtep
name br0 description management-ip 10.10.51.157 management-port 6632 tunnel-zone tzone0 connection-
state CONNECTED
```

6. Create a binding between the VTEP and a Neutron network behind a MidoNet bridge

```
midonet> bridge list
bridge bridge1 name vxlan state up
midonet> show bridge bridge1 id
85296f07-2235-4963-8160-fb66eca85675
midonet>
```

7. Add the host's IP address to the same tunnel zone as the VTEP

```
midonet> tunnel-zone tzone0 add member host host0 address 192.168.10.2
zone tzone0 host host0 address 192.168.10.2
midonet>
```

8. Create a binding between the VTEP's vlan 100 interface te-1/1/43 and the Neutron network behind the bridge1

```
midonet> vtep management-ip 10.10.51.157 binding add network-id 85296f07-2235-4963-8160-fb66eca85675
physical-port te-1/1/43 vlan 100
Internal error: The server could not comply with the request since it is either malformed or otherwise
incorrect.
midonet> vtep management-ip 10.10.51.157 binding list
management-ip 10.10.51.157 physical-port te-1/1/43 vlan 100 network-id 85296f07-2235-4963-8160-
fb66eca85675
```

9. Add the IP address of the host on the VTEP to the security group ip-address-group0

```
midonet> ip-address-group ip-address-group0 add ip address 172.168.1.1
address 172.168.1.1
midonet>
```

View the VXLAN table of SW1:

```

admin@XorPlus# run show vxlan
Egress map:
    egress_id 100009 MAC 0:c:29:23:31:9, port_id 1/1/33, vif_index 8 unicast
L3 tunnel mac map:
    vlan id 1000, ref_count 1
Port vlan map mode map & Termination admin state map:
    port id 1/1/43, ref_count 1
Tunnel Map:
    tunnel id 0X4C000200, dst_vtep 192.168.10.2, nexthops (192.168.10.2 ), ecmp_id 100009, ref_count 1
    tunnel id 0X4C000001, dst_vtep 224.0.0.1, nexthops (), ecmp_id 0, ref_count 1
Access ports:
    id 0X80000002, vpn_id 0X7000, port_id 1/1/43, vlan_id 100, egress_id 100010
Network ports:
    id 0X80000003, vpn_id 0X7000, port_id 1/1/33, egress_id 100009, tunnel_id 0X4C000200, unicast
    id 0X80000004, vpn_id 0X7000, port_id 1/1/33, egress_id 100011, tunnel_id 0X4C000200, multicast
    id 0X80000001, vpn_id 0XFFFFFFF, port_id 1/1/0, egress_id 100006, tunnel_id 0X4C000001, multicast
BFD sessions:
admin@XorPlus#

```

The show vxlan to displays information about VXLAN endpoint configuration. Display VXLAN to the remote vxlan tunnel vtep has a nexthops(192.168.10.2).

View the VXLAN mac table of SW1:

```

admin@XorPlus# run show vxlan address-table
VNID          MAC address      Type      Interface      VTEP
-----
10001         00:1e:c9:bb:bb:ce Dynamic    te-1/1/43
10001         fa:16:3e:00:0c:f3 Static     192.168.10.2
10001         fa:16:3e:28:aa:cd Static     192.168.10.2
admin@XorPlus#

```

Dump the ovsdb hardware vtep table of SW1:

```

root@XorPlus$ovsdb-client dump hardware_vtep
Arp_Sources_Local table
_uuid locator src_mac
-----

Arp_Sources_Remote table
_uuid locator src_mac
-----

Global table
_uuid          managers          switches
-----
4146166b-ad2e-4d05-857f-8ba4b3f0ac0d [bd6ac790-b304-4ed7-a77b-8ab7063b8132] [cfdcc9fa-0295-44b0-81c3-c975b3d463cb]

Logical_Binding_Stats table
_uuid bytes_from_local bytes_to_local packets_from_local packets_to_local
-----

Logical_Router table
_uuid description name static_routes switch_binding
-----

Logical_Switch table
_uuid          description name          options tunnel_key
-----

```

ele37b4a-37fe-43f2-a9f7-3a9925b6e92e "" "mn-85296f07-2235-4963-8160-fb66eca85675" {} 10001

Manager table

_uuid	inactivity_probe	is_connected	max_backoff	other_config
status		target		
bd6ac790-b304-4ed7-a77b-8ab7063b8132	30000	true	3000	{}
sec_since_connect="13921", state=ACTIVE} "ptcp:6632"				

Mcast_Macs_Local table

MAC	_uuid	ipaddr	locator_set	logical_switch
-----	-------	--------	-------------	----------------

Mcast_Macs_Remote table

MAC	_uuid	ipaddr
locator_set		logical_switch

unknown-dst	8c6b4993-7be8-4d85-811b-3255e15d2f92	""	6d275247-2c1b-4c79-8f08-b17d93bd1e32	ele37b4a-37fe-43f2-a9f7-3a9925b6e92e
-------------	--------------------------------------	----	--------------------------------------	--------------------------------------

Physical_Locator table

_uuid	dst_ip	encapsulation_type
-------	--------	--------------------

d983943f-c791-4431-89a2-ec6a531a4d15	"192.168.10.1"	"vxlan_over_ipv4"
09c0f3c2-d42a-406b-8644-3bffc472a247	"192.168.10.2"	"vxlan_over_ipv4"

Physical_Locator_Set table

_uuid	locators
-------	----------

6d275247-2c1b-4c79-8f08-b17d93bd1e32	[09c0f3c2-d42a-406b-8644-3bffc472a247]
--------------------------------------	--

Physical_Port table

_uuid	description	name	port_fault_status
vlan_bindings	vlan_stats		

35f008a2-e248-4330-a1e5-85f3f843bc68	""	"te-1/1/43"	[]	{100=ele37b4a-37fe-43f2-a9f7-3a9925b6e92e} {}
--------------------------------------	----	-------------	----	---

Physical_Switch table

_uuid	description				
management_ips	name	ports	switch_fault_status	tunnel_ips	tunnels

cfdcc9fa-0295-44b0-81c3-c975b3d463cb	""	["10.10.51.157"]	"br0"	[35f008a2-e248-4330-a1e5-85f3f843bc68, 3d5eae61-46bc-4e3c-84f3-06aed7961ff5, c30d1ef4-b54a-4946-bd8c-460af234875e]	["192.168.10.1"]	[]
--------------------------------------	----	------------------	-------	--	------------------	----

SSL table

_uuid	bootstrap_ca_cert	ca_cert	certificate	external_ids	private_key
-------	-------------------	---------	-------------	--------------	-------------

Tunnel table

_uuid	bfd_config_local	bfd_config_remote	bfd_params	bfd_status	local	remote
-------	------------------	-------------------	------------	------------	-------	--------

Ucast_Macs_Local table

MAC	_uuid	ipaddr
locator		logical_switch

"00:1e:c9:bb:bb:ce"	f431e446-6c1c-4842-ad3a-19cd04a54952	""	d983943f-c791-4431-89a2-ec6a531a4d15	ele37b4a-37fe-43f2-a9f7-3a9925b6e92e
---------------------	--------------------------------------	----	--------------------------------------	--------------------------------------

Ucast_Macs_Remote table

MAC	_uuid	ipaddr
locator		logical_switch

```

-----
"fa:16:3e:00:0c:f3" 0c6732d3-9e72-4444-9d8f-07abde993aa7 " " 09c0f3c2-d42a-406b-8644-3bffc472a247 e1e37b4a-
37fe-43f2-a9f7-3a9925b6e92e
"fa:16:3e:28:aa:cd" b1160e14-8e35-4293-a987-59ddc29f7304 " " 09c0f3c2-d42a-406b-8644-3bffc472a247 e1e37b4a-
37fe-43f2-a9f7-3a9925b6e92e
root@XorPlus$

```

Ping VM1 and VM2 on the Host:

```

root@Dev-45:~# ping 172.168.1.2 -c 5
PING 172.168.1.2 (172.168.1.2) 56(84) bytes of data.
64 bytes from 172.168.1.2: icmp_req=1 ttl=64 time=3.92 ms
64 bytes from 172.168.1.2: icmp_req=2 ttl=64 time=1.51 ms
64 bytes from 172.168.1.2: icmp_req=3 ttl=64 time=1.47 ms
64 bytes from 172.168.1.2: icmp_req=4 ttl=64 time=1.59 ms
64 bytes from 172.168.1.2: icmp_req=5 ttl=64 time=1.57 ms

--- 172.168.1.2 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 1.476/2.015/3.920/0.954 ms
root@Dev-45:~# ping 172.168.1.3 -c 5
PING 172.168.1.3 (172.168.1.3) 56(84) bytes of data.
64 bytes from 172.168.1.3: icmp_req=1 ttl=64 time=10.1 ms
64 bytes from 172.168.1.3: icmp_req=2 ttl=64 time=1.70 ms
64 bytes from 172.168.1.3: icmp_req=3 ttl=64 time=1.64 ms
64 bytes from 172.168.1.3: icmp_req=4 ttl=64 time=1.62 ms
64 bytes from 172.168.1.3: icmp_req=5 ttl=64 time=1.67 ms

--- 172.168.1.3 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4006ms
rtt min/avg/max/mdev = 1.623/3.365/10.189/3.412 ms
root@Dev-45:~# arp -n

```

Address	HWtype	HWaddress	Flags	Mask	Iface
172.168.1.3	ether	fa:16:3e:28:aa:cd	C		eth1.100
172.168.1.2	ether	fa:16:3e:00:0c:f3	C		eth1.100