61202880L1-29.4M
September 2014

ADTRAN

Configuration Guide

Configuring Multi-VRF in AOS

This configuration guide will aid in the setup of multiple virtual routing and forwarding (multi-VRF) instances for ADTRAN Operating System (AOS) products. The guide includes an overview of multi-VRF general concepts, detailed steps with example configurations, and command summary. The troubleshooting section outlines proper use of **clear**, **debug**, and **show** commands to verify that multi-VRF has been configured properly on the AOS product(s).

This guide consists of the following sections:

Multi-VRF Overview

Virtual routing and forwarding (VRF) on an AOS product enables a single physical router to be logically partitioned into multiple virtual router instances. Each router instance has its own route table and associated interfaces. Traffic being forwarded on one router instance is prevented from carrying over into another router instance. Each router instance is independent, which allows for IPv4 address overlap or use of the same IPv4 subnets between the router instances without any conflict. Typically, this feature is used on the customer edge (CE) router in a Multiprotocol Label Switching (MPLS) VPN network where more than one VPN is terminated at the CE router.

Hardware and Software Requirements and Limitations

Introduced in AOS 16.1 for data products and AOS A1 for voice products, support for multi-VRF is available on AOS products as outlined in the *AOS Product Feature Matrix*, located in the ADTRAN support community at https://supportforums.adtran.com.

Multi-VRF can only be configured using the command line interface (CLI). Applying AOS configuration settings using the web graphical user interface (GUI) only affects the default VRF instance.

Interfaces that can be associated with a VRF instance are ATM subinterfaces, demand, EFM group subinterfaces, Ethernet, Ethernet subinterfaces, Frame Relay subinterfaces, HDLC, PPP, multilink PPP, tunnel, and VLAN. An interface may only be associated with a single VRF instance.

As of release 18.1, AOS supports 2-byte autonomous system (AS) numbers. As of release 18.3, AOS supports 4-byte AS numbers as part of the route distinguisher for VRFs. VRF routes can be distinguished by 2-byte AS number, 4-byte AS number, or by IPv4 address.

The following table identifies all the features that currently support multi-VRF and the AOS version these features became VRF-aware. Refer to *Understanding AOS Version Changes* for more information about the ADTRAN Operating System version numbering system.

Table 1. AOS Features Supporting Multi-VRF

AOS Feature Name	AOS Version
AAA - Radius	R10.7.0
AAA - TACACS+	R10.7.0
Access Lists IPv4	16.01.01
Access Lists IPv6	18.01.01
ARP - Voice Products	A1.01.00
ARP	16.01.01
BGP	18.03.01
Crypto IPv6	R10.7.0
DHCP - Voice Products	A1.01.00
DHCPv4	17.01.01
DHCPv6	R10.1.0

Table 1. AOS Features Supporting Multi-VRF *(Continued)*

AOS Feature Name	AOS Version
DNS Lookup/DNS Proxy	R10.3.0
Firewall - Voice Products	A1.01.00
Firewall IPv4	17.01.01
Firewall IPv6	18.01.01
FTP Server	R10.7.0
HTTP/HTTPS Server	R10.7.0
IP Flow	16.01.01
IP Route IPv4	16.01.01
IP Route IPv6	18.01.01
IPSec IPv6	R10.7.0
Local Policy Route-Maps	16.01.01
Network Monitoring	R11.2.0
NTP Client/Server	R10.7.0
OSPFv3	R10.8.0
Ping IPv4	16.01.01
Ping IPv6	18.01.01
SNMP Server	R10.8.0
SSH Server	R10.3.0
Static Host IPv4	16.01.01
Static Host IPv6	18.03.01
Syslog Messages	R11.4.0
Telnet Server	R10.3.0
Telnet Client	16.01.01
TFTP IPv4/IPv6	R10.3.0
Traceroute	16.01.01
UDP Relay	R10.10.0

Session Initiation Protocol (SIP) became VRF-aware with AOS release R10.5.0. SIP can be configured on either the default VRF or named VRF, but can only have one instance running at a time. SIP does not support multiple VRF instances at the same time.

MGCP is restricted to the default VRF. Therefore, SIP can only be used on the default VRF when used in combination with MGCP.

There are multiple features that are available only on the default unnamed VRF. For these features, traffic entering on a nondefault VRF will be ignored. Traffic for these features originating from the unit will only use the default VRF. The restricted features are:

- Media Gateway Control Protocol (MGCP)
- Multicast routing
- Protocol Independent Multicast (PIM) sparse
- Routing Information Protocol (RIP)
- Simple Network Time Protocol (SNTP) client and server

Configuring Multi-VRF in AOS

The following steps are required to implement multi-VRF in AOS:

1. Create a nondefault VRF instance.
2. Configure interfaces and assign them to the VRF instance.
3. Configure routing for the VRF instance.

The following sections detail each of the steps listed above with an explanation and the required syntax for configuration.

Step 1: Create a Nondefault VRF Instance

Since all AOS products supporting multi-VRF have a default (unnamed) VRF instance already in place, it is necessary to create and name a nondefault VRF instance in order to use this feature. Multiple VRF instances provide separation of traffic on one router, and must share the same available hardware resources. Only eight configured VRFs are allowed in addition to the default (unnamed) VRF.

To activate the nondefault VRFs, a route distinguisher must also be defined. A route distinguisher is an 8-byte number that is prepended to an IPv4 address, allowing differentiation between overlapping addresses. The combination of the route distinguisher plus the IPv4 address is referred to as the VPNv4 address.

From the Global Configuration mode, create a VRF instance with a route distinguisher using the following command:

(config)#**vrf** <*name*> **route-distinguisher [as-2byte** <*ASN:nn*> **| as-4byte** <*ASN:nn*> **| ip** <*ipv4 address:nn*>**]**

All nondefault VRFs must have a route distinguisher specified. The parameters **as-2byte** <*ASN:nn*> and **as-4byte** <*ASN:nn*> specify the AS number-relative route distinguisher as a 16-bit AS number and a 32-bit arbitrary number (2 byte), or that the AS number-relative route distinguisher is a 32-bit AS number and a 16-bit arbitrary number (4 byte). Route distinguishers are entered in the *ASN:nn* format, where *ASN* is either the 16-bit (2 byte) or 32-bit (4 byte) AS number, and *nn* is either a 32-bit arbitrary number (2 byte) or a 16-bit arbitrary number (4 byte). The **ip** <*ipv4 address:nn*> parameter specifies an IPv4 address-relative route distinguisher, which consists of an IPv4 address and a 16-bit arbitrary number (nn). The route distinguisher 0:0 or 0.0.0.0:0 is reserved for the default (unnamed) VRF instance and cannot be reassigned.

Step 2: Configure Interfaces and Assign to the VRF Instance

An IP interface can be tied to only one VRF instance. Incoming IPv4 traffic on an interface must be forwarded using that interface's VRF forwarding table. An interface will only forward IPv4 traffic that matches its associated VRF. IPv4 addresses between interfaces on different VRFs are allowed to overlap, but they cannot overlap between interfaces on the same VRF. By default, interfaces are assigned to the default (unnamed) VRF instance. The interfaces that will support VRF are listed in *Hardware and Software Requirements and Limitations on page 2*.

It is important to realize that changing the VRF association for an interface will clear all its IP-related settings. These settings will have to be reconfigured once the VRF is changed. The following list highlights most of the settings that are cleared or will need to be reconfigured when you change the interface's VRF association:

- Access control policy settings
- Dynamic domain naming system (DNS) settings
- Fast forwarding engine (FFE) settings
- IPv4 helper address
- Policy-based routing (PBR) settings
- Primary and secondary IPv4 addresses
- Proxy Address Resolution Protocol (ARP) enable/disable
- Route cache enable/disable
- Source interface for VRF aware applications

From the desired interface command prompt, assign the interface to the nondefault VRF instance by following these steps.

1. Enter the appropriate interface command mode:

   (config)#**interface** <*interface*>

   The <*interface*> parameter specifies an interface in the format <*interface type [slot/port | slot/port.subinterface id | interface id | interface id.subinterface id]*>. For example, for an ATM subinterface, use **atm 1.1**; for a demand interface, use **demand 6**; for an Ethernet subinterface, use **eth 0/1.1**; for a PPP interface, use **ppp 1**; and for a VLAN interface, use **vlan 33**.

2. Assign the interface to the VRF instance using the following command:

   (config-interface)#**vrf forwarding** <*name*>

The <*name*> parameter specifies the name of the VRF instance.

> *In AOS firmware release 17.8, the* ***ip*** *keyword was removed from the* ***vrf forwarding*** *command.*

In the following example, the Ethernet subinterface **eth 0/1.1** is associated with the VRF named **RED**. It is assumed that the VRF instance RED was created in Step 1:

```
>enable
#configure terminal
(config)#interface eth 0/1.1
    (config-eth 0/1.1)#vrf forwarding RED
    (config-eth 0/1.1)#ip address 192.168.24.1 /21      <-------- assigns IPv4 address to the interface
                                                                  since associating the interface with a VRF
                                                                  instance will clear any previous IP-related
                                                                  settings.
```

Step 3: Configure Routing for the VRF Instance

Each VRF instance has its own separate forwarding table. The VRF must already be created before static routes can be defined in the forwarding table. Variables for this command are explained in *Table 2 on page 6*.

From the Global Configuration mode, configure static routes for each VRF instance:

(config)#**ip route vrf** <*name*> <*ipv4 address*> <*subnet mask*> [<*interface*> | <*ipv4 address*> | **null 0**]
<*distance*> **track** <*track name*> **tag** <*number*>

Table 2. Static Route Command Variables

Variable	Explanation
<*name*>	Specifies the name (up to 79 alphanumeric characters) of the previously created VRF instance.
<*ipv4 address*>	Specifies the network address to add to the route table. IPv4 addresses should be expressed in dotted decimal notation (for example, **10.10.10.1**).
<*subnet mask*>	Specifies the subnet mask that corresponds to a range of IPv4 addresses (network) or a specific host. Subnet masks can be expressed in dotted decimal notation (for example, **255.255.255.0**) or as a prefix length (for example, ***/24***). Valid prefix lengths are **0** to **32**.
[<*interface*> **|** <*ipv4 address*>]	Specifies the next hop IPv4 address or an egress interface in the unit. Use the **ip route** <*ipv4 address*> <*subnet mask*> **?** command to display a complete list of egress interfaces. IPv4 addresses should be expressed in dotted decimal notation (for example, **10.10.10.1**).
null 0	Routes traffic destined for the specified network to the null interface. The router drops all packets destined for the null interface. Used to allow the router to advertise a route, but not forward traffic to the route.

Table 2. Static Route Command Variables *(Continued)*

Variable	Explanation
<*distance*>	Optional. Specifies an administrative distance associated with a particular router used to determine the best route when multiple routes to the same destination exist. The smaller the administrative distance, the more preferable the route. Range is **1** to **255**.
track <*track name*>	Optional. Enables tracking on the indicated route. Once the named track enters a fail state, the route specified by the command is disabled and traffic will no longer be routed using that route. For more information on configuring tracks, refer to *Configuring Network Monitor in AOS*, available online at https://supportforums.adtran.com.
tag <*number*>	Optional. Specifies a number to use as a tag for this route. Route tags are used to internally label and filter routes when dynamically redistributing routes into a routing protocol (such as RIP/OSPF/BGP). Range is **1** to **65535**.

In the following example, a static route is configured for the VRF named **RED** on the PPP interface:

>**enable**
#**configure terminal**
(config)#**ip route vrf RED 10.0.0.0 255.0.0.0 ppp 1**

Optional Settings

The following additional settings are optional, but could be necessary for your situation. These instructions only guide you with settings specific to multi-VRF. For more specific details about these features, refer to the appropriate configuration guide or the *AOS Command Reference Guide* available online at https://supportforums.adtran.com.

Enabling a Firewall

Firewall settings are applied globally across all VRF instances and cannot be changed on an individual VRF basis. However, the IPv4 firewall can be enabled or disabled independently for each VRF. For example, if session policy timeouts are set to 900 seconds, this will be the setting on all VRF instances with the firewall enabled. The maximum number of firewall associations on the unit is shared among all VRF instances. The maximum number of access control policy (ACPs) (20 user defined) is shared among all VRF instances. ACPs can be applied to multiple interfaces in different VRFs.

To enable the IPv4 firewall for a VRF, use the following command from the Global Configuration mode:

(config)#**ip firewall vrf** <*name*>

*Executing the command **ip firewall** (without the **vrf** keyword) will enable the IPv4 firewall for the default VRF only.*

In the following example, the IPv4 firewall is enabled on the VRF named **RED**:

>**enable**
#**configure terminal**
(config)#**ip firewall vrf RED**

Further details for configuring firewalls are provided in the *Configuring the Firewall (IPv4) AOS* and the *AOS Command Reference Guide* available online at https://supportforums.adtran.com.

Applying a Route Map to Router Traffic

You can also configure routing policies for traffic generated by each VRF instance. This can be useful for certain AOS features that are supported on a VRF instance other than the default. The route map will only apply to traffic originating on the router from the specified VRF instance. There is only one local policy for each VRF instance.

Before a route map can be specified, it must first be defined using the **route-map** command. For more information, refer to the *AOS Command Reference Guide* available online at https://supportforums.adtran.com. Once the route map is defined, apply the route map to local router traffic (be sure to specify the correct VRF instance) with this Global Configuration mode command:

(config)#**ip local policy route-map** *<mapname>* **vrf** *<name>*

Using NAT with Multiple VRF Instances

Traffic can change from one VRF instance to another by using a network address translation (NAT) entry in an ACP. The destination VRF instance is determined by the ACP entry's destination. When the destination is specified as an IPv4 address, the VRF name must be indicated in the ACP entry. However, when the destination is specified as an interface, the VRF instance for that interface is implied and is not necessary in the ACP entry.

Additional information is available about configuring NAT in the AOS Command Reference Guide available online at https://supportforums.adtran.com.

To use NAT with multiple VRF instances, use the following steps:

1. Create an IPv4 ACP and enter its configuration mode:

   (config)#**ip policy-class** *<ipv4 acp name>*

2. Add a NAT entry for traffic that matches the named IPv4 access control list (ACL). This command also specifies whether to use source NAT or destination NAT, depending upon your application. Using the **nat source list** command translates the source IPv4 address to a specified IPv4 address (or to the primary IPv4 address of the specified interface). Using the **nat destination list** command translates the destination IPv4 address to a specified IPv4 address.

a. To specify the source IPv4 address, use the following command (it is important to include the VRF name):

(config-policy-class)#**nat source list** <*ipv4 acl name*> [**address** <*ipv4 address*> **vrf** <*name*>] **overload** [**no-alg**] [**policy** <*ipv4 acp name*>]

To specify the primary IPv4 address of an interface, use the following command:

(config-policy-class)#**nat source list** <*ipv4 acl name*> [**interface** <*interface*>] **overload** [**no-alg**] [**policy** <*ipv4 acp name*>]

As shown above, there are two different ways to configure source NAT within an ACP: **nat source list** <*ipv4 acl name*> **address** or **nat source list** <*ipv4 acl name*> **interface**. The VRF must be entered when using the **address**. If the **interface** is used instead of the **address**, the VRF is not needed.

In the following example, NAT is configured to translate the source IPv4 address:

(config-policy-class)#**nat source list MATCHALL address 208.61.209.1 vrf RED overload**

In the following example, NAT is configured the source IPv4 address to the primary IPv4 address of the **ppp 1** interface:

(config-policy-class)#**nat source list MATCHALL interface ppp 1 overload**

b. To specify the IPv4 address to which to translate the original destination IPv4 address (and optionally port), use the following command (including the VRF name):

(config-policy-class)#**nat destination list** <*ipv4 acl name*> **[address** <*ipv4 address*> **vrf** <*name*>] [**port** <*number*>] [**no-alg**]

The IPv4 firewall must be enabled on all VRFs involved in NAT. Disabling the IPv4 firewall on a VRF prevents NAT, attack checking, and any other firewall processing from being applied to traffic.

3. Apply the IPv4 ACP to the interface. To apply the IPv4 ACP to an interface, enter the **access-policy** command from the appropriate interface configuration mode.

(config-interface)#**ip access-policy** <*ipv4 acp name*>

*In AOS firmware release 17.9, the command syntax for the **access-policy** command was changed to include the **ip** keyword.*

In the following example, an IPv4 ACP named **TRUSTED** is created. NAT is enabled for traffic that matches the IPv4 ACL **MATCHALL**, and the source IPv4 address is translated to **208.61.209.1** on VRF **RED**. The ACP is applied to the Ethernet 0/1 interface, which lies in the default VRF implicitly.

In this example, it is assumed that the interface assigned the IPv4 address **208.61.209.1** is assigned to VRF **RED**. Traffic entering the Ethernet 0/1 interface is checked by the **TRUSTED** ACP and matched by the **MATCHALL** ACL (not shown). The traffic then transitions to VRF **RED** using a translated address of **208.61.209.1**.

>**enable**
#**configure terminal**
(config)#**ip policy-class TRUSTED**
(config-policy-class)#**nat source list MATCHALL address 208.61.209.1 vrf RED overload**
(config-policy-class)#**exit**
(config)#**interface ethernet 0/1**
(config-eth 0/1)#**ip access-policy TRUSTED**

An IPv4 ACP entry can define a destination ACP using the **policy** keyword. Because the destination ACP is examined before the packet changes VRF instances, the defined ACP must be applied to the default route's destination interface in the source VRF instance. The actual outgoing interface in the destination VRF instance does not need to have the same ACP, or even need an ACP applied to it. Defining the default route is explained in Step 4.

In the following example, the NAT source statement belongs in the **Private** ACP, and defines **PUBLIC-1** as the destination IPv4 ACP with the **policy** keyword. Both the **Private** and **PUBLIC-1** ACPs must be applied to interfaces in the same VRF instance for the rule to be processed. IPv4 ACP **Private** must be applied to the interface where incoming traffic is to be translated. IPv4 ACP **PUBLIC-1** must be applied to the route destination interface.

(config-vlan 1)#**ip policy-class Private**
(config-policy-class)#**nat source list TRAFFIC interface ppp 1 overload policy PUBLIC-1**

4. Define the default route on the source VRF instance. A default route is required in the IPv4 route table of the source VRF instance since the first route lookup is performed exclusively from this route table. The IPv4 firewall will drop the packet if no route can be found for the destination of the packet. If a default route exists out of the same interface on which traffic enters, **ip firewall check reflexive-traffic** must be enabled. Conversely, if reflexive traffic is not allowed by the firewall, default routes cannot go out of the same interface on which traffic enters. If a destination interface does not already exist on the source VRF instance, a loopback interface can be used as the route destination.
Use the following command to define the default route:

(config)#**ip route vrf** <*name*> <*ipv4 address*> <*subnet mask*> <*interface*>

Use the following command to create a loopback interface:

(config)#**interface loopback** <*interface id*>

Use the following command to place the interface in the correct VRF:

(config-loop 1)#**vrf forwarding** <*name*>

Use the following command to assign an arbitrary IPv4 address to the interface:

(config-loop 1)#**ip address** <*ipv4 address*> <*subnet mask*>

Use the following command to configure the default route to the loopback interface:

(config)#**ip route vrf** *<name> <ipv4 address> <subnet mask>* **loop** *<interface id>*

In the following example, a loopback interface 1 is created and assigned to the VRF **LAN**. The IPv4 address **172.16.1.1** is assigned to the interface. The last line defines the default route to the loopback interface 1 on VRF **LAN**:

(config)#**interface loopback 1**
(config-loop 1)#**vrf forwarding LAN**
(config-loop 1)#**ip address 172.16.1.1 255.255.255.255**
(config-loop 1)#**exit**
(config)#**ip route vrf LAN 0.0.0.0 0.0.0.0 loop 1**

Using DHCP Server on Nondefault VRFs

The AOS product can be configured as a Dynamic Host Configuration Protocol version 4 (DHCPv4) server, assigning IPv4 addresses when requested by network devices. The DHCPv4 server provides address resources to all VRF instances through the use of DHCPv4 address pools. Each DHCPv4 address pool can be associated with a single VRF instance. Optionally, excluded IPv4 addresses and IPv4 address ranges can be added to the DHCPv4 address pool to restrict certain addresses from being assigned.

Create the address pool and enter the IPv4 DHCP Pool Configuration mode:

(config)#**ip dhcp-server pool** *<name>*

Next, associate the pool to a nondefault VRF instance:

(config-dhcp)#**vrf** *<name>*

At this point, you can further configure the settings for this particular DHCPv4 address pool. When a pool is assigned to a VRF instance, it loses any previously configured settings. For more detailed instructions on configuring DHCPv4, refer to the technical support notes located in ADTRAN's support community, or the *AOS Command Reference Guide* available online at https://supportforums.adtran.com.

Optionally, you can define an IPv4 address or range of addresses to exclude from being assigned by the DHCPv4 server. This command is entered from the Global Configuration mode.

Specify an IPv4 address to exclude:

(config)#**ip dhcp-server excluded-address vrf** *<name> <ipv4 address>*

or specify an IPv4 address range to exclude:

(config)#**ip dhcp-server excluded-address vrf** *<name> <start ipv4 address>*
<end ipv4 address>

In the following example, a DHCPv4 address pool is created named **PRIVATE** and its configuration mode is entered. This pool is then associated to the VRF instance named **RED**. When network devices request an IPv4 address and are connected through an interface on the **RED** VRF instance, the DHCPv4 server will lease an IPv4 address from the **PRIVATE** address pool. The address given out is defined by the network statement as any address from **10.22.199.1** to **10.22.199.254**.


```
>enable
#configure terminal
(config)#ip dhcp-server pool PRIVATE
(config-dhcp)#vrf RED
WARNING!!! All settings for this pool have been removed
(config-dhcp)#network 10.22.199.0 255.255.255.0
```

In the following example, the IPv4 addresses in the range **10.20.199.1** to **10.20.199.9** are added to the excluded address list for any DHCPv4 pool associated with the VRF **RED**, which in this case is the **PRIVATE** address pool. The configured excluded addresses will not be issued on this VRF instance.

```
>enable
#configure terminal
(config)#ip dhcp-server excluded-address vrf RED 10.20.199.1 10.20.199.9
```

Example Configurations

The following examples describe some of the common real-world applications of multi-VRF. All examples are configured using the command line interface (CLI). The configuration parameters entered in these examples are sample configurations only. You should configure these applications in a manner consistent with the needs of your particular network. CLI prompts have been removed from the configuration examples to provide you with a method of copying and pasting configurations directly from this guide into the CLI. You should not copy these configurations without first making the necessary adjustments to ensure they will function properly in your network.

Example 1: Multi-VRF with External IPv4 Firewall

This example illustrates a multi-VRF scenario using one CE router, providing customer access to the Internet on a nondefault VRF named BLUE. The default VRF provides traffic routing for all three customer sites. All customer traffic must pass through an external firewall that is provided by a third-party firewall application.

Example 1 can also be configured using PBR. Additional information is available in Configuring Policy-Based Routing in AOS online at https://supportforums.adtran.com.

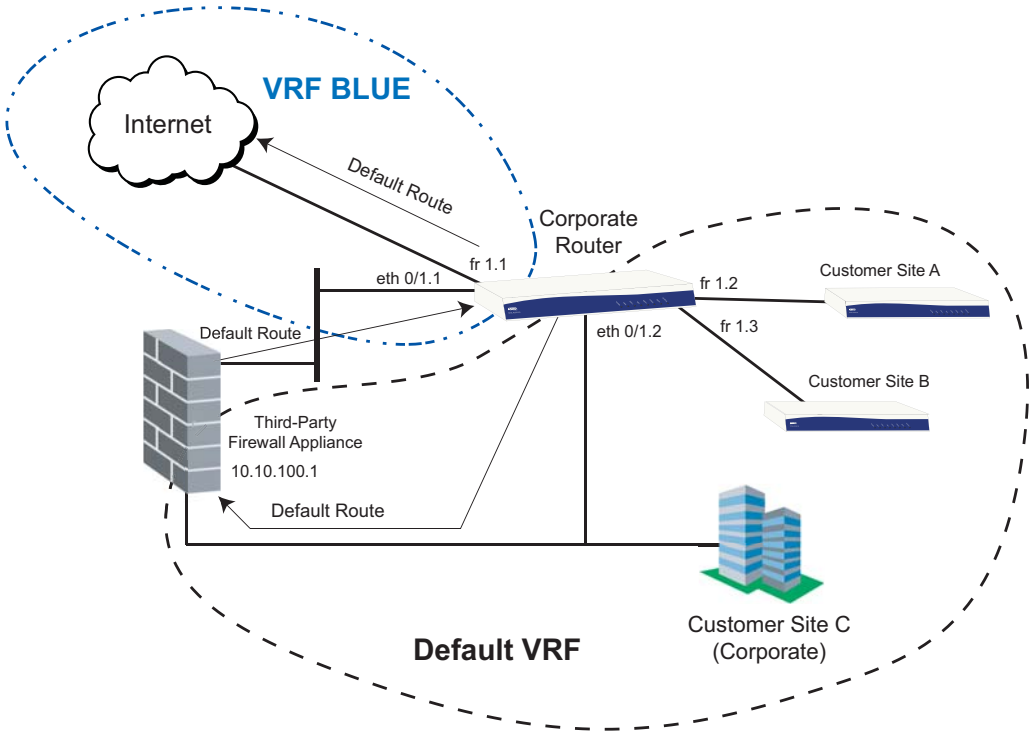

Figure 1. Example of External Firewall Application

The following commands are entered on the router labeled **Corporate** in *Figure 1*. Some configuration items not directly related to multi-VRF have been omitted.

Corporate Router

```
vrf BLUE route-distinguisher as-2byte 2:2
!
interface eth 0/1.1
   vrf forwarding BLUE
   ip address 10.10.1.1 255.255.255.0
!
interface fr 1.1
   vrf forwarding BLUE
   ip address 10.10.2.1 255.255.255.0
!
ip route vrf BLUE 0.0.0.0 0.0.0.0 fr 1.1
ip route 0.0.0.0 0.0.0.0 10.10.100.1
!
```

> **NOTE** *Each interface in this example has the IPv4 address assigned to the interface after applying the VRF. This is necessary since associating the interface with a VRF instance will clear any previous IP-related settings.*

> **NOTE** *The interfaces associated with the default VRF do not require any additional association. All interfaces automatically reside in the default VRF. In Example 1, this applies to interfaces **fr 1.2**, **fr 1.3**, and **eth 0/1.2**.*

> **NOTE** *In AOS firmware release 17.8, the syntax for the commands **vrf** <name> and **vrf forwarding** was changed to remove the **ip** keyword.*

Example 2: Multi-VRF with Separation of Traffic Using Two VRFs

This example illustrates a multi-VRF scenario using one CE router, providing separation of customer traffic through two nondefault VRF instances. The VRF instances are named ENGINEERING and SALES. Customer Site A receives its data traffic through the VRF **ENGINEERING** while Customer Site B receives its data traffic through the VRF **SALES**. This allows for duplication and reuse of the same IPv4 addresses on both VRFs while keeping the traffic separate.

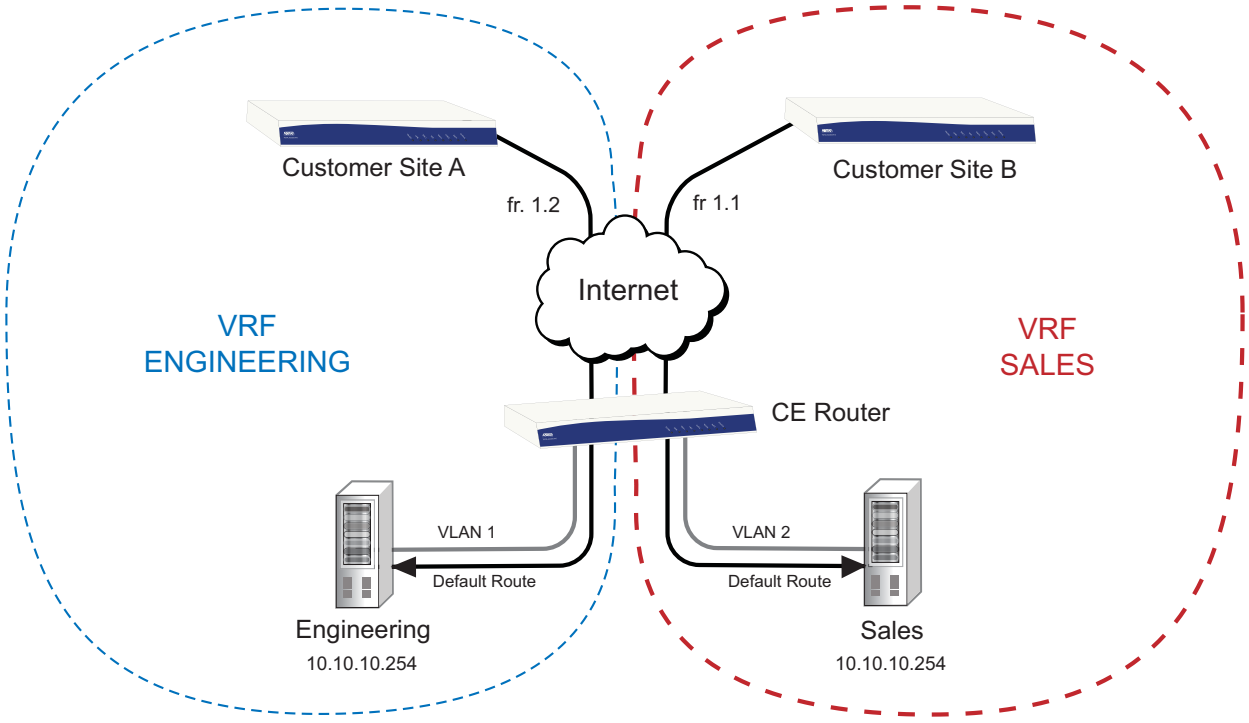

Figure 2. Example of Traffic Separation Using Two VRFs

The following commands are entered to configure multi-VRF for Example 2. Some configuration items not directly related to multi-VRF have been omitted.

CE Router

```
vrf ENGINEERING route-distinguisher as-2byte 2:2
vrf SALES route-distinguisher as-2byte 3:3
!
interface switchport 0/1
   no shutdown
   switchport access vlan 1
!
interface switchport 0/2
   no shutdown
   switchport access vlan 2
!
interface vlan 1
   vrf forwarding ENGINEERING
   ip address 10.10.10.1 255.255.255.0
!
interface vlan 2
   vrf forwarding SALES
   ip address 10.10.10.1 255.255.255.0
!
interface fr 1.1
   vrf forwarding SALES
   ip address 10.10.20.1 255.255.255.252
!
interface fr 1.2
   vrf forwarding ENGINEERING
   ip address 10.10.20.1 255.255.255.252
!
ip route vrf ENGINEERING 0.0.0.0 0.0.0.0 10.10.10.254
ip route vrf SALES 0.0.0.0 0.0.0.0 10.10.10.254
!
```

Each interface in this example has the IPv4 address assigned to the interface after applying the VRF. This is necessary since associating the interface with a VRF instance will clear any previous IP-related settings.

*The command **switchport access vlan 1** is enabled by default. Therefore, it will not appear in the output when the **show running-config** command is issued.*

*In AOS firmware release 17.8, the syntax for the commands **vrf** <name> and **vrf forwarding** was changed to remove the **ip** keyword.*

Example 3: Multi-VRF with NAT

In this example, multi-VRF is used to isolate different networks from each other, as well as provide Internet access using NAT across VRF instances (see *Figure 3 on page 17*). The VRF **CORP** (corporate network) and the VRF **GUEST** (guest network) are isolated from each other. Using separate VRF instances allows the same subnet (192.168.1.0 /24) to be used in each instance without causing a problem with overlapping addresses. Only hosts in the corporate network are able to manage the AOS router since its IPv4 firewall policy (**ip policy-class CORP**) allows only administrative access.

Both the corporate and guest networks automatically assign IPv4 addresses using DHCPv4. The **GUEST** and **CORP** DHCPv4 pools are configured in VRF **GUEST** and **CORP**, respectively. Hosts using DHCPv4 on the corporate network will be directed to an internal DNS server. Hosts on the guest network will use a DNS server assigned by the Internet service provider (ISP). When using multi-VRF, the AOS router should not be used as the DNS server (using domain-proxy), unless the name server to which the router is a proxy, is also located in the same VRF instance as the hosts sending the DNS requests. Consequently, the DHCPv4 pools assign an external DNS server instead of advertising the AOS router as the DNS server.

The Internet and other hosts on the demilitarized zone (DMZ) are accessible from the corporate and guest networks by traversing VRF instances to the VRF **INTERNET**. This access is accomplished with the NAT statements in each IPv4 ACP (**ip policy-class CORP** and **ip policy-class GUEST**). Both the guest and corporate networks lack a public interface that can be used for Internet access, making it necessary to NAT traffic to VRF **INTERNET**. In order for their traffic to have an exit point out of their current VRF, a default route is needed. Since there is no valid interface, a default route is created directing traffic to a loopback interface in the same VRF instance with an arbitrary IPv4 address. This method allows the firewall to process packets according to the rules in the ACP and allow NAT to take place. A route to the null 0 interface or to an IPv4 address assigned to any interface in the same VRF will also work, but only if there are no firewall rules containing a catchall IPv4 ACL directed to the self IPv4 ACP used for administrative access.

VRF **INTERNET** facilitates Internet access, as well as creates a DMZ where a public web server resides. Since the server is in a separate VRF instance from the guest and corporate networks, any traffic sourced from the DMZ is automatically isolated to prevent it from accessing internal resources. Traffic sourced from any hosts in the DMZ only have the ability to communicate with other hosts in the DMZ or out to the Internet. However, traffic sourced from the corporate or guest networks is able to communicate with hosts in the DMZ by being translated into VRF **INTERNET**. The AOS IPv4 firewall allows responses from the DMZ back to the corporate and guest networks only if it is sourced from the corporate or guest networks. For instance, anyone on the corporate or guest network is able to access the web server but, if the web server is compromised by a hacker, it will not have access to anything on corporate or guest networks because it is in a separate VRF instance. There is no NAT statement in the DMZ firewall ACP that allows the VRF instance to change.

> *Firewall rules can be used to create a DMZ rather than using the multi-VRF feature. For more information on configuring a DMZ, refer to Configuring a DMZ in AOS available online at https://supportforums.adtran.com.*

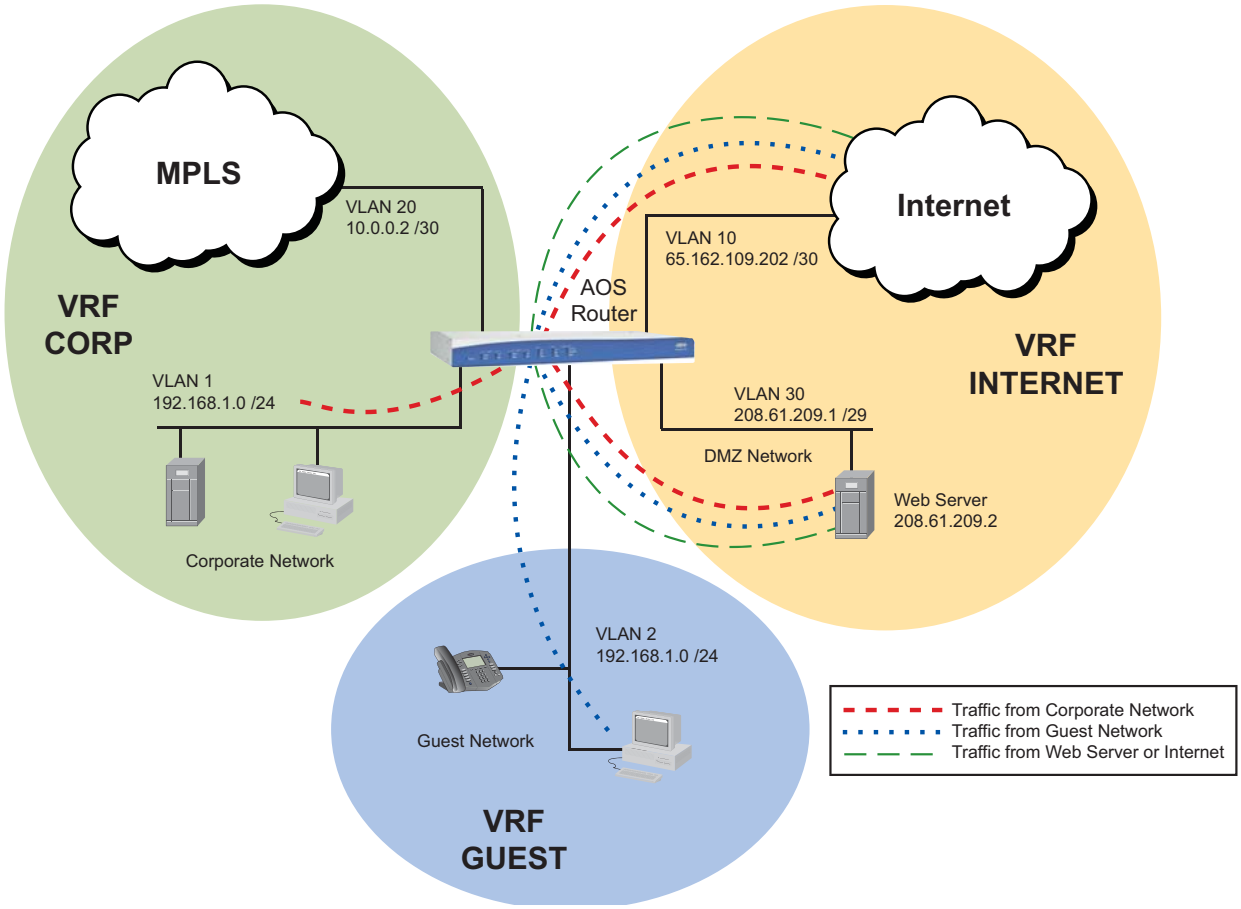

Figure 3. Example of Traffic Separation Using NAT

The following commands are entered on the AOS router shown in *Figure 3*. Some configuration items not directly related to multi-VRF have been omitted.

```
!
vrf INTERNET route-distinguisher as-2byte 1:1
vrf GUEST route-distinguisher as-2byte 2:2
vrf CORP route-distinguisher as-2byte 3:3
!
!
ip firewall vrf INTERNET
ip firewall vrf GUEST
```

```
ip firewall vrf CORP
!
!
ip dhcp-server pool "GUEST"
 vrf GUEST
 network 192.168.1.0 255.255.255.0
 dns-server 192.0.2.35
 default-router 192.168.1.1
!
ip dhcp-server pool "CORP"
 vrf CORP
 network 192.168.1.0 255.255.255.0
 dns-server 192.168.1.250
 default-router 192.168.1.1
!
interface loop 1
 vrf forwarding GUEST
 ip address 172.16.1.1 255.255.255.255
 no shutdown
!
interface loop 2
 vrf forwarding CORP
 ip address 172.16.2.1 255.255.255.255
 no shutdown
!
interface switchport 0/1
 description Internet
 no shutdown
 switchport access vlan 10
!
interface switchport 0/2
 description Corporate MPLS
 no shutdown
 switchport access vlan 20
!
interface switchport 0/3
 description DMZ
 no shutdown
 switchport access vlan 30
!
```



```
interface switchport 0/4
  description Corporate LAN
  no shutdown
  switchport access vlan 1
!
interface switchport 0/5
  description Guest LAN
  no shutdown
  switchport access vlan 2
!
interface vlan 1
  description Corporate LAN
   vrf forwarding CORP
  ip address 192.168.1.1 255.255.255.0
  ip access-policy CORP
  no shutdown
!
interface vlan 2
  description Guest LAN
  vrf forwarding GUEST
  ip address 192.168.1.1 255.255.255.0
  ip access-policy GUEST
  no shutdown
!
interface vlan 10
  description Internet
  vrf forwarding INTERNET
  ip address 65.162.109.202 255.255.255.252
  ip access-policy INTERNET
  no shutdown
!
interface vlan 20
  description Corporate MPLS
  vrf forwarding CORP
  ip address 10.0.0.2 255.255.255.252
  ip access-policy CORP
  no shutdown
!
```

```
interface vlan 30
  description Corporate DMZ
  vrf forwarding INTERNET
  ip address 208.61.209.1 255.255.255.248
  ip access-policy DMZ
  no shutdown
!
!
ip access-list standard MATCHALL
  permit any
!
ip access-list extended FROM-DMZ
  permit ip 208.61.209.0 0.0.0.7 any
!
ip access-list extended TO-DMZ
  permit ip any 208.61.209.0 0.0.0.7
!
ip access-list extended CORP-TO-MPLS
  permit ip 192.168.1.0 0.0.0.255 192.168.0.0 0.0.255.255
!
ip access-list extended MPLS-TO-CORP
  permit ip 192.168.0.0 0.0.255.255 192.168.1.0 0.0.0.255
!
!
ip policy-class CORP
  allow list MATCHALL self
  allow list CORP-TO-MPLS stateless
  allow list MPLS-TO-CORP stateless
  nat source list TO-DMZ interface vlan 30 overload
  nat source list MATCHALL interface vlan 10 overload
!
ip policy-class DMZ
  allow list FROM-DMZ stateless
!
ip policy-class GUEST
  nat source list TO-DMZ interface vlan 30 overload
  nat source list MATCHALL interface vlan 10 overload
!
```

*The remote networks on the MPLS are in the 192.168.0.0 /16 subnet. The extended ACL **CORP-TO-MPLS** and **MPLS-TO-CORP** permit traffic to these remote networks.*


```
ip policy-class INTERNET
  allow list TO-DMZ policy DMZ
!
!
ip route vrf GUEST 0.0.0.0 0.0.0.0 loop 1
ip route vrf CORP 0.0.0.0 0.0.0.0 loop 2
ip route vrf INTERNET 0.0.0.0 0.0.0.0 65.162.109.201
ip route vrf CORP 192.168.0.0 255.255.0.0 10.0.0.1
!
```

Each interface in this example has the IPv4 address assigned to the interface after applying the VRF. This is necessary since associating the interface with a VRF instance will clear any previous IP-related settings.

*The command **switchport access vlan 1** is enabled by default. Therefore, it will not appear in the output when the **show running-config** command is issued.*

*In AOS firmware release 17.8, the syntax for the commands **vrf** <name> and **vrf forwarding** was changed to remove the **ip** keyword.*

Command Summary Tables

This table summarizes the commands necessary for basic configuration of multi-VRF:

Table 3. Basic Configuration Commands

Step	Command	Explanation
Step 1	(config)#**vrf** <*name*> **route-distinguisher [as-2byte** <*ASN:nn*> **| as-4byte** <*ASN:nn*> **| ip** <*ipv4 address:nn*>**]**	Create a nondefault VRF instance and assign a route distinguisher.
Step 2	(config)#**interface** <*interface*> (config-interface)#**vrf forwarding** <*name*>	Assign interfaces to the VRF instance. An interface can only be assigned to one VRF, but multiple interfaces can be assigned to the same VRF. An interface will only forward IP traffic that matches its associated VRF instance.This command must be executed from the interface configuration mode. It is important to realize that changing the VRF association for an interface will clear all its IP-related settings.
Step 3	(config)#**ip route vrf** <*name*> <*ipv4 address*> <*subnet mask*> [<*interface*> | <*ipv4 address*> | **null 0**] <*distance*> **track** <*track name*> **tag** <*number*>	Define static routes in the forwarding table for the VRF instance. An explanation for each of the variables in this command is available in *Table 2 on page 6*.

The following table summarizes optional commands that can be used to configure multi-VRF settings on your AOS product.

Table 4. Optional Configuration Commands

Options	Command	Explanation
Permit Telnet/SSH Sessions	(config)#**line** [**ssh** <*0-4*> | **telnet** <*0-4*>] [<*0-4*>] (config-line)#**ip access-class** <*ipv4 acl name*> **in any-vrf**	Applies an existing IPv4 ACL to limit sessions from any VRF instance based on traffic matching the ACL. This command must be executed from the Line (Telnet) or Line (SSH) Interface Configuration mode.
Enable the Firewall	(config)#**ip firewall vrf** <*name*>	Enables the IPv4 firewall on the specified VRF instance.
Apply a Route Map	(config)#**ip local policy route-map** <*mapname*> **vrf** <*name*>	Applies a route map to traffic generated by the router from a specific VRF instance.

Table 4. Optional Configuration Commands *(Continued)*

Options	Command	Explanation
Enable NAT for Source Traffic	(config)#**ip policy-class** <*ipv4 acp name*> (config-policy-class)#**nat source list** <*ipv4 acl name*> [**address** <*ipv4 address*> **vrf** <*name*>] [**interface** <*interface*>] **overload** [**no-alg**] [**policy** <*ipv4 acp name*>]	Creates an ACP and enters its configuration mode. Enables NAT for traffic that matches the named IPv4 ACL and changes the source address, which resides on the specified VRF instance. If the **interface** parameter is used, the traffic is translated to the VRF of the named interface.
Enable NAT for Destination Traffic	(config)#**ip policy-class** <*ipv4 acp name*> (config-policy-class)#**nat destination list** <*ipv4 acl name*> **address** <*ipv4 address*> [**vrf** <*name*>] [**port** <*number*>] [**no-alg**]	Creates an ACP and enters its configuration mode. Enables NAT for traffic that matches the named IPv4 ACL and changes the destination address, to the specified address, which resides in the specified VRF.
Create and Associate a DHCP Server Pool	(config)#**ip dhcp-server pool** <*name*> (config-dhcp)#**vrf** <*name*>	Associates a DHCPv4 server pool with the specified VRF instance. This command must be executed from the specific pool's configuration mode, which appears in the first line of the example.
Specify DHCP Server Pool Addresses to Exclude	(config)#**ip dhcp-server excluded-address vrf** <*name*> **[**<*ipv4 address*> |<*start ipv4 address*> <*end ipv4 address*>**]**	Adds IPv4 addresses to the DHCPv4 server pool to be excluded from assigned client leases. This only adds the excluded IPv4 address to the pool associated with the specified VRF instance.

Troubleshooting

After configuring multi-VRF, several different commands can be issued from the Enable mode in the CLI to assist in troubleshooting. The troubleshooting commands that use multi-VRF instances are provided in the following tables. *Table 5* contains the **clear** commands, *Table 6 on page 25* contains the **debug** commands, and *Table 7 on page 25* contains the **show** commands.

Troubleshooting Commands

Table 5. AOS Multi-VRF Clear Commands

Command	Explanation
clear host [**vrf** <*name*>] [<*hostname* | *>]	Clears one or all host names for the specified VRF from the DNS cache. If no VRF instance is specified, the command clears host names for the default (unnamed) VRF instance. The * clears all hosts from the host table.
clear ip cache [**vrf** <*name*>] [**counters**]	Clears all fast-cache entries for a particular VRF instance. If no VRF instance is specified, it clears cache entries for the default (unnamed) VRF instance. The optional **counters** keyword resets all the counters in the cache table.
clear ip dhcp-server binding [**vrf** <*name*> | * | <*ipv4 address*>]	Removes one or all address lease(s) for the specified VRF. If no VRF instance is specified, all leases are removed for the default (unnamed) VRF instance. The * clears all leases for the VRF instance. The <*ipv4 address*> parameter clears a specific DHCPv4 binding associated with an IPv4 address.
clear ip policy-sessions [**any-vrf** | **vrf** <*name*> <*ipv4 acp name*> | <*ipv4 acp name*>] [**ahp** | **esp** | **gre** | **icmp** | **tcp** | **udp** | <*protocol*>] <*ipv4 source*> <*source port*> <*ipv4 destination*> <*destination port*> **[destination | source]** <*nat ipv4*> <*nat port*>	Clears all IPv4 policy class sessions or a specific session for the specified VRF instance. If no VRF instance is specified, sessions are removed for the default (unnamed) VRF instance. Using the **any-vrf** parameter in this command will clear sessions for all VRF instances.
clear ip route [**vrf** <*name*>] [<*ipv4 address*> <*subnet mask*> | *]	Removes all routes in the specified VRF instance. If no VRF instance is specified, all routes are removed from the default (unnamed) VRF instance. The * clears all routes from the route table.

Table 6. AOS Multi-VRF Debug Commands

Command	Explanation
debug firewall [**vrf** <*name*>]	Displays debug messages associated with the IPv4 firewall for the specified VRF instance. If no VRF instance is specified, only events for the default (unnamed) VRF instance are displayed. If NAT is being used to change VRF instances, activate debugging for both VRF instances.
debug ip dhcp server [**vrf** <*name*>]	Displays debug messages associated with DHCPv4 server operation for the specified VRF instance. If no VRF instance is specified, only events for the default (unnamed) VRF instance are displayed.
debug ip icmp	Displays the Internet Control Message Protocol (ICMP) messages that are destined for the router or originated by the router. If the ICMP messages are destined for or originated from a nondefault VRF instance, then the VRF name is indicated in the output. If no VRF instance is named in the output, the default VRF is assumed.
debug ip packet [**vrf** <*name*> | **any-vrf** | <*ipv4 acl name*>] [**detail** | **dump**]	Displays per-packet information for the specified VRF or IPv4 ACL. If no VRF is specified, only information for the default (unnamed) VRF is displayed. Information for all VRF instances can also be displayed by adding the **any-vrf** parameter to the command.
debug ip policy <*ipv4 acl name*>	Displays IPv4 policy-based routing events.
debug ip routing [**vrf** <*name*>]	Displays IPv4 routing table events for the specified VRF. If no VRF is specified, only events for the default (unnamed) VRF are displayed.
debug ip tcp	Displays debug Transmission Control Protocol (TCP) events for the default VRF and nondefault VRF instances on the router.
debug ip udp	Displays debug User Datagram Protocol (UDP) send and receive events for the default VRF and nondefault VRF instances on the router.

Table 7. AOS Multi-VRF Show Commands

Command	Explanation
show arp [**vrf** <*name*>] [**realtime**] **show ip arp** [**vrf** <*name*>] [**realtime**]	Both of these commands display the ARP table for the specified VRF. If no VRF is given, output for the default (unnamed) VRF is displayed.
show hosts [**vrf** <*name*>]	Displays the DNS host name cache for a specific VRF. If no VRF is given, output for the default (unnamed) VRF is displayed.
show ip cache [**vrf** <*name*>]	Displays the fast-cache entries for a specific VRF. If no VRF is given, output for the default (unnamed) VRF is displayed.

Table 7. AOS Multi-VRF Show Commands *(Continued)*

Command	Explanation
show ip local policy [**vrf** <*name*>]	Displays the local PBR settings for a specific VRF. If no VRF is given, output for the default (unnamed) VRF is displayed.
show ip route [**vrf** <*name*>] [**bgp** | **connected** | **ospf** | **rip** | **static** | **summary** | **table** | <*ipv4 address*> | <*subnet mask*>]	Displays the route table for a specific VRF. If no VRF is given, output for the default (unnamed) VRF is displayed.
show vrf <*name*>	Displays all the configured VRFs, route distinguishers, and the interfaces associated with each. If a VRF is specified, only information about that one VRF is displayed.
show vrf interfaces <*name*>	Displays all the IP interfaces, their primary IPv4 address, their current status, and their VRF. If a VRF is specified, only interfaces in that VRF are displayed.
show tcp info	Displays TCP information, including the state, local IPv4 address, local port, remote IPv4 address, remote port, interface, and the VRF instance if it does not belong to the default VRF.
show udp info	Displays UDP session information, including the local IPv4 address and mask, remote IPv4 address and mask, socket, and VRF instance, if applicable.
show ip dhcp-server binding [**vrf** <*name*> | <*ipv4 address*>]	Displays the DHCPv4 lease table for the specified VRF instance. If no VRF is given, output for the default (unnamed) VRF is displayed.
show ip policy-sessions [**any-vrf** | **vrf** <*name*>] <*ipv4 acp name*> [**include-deleted**]	Displays a list of current IPv4 ACP associations for the specified VRF instance or specified IPv4 ACP. If no VRF is given, output for the default (unnamed) VRF is displayed. The active and deleted ACP associations can also be displayed by adding the optional **include-deleted** parameter to this command. Using **any-vrf** will display all associations for all VRFs.

Utility Commands

The following commands (**ping**, **telnet**, and **traceroute**) are utility commands that can be issued from the Enable or Basic mode in the CLI to assist in troubleshooting.

The **ping** command is used to send ICMP echo requests to an IPv4 address on the specified VRF. If no VRF is specified, the default (unnamed) VRF is assumed. This command is commonly used in troubleshooting to verify IPv4 connectivity to the destination.

#**ping vrf** <*name*> <*ipv4 address | hostname*>

The **telnet** command is used to open a Telnet session to another system on the network. The addition of the VRF name specifies the location of the IPv4 address or host name. If no VRF is specified, the default VRF is assumed.

#**telnet vrf** <*name*> <*ipv4 address | hostname*> **port** <*tcp port*>

The **traceroute** command displays IPv4 routes for a packet en-route to a specified destination on the specified VRF. If no VRF is specified, then the default VRF is assumed. This command is commonly used in troubleshooting to verify the traffic path through the IPv4 network.

#**traceroute vrf** *<name> <ipv4 address | hostname>* **source** *<ipv4 address>*

Additional Resources

There are additional resources available to aid in configuring your AOS unit. Many of the topics discussed in this guide are complex and require additional understanding. The documents listed below are available online at ADTRAN's Support Forum at https://supportforums.adtran.com.

- *AOS Command Reference Guide*
- *Configuring AAA in AOS*
- *Configuring a DMZ in AOS*
- *Configuring BGP in AOS 18.03/R10.1.0 or later*
- *Configuring DHCP in AOS*
- *Configuring the Firewall (IPv4) AOS*
- *Configuring IP Access Control Lists (ACLs) in AOS*
- *Configuring IPv6 in AOS*
- *Configuring Policy-Based Routing in AOS*
- *Configuring SNMP in AOS*
- *Understanding AOS Version Changes*