Introduction to VRF

VRF (Virtual Routing and Forwarding) is a technology that virtualizes a single physical routing device into multiple virtual routing devices, each of them being (relatively) independent of each other, allowing for overlapping subnets, separate routing tables to make Layer 3 segregated, separate ARP tables and separate sets of Layer 3 VLAN interfaces assigned to each VRF.

Figure 1 Multiple VRF Process Modules on One Pica8 Switch

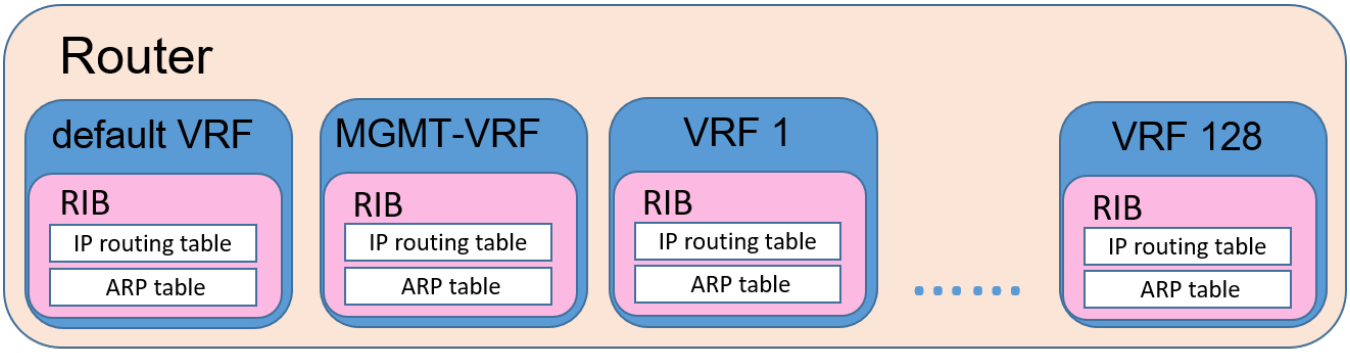

Figure 1 shows multiple VRF process modules on one Pica8 switch. Pica8 switches support multiple VRF instances: one default VRF, one management VRF and multiple user-defined VRFs. IP routing and traffic is separated between the VRFs. By default, PICOS starts up with only default VRF, which cannot be deleted. User can create other VRFs based on the requirements for route separation. The command **set system management-vrf enable** <true | false> can be used to enable management VRF and **set ip vrf** <*vrf-name*> [**description** <*string*>] can be used to create the user-defined VRFs. A maximum of 128 user-defined VRFs can be created on a Pica8 switch.

Currently, default VRF supports static routing and other routing protocols (BGP, OSPF, RIP, PIM), the user-defined VRFs support static routing and OSPF routing protocol. However, management VRF supports neither static routing nor any other dynamic routing protocols.

The following sections describe the VRF characteristics and application scenarios, then details how to use default VRF, management VRF and user-defined VRFs.

VRF Characteristics

- Each VRF has an independent routing table and ARP table to implement independent routing and forwarding functions.
- Each VRF has an independent address space. This allows address overlapping between different VRFs without address conflicts occurring on the same device.
- Users in the same VRF can communicate with each other, but users in different VRFs cannot communicate with each other.
- As with the introduction of management VRF, Out of Band (OOB) management flow is completely separate from data flow, which enhances the security of the management network.
- Default VRF supports static routing and other routing protocols (BGP, OSPF, RIP, PIM), the user-defined VRFs support static routing and OSPF protocol. However, Management VRF does not support any routing protocols, including static routing.

Application Scenarios

This document lists several use cases which can be deployed with VRFs as follows,

- User can deploy VRF function to solve the problem of insufficient IP addresses, as different VRFs have different address spaces which allows address overlapping between different VRFs.
- User can deploy VRF function to achieve traffic isolation of different users and increase data communication security, as the communication between different tenants is segregated in different VRFs.
- User can enable management VRF to separate the Out of Band management traffic from the data traffic, thus to enhance the security of the management network.
- VRF virtualizes a single physical routing device into multiple virtual routing devices; this can save hardware costs.

Default VRF

By default, all the L3 VLAN interfaces and Ethernet0 management interface, and their IP route tables share one VRF - the default VRF. When a VLAN interface is created, it is in the default VRF if not explicitly bound to any VRF. However, user can bind an existing Layer 3 VLAN interface to a user-defined VRF by using command **set l3-interface vlan-interface** <*interface-name*> **vrf** <*vrf-name*>.

Ethernet0 is used for Out of Band (OOB) management, the L3 VLAN interfaces can be used for in-band (IB) management or transmission of data traffic.

Default VRF supports static routing and other routing protocols (BGP, OSPF, RIP, PIM).

In-band Management Interface

By default, the user cannot remotely log in and manage the switch through an L3 VLAN interface. PICOS provides in-band management in default VRF. In-band management provides a method of access to the switch even if the Ethernet0 interface is down. You can enable in-band management function by setting **set system inband vlan-interface** <*vlan-interface*> to perform the SSH, TELNET, SNMP and HTTP services through any one of the VLAN interfaces in the default VRF.

For example,

Set VLAN interface VLAN400 in the default VRF as the in-band management port.

```
admin@Xorplus# set system inband vlan-interface VLAN400
admin@Xorplus# commit
```

In-band management provides a method of access to the switch even if the Ethernet0/1 interface is down.

NOTE:

- Only the L3 VLAN interface in the default VRF can be set as the in-band management port. If a VLAN interface has been set as an in-band management port, it cannot be bound to other VRFs.
- A maximum of four L3 VLAN interfaces in the default VRF can be set as the in-band management ports by using **set system inband vlan-interface**<*vlan-interface*> command.
- By default, all the management services (including **802.1X / OVSDB management protocol / SNMP trap / sFlow / syslog / NTP / TACACS+ /RADIUS**) run in default VRF. If management VRF is enabled, you have to consider which VRF (default VRF or management VRF) will be used to run management services, see the **Management Services** for details.

Management VRF

By default, PICOS starts up with only default VRF, management VRF function is disabled. To enhance the security of the management network, and prevent attacks by illegal users, users can use command **set system management-vrf enable** true to enable management VRF.

Once management VRF is enabled, a VRF with fixed name **mgmt-vrf** is created automatically by the system, and the Eth0 management interface is automatically moved from the default VRF to the management VRF. As long as the management VRF is not disabled, Eth0 is always in the management VRF. Eth0 will return from the management VRF to the default VRF only when management VRF is disabled by setting **set system management-vrf enable** false.

Management VRF is dedicated to transmit the OOB management traffic. Other VRFs are used to transmit the data traffic, thus separating the OOB management traffic from the data traffic effectively.

The management services (including **802.1X / OVSDB management protocol / SNMP trap / sFlow / syslog / NTP / TACACS+/RADIUS**) are running in default VRF by default. However, when management VRF is enabled and Eth0 is used for management services, the management services need to be manually moved from the default VRF into the management VRF by using relevant CLI commands.

For example, if you want to enable management VRF and use the Eth0 management interface for NTP service, the VRF-related configurations are as follows. For more details, see **Management Services.**

```
admin@Xorplus# set system management-vrf enable true
admin@Xorplus# set system ntp vrf vrf mgmt-vrf
admin@Xorplus# commit
```

Management VRF configuration notes:

- Only Eth0 management interfaces can run in the management VRF, no other VLAN interface or loopback interface can run in the management VRF.
- Management VRF supports neither static routing nor any other dynamic routing protocol.

Management Services

To properly use management services, the following descriptions in this section provide a guidance.

By default, all the management services (including **802.1X / OVSDB management protocol / SNMP trap / sFlow / syslog / NTP / TACACS+/RADIUS**) run in default VRF. If management VRF is enabled, Eth0 interface is automatically moved from the default VRF to the management VRF, you have to decide which VRF (default VRF or management VRF) will be used to run the management services.

1. syslog/NTP/TACACS+/RADIUS/SNMP Trap/sFlow/NAC/OVSDB

Management services (referring to syslog/NTP/TACACS+/RADIUS/SNMP Trap/sFlow/NAC /OVSDB) are bound to default VRF at system startup. They can move from default VRF to management VRF and vice versa by using the following CLI commands.

set system syslog vrf <**mgmt-vrf** | **default**>

set system ntp vrf <mgmt-vrf | default>

set system aaa radius vrf <mgmt-vrf | default>

set system tacacs-plus radius vrf <mgmt-vrf | default>

set protocols snmp trap-group vrf <mgmt-vrf | default>

set protocols sflow collector <*ip-address*> **vrf <mgmt-vrf | default>**

set protocols dot1x aaa vrf <mgmt-vrf | default>

set protocols ovsdb controller <*controller-name*> **vrf <mgmt-vrf | default>**

These management services cannot be bound to a user-defined VRF.

Whether the management service is in the default VRF or the management VRF, you need to ensure that the server relevant to the management service (e.g. syslog server, TACACS+ server) is route reachable in the VRF running the management service.

For example,

If you want to enable management VRF and use the Eth0 management interface for NTP service, NTP service needs to be manually moved from the default VRF into the management VRF by using CLI command **set system ntp vrf mgmt-vrf**. The VRF-related configurations are as follows:

```
admin@Xorplus# set system management-vrf enable true
admin@Xorplus# set system ntp vrf mgmt-vrf
admin@Xorplus# commit
```

2. ssh/scp/tftp/ping/traceroute/apt-get

When executing ssh/scp/tftp/ping/traceroute/apt-get commands, the system looks for the next-hop route in default VRF by default. If management VRF is enabled, and Eth0 is used as the route interface, you have to add the VRF parameter in the command to specify that finding the next-hop route in the management VRF.

- At the Linux prompt

If management VRF is enabled, and you want to find the next-hop route in management VRF when running the commands **traceroute/SCP/ping/apt get /SSH** at Linux prompt, that is, using Eth0/1 management interface as the route interface, you have to add **ip vrf exec mgmt-vrf** before the commands.

The example format of these commands is shown below:

sudo ip vrf exec <**mgmt-vrf|***vrf-name*> **traceroute** 10.10.51.11

sudo ip vrf exec <**mgmt-vrf|***vrf-name*> **scp** admin@10.10.51.18:/home/Pica8.pm

sudo ip vrf exec <mgmt-vrf|*vrf-name*> ping 10.10.51.1

sudo ip vrf exec <mgmt-vrf|*vrf-name*> apt-get update

sudo ip vrf exec <mgmt-vrf|*vrf-name*> ssh <*ip-address*>

ip vrf exec <**mgmt-vrf|***vrf-name*> is added to specify which VRF to run the command in. If not specified, find the next hop routing information from the default VRF.

- At PICOS CLI prompt

If management VRF is enabled, and you want to find the next-hop route in management VRF when running the commands ping/traceroute/tftp at PICOS CLI prompt, that is, using Eth0/1 management interface as the route interface, you need to add **vrf mgmt-vrf** before the commands.

ping/traceroute/tftp commands are shown below,

ping <ip-address> [<packets>] [**vrf** <**mgmt-vrf** | *vrf-name*>] [**source** <source-ip-address >] [**deadline** <deadline-time>] [**ttl** <ttl-value>] [**interval** <interval-value>] [**pattern** <pattern-value>] [**size** <size-value>] [**tos** <tos-value>]

traceroute <ipv4-address> [**vrf** <**mgmt-vrf** | *vrf-name*>]

file tftp get remote-file <*remote-file-path*> [**local-file** *local-file-path*] **ip-address** <*ip-address*> [**vrf** <**mgmt-vrf** | *vrf-name*>]

[**vrf** <**mgmt-vrf** | *vrf-name*>] is used to specify which VRF to run the command in. If no VRF is specified, find the next hop routing information from the default VRF. For details about the usage of each parameter, see section Command Reference.

User-defined VRF

By default, all the L3 VLAN interfaces and their IP route tables are in the default VRF. User can create multiple new VRFs with the VRF definition command **set ip vrf** <*vrf-name*> [**description** <*string*>], and then add the Layer 3 VLAN interfaces and their route settings to different VRFs to run services in these VRFs. The system segregates the IP routing table, ARP table, hardware forwarding table, and host hardware forwarding table of different VRFs on one customer edge (CE) device. A maximum of 128 user-defined VRFs can be created.

Figure 2 Networking diagram of VRF

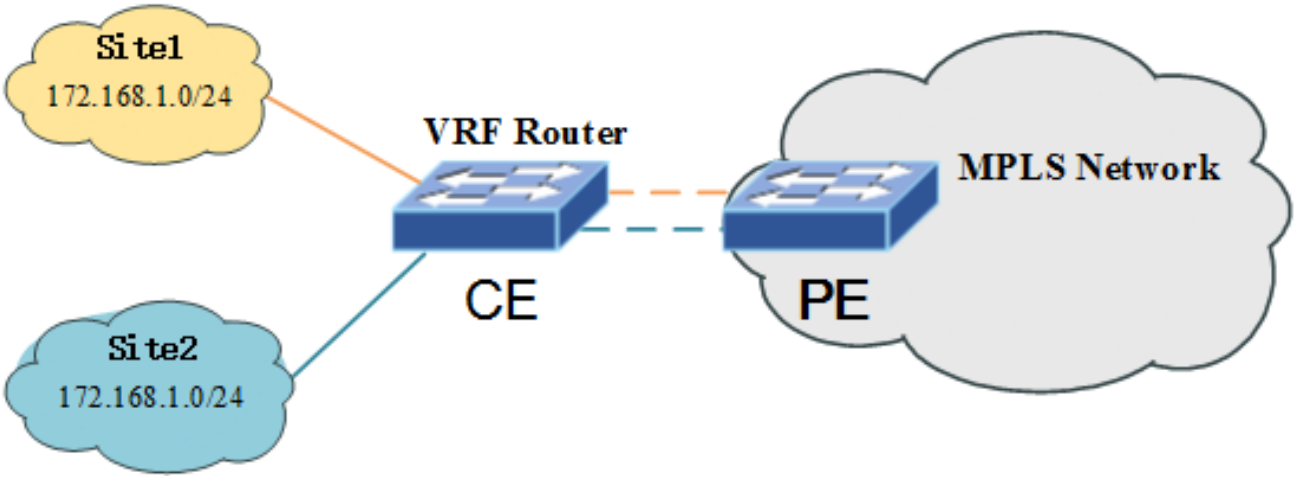

In **Figure 2**, when implementing VRF function on the CE device for different L3 VLAN access interfaces, users from Site1 and Site2 can use overlapping IP addresses when accessing the internet through the CE and have segregated users' routing spaces on the CE device.

When the CE switch receives the data packets, it looks up the IP routing table divided by the VRF, which is bound to the ingress Layer 3 VLAN interface, and then forwards the data packets based on the routing entry in this VRF.

VRF implements data traffic segregation among different customers while sharing the same physical router device, that is to say, users in the same VRF could communicate with each other, but it could not communicate with each other in different VRFs.

NOTE:

- PICOS supports only VRF-Lite, a lighter version of VRF, referring to VRF without MPLS.
- The user-defined VRFs support static routing and OSPF routing protocol.