6AOSCG0073-29C
April 2016

ADTRAN

Configuration Guide

Configuring RapidRoute in AOS

This configuration guide provides an overview of RapidRoute functionality, feature-specific RapidRoute information, and command line interface (CLI) configuration examples for the RapidRoute flow bundling and service assurance features in ADTRAN Operating System (AOS) router products.

This guide contains the following sections:

Overview of RapidRoute

RapidRoute is the fast forwarding engine (FFE) used by AOS products to optimize the packet routing process for packets that go through repetitive sets of rules and procedures before being routed to their destination. The packet processing architecture classifies packets into flows based upon the IP protocol used by the packet, the source and destination IP address of the packet, the packet's protocol-specific information (such as source and destination Transmission Control Protocol (TCP)/User Datagram Protocol (UDP) port numbers), the packet's ingress interface, and the IP traffic class bits. Packet flows are defined as the unidirectional representation of a conversation between two IP hosts. Using this information, each configured ingress interface maintains a traffic flow table, with identifiers that are the same as those in the firewall association table, to provide continuity between the flow entry and the firewall's association selector. When packets are received, they can match a previous session and bypass duplicate processing because the router has cached the final routing decision for the packet. Only the initial packet in a flow is classified, and only once, rather than every time it is used by an AOS feature.

RapidRoute accelerates traffic flow classification, lookup, modification, and forwarding in order to achieve predictable and high throughput performance results for features such as IP routing, IP firewall, IP security (IPsec), and static packet filters. RapidRoute functions similarly in both Internet Protocol version 4 (IPv4) and IPv6, and is enabled by default on all IP interfaces.

RapidRoute Table Entries

The packet flow entries created when RapidRoute is enabled are only added to the traffic flow table when the first packet of a flow fails an FFE lookup and then travels through the regular data path. FFE entries can be one of three types: valid, ineligible, or rejected. Valid entries result in quickly forwarded matching packets, and are sent out the egress interface. Rejected entries result in quickly discarded matching packets. Rejected entries are usually added when there is no route to the destination, or the traffic flow is blocked by a static filter. Ineligible entries result in packets flowing through the regular data path. Packets that enter RapidRoute, but then travel through the regular data path are referred to as exception packets. Exception packets are any packets that RapidRoute cannot accelerate. The following criteria could make RapidRoute unable to process a packet:

- An IP packet does not match any existing FFE entry
- An IP packet matches an ineligible entry
- An IP packet's Next Header field is not a supported upper layer protocol
- An IP packet contains any header extensions
- The IP header is malformed
- The IP packet length does not match the packet length in the IP header
- The IP hop count is 1 or less

RapidRoute only removes entries from the FFE table when they have expired, or there have not been any packets that match an entry for its configured inactive timeout interval. You can manually remove FFE table entries by issuing a **clear** command from the Enable mode (refer to *Clearing RapidRoute Entries on page 5*).

The **clear** command removes all matching FFE entries. In addition, other routing or feature changes can cause FFE entries to be removed. The following are instances that cause FFE entries to be cleared:

- When a forwarding route changes or is removed from the route table, all FFE entries that might have been covered by this route are cleared.

- Whenever a Neighbor Discovery (ND) cache entry is removed or otherwise loses its IPv6 address to Medium Access Control (MAC) address binding, all FFE entries that use the ND cache entry's address as the gateway IPv6 address are cleared.
- Whenever an IP policy session is cleared from the firewall, all FFE entries that are associated with this firewall association are also cleared.
- If a route map is applied to a RapidRoute-enabled interface or an already-applied route map changes, then all FFE entries on that interface are cleared.
- If an access control list (ACL) that is applied to a RapidRoute-enabled interface as an inbound access group changes, then all FFE entries on that interface are cleared. If an ACL that is applied to any egress interface as an outbound access group changes, then all FFE entries on that interface, or that egress from that interface, are cleared.
- Whenever an interface's status changes from active to inactive, all FFE entries in that interface's FFE table are cleared. In addition, all FFE entries in other FFE tables that use this interface as the egress interface are cleared.
- Whenever an interface adds, removes, or modifies one of its IP addresses, then all the FFE entries on the interface are cleared.
- Whenever the maximum transmission unit (MTU) changes on an interface, all the FFE entries on that interface are cleared.
- Whenever IP forwarding is globally disabled, all FFE entries on all interfaces are cleared.

Hardware and Software Requirements and Limitations

RapidRoute was introduced in AOS firmware release 13.0, and is available on AOS products as outlined in the *AOS Product Feature Matrix* (available online at https://supportforums.adtran.com). As of AOS firmware release R10.4.0, RapidRoute is enabled on all supported IP interfaces by default.

RapidRoute flow bundling and service assurance features are available on AOS products running firmware release R11.10.0, as outlined in the *AOS Product Feature Matrix* (available online at https://supportforums.adtran.com). These features are available for both IPv4 and IPv6.

Both RapidRoute service features interact with many other AOS features; *Table 1 on page 12* outlines these interactions. A working knowledge of all other AOS features is assumed. For your convenience, a list of available documentation regarding AOS features that interface with RapidRoute are included in *Additional Resources on page 35*.

Flow bundling is enabled by default and does not require a command to enable it. In AOS firmware release R11.10.1, however, global commands were added to allow IPv4 and IPv6 flow bundling to be disabled if necessary.

Configuring RapidRoute in AOS

RapidRoute is configured in a similar manner for both IPv4 and IPv6. Both protocols rely on a global configuration of RapidRoute, as well as RapidRoute configuration on a per-interface basis. The following sections outline the commands used to configure RapidRoute from the CLI. Web configuration of RapidRoute is not supported.

Global RapidRoute Configuration (Optional)

RapidRoute global configuration consists of optionally configuring the maximum number of FFE entries allowed at a given time, the timeout value for FFE entries, and the limit of exception packets. RapidRoute for IPv4 also allows you to enable RapidRoute for IPsec security associations. To configure the global parameters of RapidRoute, follow these steps:

1. To specify the global maximum number of FFE entries allowed at a given time, enter the **ip ffe max-entries** <*value*> or **ipv6 ffe max-entries** <*value*> command from the Global Configuration mode. The **ip** parameter denotes RapidRoute is used with the IPv4 protocol, whereas the **ipv6** parameter denotes RapidRoute is used with the IPv6 protocol. The <*value*> parameter is the maximum number of entries allowed. Valid range is **1** to **500000** entries, with a default value of **16384** entries. Using the **no** form of this command returns the maximum number of entries to the default value. To change the maximum number of entries allowed for IPv6 RapidRoute, enter the command as follows:

   (config)**#ipv6 ffe max-entries 20000**

2. To specify the timeout values for FFE entries, based on IP protocol, enter the **[ip | ipv6] ffe timeout [ah | esp | gre | icmp | other | tcp | udp]** <*max timeout*> **[**<*inactive timeout*>**]** command from the Global Configuration mode. The **ip** and **ipv6** parameters specify whether RapidRoute is being used with IPv4 or IPv6, respectively. The **ah**, **esp**, **gre**, **icmp**, **tcp** and **udp** parameters specify the AH, ESP, GRE, ICMP, TCP, and UDP protocols, and the **other** parameter specifies a different protocol than those listed. The <*max timeout*> parameter is the maximum time (in seconds) that a given FFE entry stays in the table before being removed. The entry is removed when this time interval expires, even if traffic is matching the FFE entry. Valid range is **60** to **86400** seconds, with a default value of **1800** seconds. The optional <*inactive timeout*> parameter specifies the maximum time (in seconds) that a given FFE entry stays in the table if no traffic is matching it. Valid range is **10** to **86400** seconds, with a default value of **15** seconds. Using the **no** form of this command returns the timeout values to the default settings. To change the timeout values for IPv6 FFE entries, enter the command as follows:

   (config)**#ipv6 ffe timeout tcp 2000 30**

3. To specify a limit to the number of unhandled FFE exception packets allowed at any given time, enter the **[ip | ipv6] ffe limit exceptions** <*number*> command from the Global Configuration mode. Exception packets are any packets that RapidRoute cannot handle, for example, traffic that matches ineligible entries, fragmented packets, packets with header errors, or the first packet in a given traffic flow that is used to build an FFE entry. Once the limit of unhandled FFE exception packets is reached, subsequent exception packets are dropped until the previously unhandled exceptions are dealt with. The **ip** and **ipv6** parameters specify whether RapidRoute is being used with IPv4 or IPv6, respectively. The <*number*> parameter specifies the maximum number of unhandled FFE exception packets allowed at a given time. Valid range is **1** to **1024**, with a default value of **128** packets. Using the **no** form of this command returns the limit to the default value. To change the maximum number of allowable IPv6 exception packets, enter the command as follows:

   (config)**#ipv6 ffe limit exceptions 200**

4. To enable RapidRoute for IPv4 IPsec security associations, enter the **ip crypto ffe [max-entries** <*value*>**]** command from the Global Configuration mode. The optional **max-entries** <*value*> parameter specifies the maximum number of entries per inbound (decrypting) IPsec security association. Valid range is **1** to **8192** entries, with a default value of **4096** entries. Using the **no** form of this command disables RapidRoute for IPv4 IPsec security associations. By default, RapidRoute is not enabled for IPv4 IPsec.

To enable RapidRoute for IPv4 IPsec security associations, enter the command as follows:

(config)**#ip crypto ffe**

> NOTE *When using virtual private networks (VPN) ADTRAN recommends configuring RapidRoute for IPsec security associations using the **ip crypto ffe** command.*

Interface RapidRoute Configuration

RapidRoute is enabled by default at the interface. You can choose to disable (or later re-enable) RapidRoute using the **[no] [ip | ipv6] ffe [max-entries** <*value*>**]** command from the interface's configuration mode. Optionally, you can specify the maximum number of FFE entries supported by the interface by using the **max-entries** <*value*> parameter. Valid range is **1** to **500000** entries, with a default value of **4096**. Using the **no** form of this command disables RapidRoute on the interface. To disable RapidRoute on an IPv4 enabled interface, enter the interface's configuration mode and enter the command as follows:

(config)**#interface eth 0/1**
(config-eth 0/1)**#no ip ffe**

> NOTE *ADTRAN recommends that RapidRoute remain enabled on all IP interfaces.*

Clearing RapidRoute Entries

You can manually clear RapidRoute entries in the FFE table by entering the **clear [ip | ipv6] ffe [**<*interface*>**]** command from the Enable mode. This command clears all the FFE entries in a given ingress interface. If no interface is specified, then all IPv4 (**ip** keyword) or IPv6 (**ipv6** keyword) FFE entries in the AOS device are cleared. Specify interfaces in the <*interface*> <*slot/port* | *interface id*> format. To clear all IPv6 FFE entries, enter the command as follows:

>enable
#clear ipv6 ffe

Viewing RapidRoute Statistics

You can view RapidRoute statistics by entering the **show [ip | ipv6] ffe [destination** <*ip address*>**] [destination-port** <*port*>**] [egress** <*interface*>**] [icmp-type [echo | reply |** <*type*>**]] [ingress** <*interface*>**] [peak [history]] [protocol [ah | esp | fragment | gre | icmp | tcp | udp |** <*protocol*>**]] [source** <*ip address*>**] [source-port** <*port*>**] [type [ineligible | valid | rejected]] [wildcard [interface** <*interface*>**]] [details] [summary]** command from the Enable mode. This command can display the all current FFE entries for every interface that has RapidRoute enabled, or you can use the optional parameters to limit the output to FFE entries that match the entered criteria. The **details** keyword specifies that more details about each entry are given, and the **summary** keyword specifies that summarized statistics are displayed. The optional filtering parameters for the output are the following:

- The **ip** and **ipv6** keywords specify whether the output displayed is for IPv4 or IPv6 RapidRoute statistics.
- The **destination** <*ip address*> parameter filters output by a destination IP address. IPv4 addresses should be expressed in dotted decimal notation (**x.x.x.x**), for example, **10.10.10.1**. IPv6 addresses should be expressed in colon hexadecimal notation (**X:X:X:X::X**), for example, **2001:DB8:1::1**.
- The **destination-port** <*port*> parameter filters output by destination TCP or UDP port. Ports range from **0** to **65535**.
- The **egress** <*interface*> parameter displays FFE entries for an egress interface. Specify interfaces in the <*interface*> <*slot/port* | *interface id*> format.
- The **icmp-type** parameter displays FFE entries using a specific ICMP type. You can select **echo** to display ICMP echo entries, **reply** to display ICMP reply entries, or display another ICMP type by entering a number between **0** to **255**.
- The **ingress** <*interface*> parameter displays FFE entries for an ingress interface. Specify interfaces in the <*interface*> <*slot/port* | *interface id*> format.
- The **peak** parameter displays the current and peak count of RapidRoute sessions. Information is displayed for each eligible interface and the global values. The optional **history** parameter displays a graphical presentation of the peak global RapidRoute count per second for the last 60 seconds, the peak and average global RapidRoute count per minute for the last 60 minutes, and the peak and average global RapidRoute count per hour for the last 72 hours. Data is presented as a percentage of the value configured with the **[ip | ipv6] ffe max-entries** <*value*> command.

  NOTE *The **peak** parameter cannot be used in conjunction with any other parameters of this command.*

- The **protocol** parameter displays FFE entries that use a specific protocol. You can specify AH (**ah**), ESP (**esp**), FRAG (**fragment**), GRE (**gre**), ICMP (**icmp**), TCP (**tcp**), UDP (**udp**), or you can specify another protocol by entering a number between **0** and **255**.
- The **source** <*ip address*> parameter filters output by a source IP address. IPv4 addresses should be expressed in dotted decimal notation (**x.x.x.x**), for example, **10.10.10.1**. IPv6 addresses should be expressed in colon hexadecimal notation (**X:X:X:X::X**), for example, **2001:DB8:1::1**.
- The **source-port** <*port*> parameter filters output by source TCP or UDP port. Ports range from **0** to **65535**.
- The **type** parameter filters output by specific entry type. You can specify only ineligible entries are displayed (**ineligible**), only rejected entries are displayed (**rejected**), or only valid entries are displayed (**valid**).
- The **wildcard** parameter displays the field wildcards used on all IP interfaces for RapidRoute flow bundling. Field wildcards include: source IP address, IP precedence, IP Differentiated Services Code Point (DSCP), Layer 4 IP protocol, TCP/UDP source and destination ports, Internet Control Message Protocol (ICMP) type, Encapsulating Security Payload (ESP) Security Parameter Index (SPI), and Generic Routing Encapsulation (GRE) tunnel key. If the optional **interface** <*interface*> parameter is used, output is limited to wildcards used on the specified interface. Specify interfaces in the <*interface*> <*slot/port* | *interface id*> format.

  NOTE *The **wildcard** parameter cannot be used in conjunction with any other parameters of this command.*

Understanding Show Command Output

RapidRoute statistics are displayed using the **show** command as described in the previous section. There are several types of valuable information included in the **show [ip | ipv6] ffe** command, such as information about all of the current IP traffic running through the AOS unit that includes all of the sessions associated with each interface and their protocol, ports, and IP addresses, how long the session has been in use, how many packets have hit and not hit the policy, and the flags associated with the policy. The following section describes each type of output displayed after entering this command.

To display RapidRoute information for IPv4, enter the command from Enable mode as follows:

>enable
#show ip ffe

```
Timeout          TCP       UDP       ICMP        AH        ESP        GRE      Other
Age:          30m 0s    30m 0s    30m 0s    30m 0s    30m 0s    30m 0s    30m 0s
Inactive:        15s       15s       15s       15s       15s       15s       15s

Exceptions: 0/217/0 (current/max/drops)

Type:  * valid, ! ineligible, - rejected
Flags: F firewall, N NAT, T altered ToS, D don't fragment, I IPsec
       H hardware assist, i ingress-filter, e egress-filter

--------------------------------
Ingress: giga-eth 0/5.1
    0 hits, 89021 misses, 0 drops

T Proto ToS        Age       Used      Drops Flags
  Source
  Destination
  ---------------
! any   any        27s         49          0
  any
  10.100.13.25
Number of entries: 1 of 1  (4096 maximum)

--------------------------------
Ingress: system-management-evc
    0 hits, 165059 misses, 0 drops

T Proto ToS        Age       Used      Drops Flags
  Source
  Destination
  ---------------
! udp   dscp:0     23s          8          0
  10.255.10.1:123
  10.42.157.2:123
! udp   dscp:0      4s          0          0
  10.100.254.4:1812
  10.42.157.2:42238
Number of entries: 2 of 2  (4096 maximum)

--------------------------------
Total number of entries: 3 of 3  (32768 maximum)
```


Timeout Information

The **Timeout** section of the output displays the **Age** timeout information (maximum timeout), as well as the **Inactive** timeout information.

The **Age** field, or maximum timeout, displays how long an entry remains in the traffic flow table while the session is being used and active (before the session is refreshed). This setting keeps infrequent sessions out of the FFE table so that router functionality is not diminished. For example, if you are running an FTP transfer, which uses TCP port 21, the first packet of the flow takes the slow path through the router, but establishes a RapidRoute session. Each subsequent packet in the flow takes the fast path. Once 30 minutes has elapsed (the default age timeout), the session is eliminated. If the transfer is still running, the next packet after the 30 minute mark creates a new session.

The **Inactive** field, or inactive timeout, displays how long the FFE table will wait before eliminating an idle session. The default value is **15** seconds. For example, if you ping something on your network, RapidRoute creates a session for it. Typically, a ping only sends a few requests, and you do not want an idle session remaining in the FFE table for the full 30 minutes.

Both timeout values are configurable so that you can specify timeouts for the various protocols you might have running through the system. For example, if you mostly use ping as a troubleshooting tool, the inactive session time could be changed to one second so that if multiple pings are issued at once, the session will not be sitting in the active session table for the full 15 seconds. Conversely, if you have a UDP program running, such as a video streaming service that usually lasts for more than an hour, you can set the age timeout to 3600 seconds.

Entry Types and Associated Flags

The next section of the command output displays the entry type (**Type**) and flags associated with the RapidRoute session (**Flags**).

The **Type** field describes the type of session: valid, ineligible, or rejected. The ***** character indicates a valid entry type which means that the session is eligible for RapidRoute and is being used currently; most entries should be of this type. The **!** character, used for ineligible entries, indicates that the packet's path is not one that can be accelerated by RapidRoute and that can contribute to RapidRoute missed entries. Traffic destined for an IP address on the router itself will always be ineligible. The **-** character, used for rejected entries, is used when the router rejected the session altogether. This can occur when the firewall blocks a session, drops a packet that violates a rule, or considers packets to be an attack like a SYN flood or SMURF attack.

The **Flag** field outlines a legend for the flags associated with each session. These flags indicate the purpose for which the session was created and where it is being used. The flags used in RapidRoute indicate the following:

- **F firewall** indicates the session has a path through the firewall.
- **N NAT** indicates the session is affected by Network Address Translation (NAT) (could be source or destination NAT or port forward).
- **T altered ToS** indicates the session is altering the type of service (ToS) byte.
- **D don't fragment** indicates a session in which the router is modifying the do not fragment bit (a policy-based routing feature).
- **I IPsec** indicates the path is running through the IPsec process.
- **H hardware assist** indicates the session is flowing through a hardware-assisted datapath.

- **i ingress-filter** indicates the session is using ingress filter rules.
- **e egress-filter** indicates the session is using egress filter rules.

Session Table Information

The remaining information in the command output is the session table that RapidRoute compiles. It includes the ingress interface of the packets, the packet hit, miss, and drop counts, the type and protocol of the flow, the ToS byte information, the source and destination IP addresses, the source and destination port, session age, how many times the session was used, how many packets were dropped from the session, and any session flags. The session table provides information about different flows that are running through the router at a particular point in time and can be used to troubleshoot specific issues or to check how many total RapidRoute sessions there are. The end of the output (per interface) displays how many entries there are, as well as the maximum entries allowed.

The packet hit counts are the number of packets that came and went through a RapidRoute-accelerated path that was already in the RapidRoute table.

The packet miss count displays how many packets have come in on the interface but have not been able to route through a RapidRoute session. There will always be misses displayed because the first packet in a session does not use RapidRoute, and sometimes there are sessions that RapidRoute does not support. Packet miss counts can be used to learn about the traffic on your network because they indicate packets that are taking the slow path through the router for some reason. Misses that exceed the number of hits on an interface can indicate a problem.

Packet misses can occur when the firewall is disabled due to the following:

- IP header options are present
- IP fragments are present
- Time to live (TTL) is less than two (indicating the packet cannot be forwarded but might be locally destined for the router)

Packet misses can occur when the firewall is enabled due to the following for TCP packets:

- IP payload length does not have enough room to hold a complete TCP header (20 bytes)
- IP payload length does not have enough room for the entire TCP payload
- No TCP flags are set
- TCP SYN, RST, URG, or FIN flags are set (commonly happens when sessions are being created or ending)
- TCP sequence number is zero

Packet misses can occur when the firewall is enabled due to the following for UDP packets:

- IP payload length does not have enough room to hold a complete UDP header (8 bytes)
- IP payload length does not have enough room for the entire UDP payload

The packet drop count displays packets that hit the interface, but RapidRoute knew that the firewall or some other process would drop those packets eventually, so RapidRoute dropped them immediately.

Summarized Output

If you have large amounts of traffic flowing through the AOS device, you will see a large amount of output when you issue the **show [ip | ipv6] ffe** command. To display a summarized version of FFE entries on every interface that has RapidRoute enabled, enter the **show [ip | ipv6] ffe summary** command. This command displays the maximum allowed entries on the interface, the current number of FFE entries, the number of packets that have been accelerated by matching a valid FFE entry, the number of packets that have not been accelerated by FFE, and the number of packets that have been dropped by FFE.

Enter the command from the Enable mode as follows:

```
>enable
#show ipv6 ffe summary
Exceptions: 0/128/0 (current/max/drops)

Ingress                 MaxEntries  Entries  Hits  Misses  Drops
----------              ----------  -------  ----  ------  -------
eth 0/1                 4096        1        1000  200     11
eth 0/2.1               4096        1        1123  211     0
eth 0/2.2               4096        1        1467  301     0
                        ----------  -------  ----  ------  -------
                Global  16384       3        3590  712     11
```

RapidRoute Services

In AOS firmware release R11.10.0, two new RapidRoute service features were introduced: flow bundling and service assurance. These services and features are available in addition to general RapidRoute functionality, and can be configured for either IPv4 or IPv6. An overview of these services is included in the following sections.

RapidRoute Flow Bundling

RapidRoute flow bundling is the ability of the AOS device to provide wildcards for certain fields within an RapidRoute flow, and then bundle those flows together, based on their wildcarded fields, into a single flow to be tracked by RapidRoute. This limits the total number of flows that need to be managed without affecting RapidRoute behavior or flows. A wildcarded field in a RapidRoute flow is the field with criteria that is not considered to be part of the flow's unique index. Flow fields considered for wildcarding include the following 11 criteria:

- Source Address
- IP Precedence (first three bits of ToS bytes)
- IP DSCP (first six bits of ToS bytes)
- Layer 4 IP Protocol
- TCP Source Port
- TCP Destination Port
- UDP Source Port
- UDP Destination Port

- ICMP Type, Code, and ID
- ESP SPI
- GRE Tunnel Key

If another subsystem (such as firewall, ACLs, quality of service (QoS), virtual routing and forwarding (VRF), etc.) requires the use of one of those fields for traffic matching, wildcards are disabled for the specific interface so that matching can occur. The wildcard criteria are independent for IPv4 and IPv6, and are determined on a per-interface basis. The ability for any one of the criteria to be wildcarded depends on the configuration of the unit.

Each interface queries the subsystems attached to it to begin generating a list of potential wildcards. All subsystems affecting ingress matching criteria are queried, and for each piece of match criteria used, the corresponding wildcard is disabled. For example, if an interface is configured with an ACL that matches on ICMP type, then ICMP type cannot be wildcarded even if the other criteria are. The more features that are enabled on the interface, the smaller the list of fields that can be wildcarded. For example, with only basic routing enabled, all wildcards are enabled for the interface.

Each interface also considers all potential egress interfaces when calculating its wildcard possibilities. If any potential egress interface (i.e. on the same VRF) has configured a feature that disables certain wildcards for egress traffic, those wildcards are disabled on all potential ingress interfaces.

Flow Bundling Operation in AOS

RapidRoute flow bundling is enabled by default, and does not require a command to enable it. It will wildcard as many items from the wildcard criteria list as possible on a per-interface basis. Each interface has its own wildcard set, as opposed to global or per-flow wildcard set configurations. Destination IP addresses are never wildcarded, and because many subsystems use the same criteria to match or mark traffic as the flow bundling system, many subsystem configurations can affect what items are wildcarded on a particular interface.

Flow bundling monitors the following subsystems to determine wildcard availability on the interface:

- Simple Forwarding
- Firewall
- QoS
- Access Groups
- Route Maps
- Tunnels
- IPsec
- Egress Interfaces
- Integrated Traffic Monitoring (ITM)
- Packet Capture
- Media Anchoring

The interaction between flow bundling and these subsystems is described in *Table 1 on page 12*. Viewing and managing wildcards is described in *Viewing and Managing RapidRoute Flow Bundling on page 13*.

In addition, Rapid Route no longer takes the explicit congestion notification (ECN) bits (the two least significant bits of the ToS/traffic class byte) into account when creating new RapidRoute flows.

RapidRoute Service Assurance

The second RapidRoute feature introduced in AOS firmware release R11.10.0 is service assurance. This feature provides better visibility into RapidRoute configurations and limitations on the AOS device. When RapidRoute flow limits are surpassed, performance issues can arise that can be difficult to diagnose. Service assurance provides CLI commands that display the current peak count and peak history of RapidRoute flows to aid in problem diagnostics. In addition, Simple Network Management Protocol (SNMP) alarms can be used to monitor the peak RapidRoute flow counts using SNMP traps based on network monitoring configurations.

Viewing, managing, and configuring RapidRoute service assurance are described in *Using RapidRoute Service Assurance on page 29*.

RapidRoute Flow Bundling Behaviors

When flow bundling is used in conjunction with certain subsystems, some wildcard functions are disabled. *Table 1* outlines these interactions.

Table 1. Flow Bundling Behavior and Its Interaction with Other Features

Feature	Flow Bundling Behavior
Firewall	When the firewall is enabled for forwarded traffic on a particular VRF, flow bundling is disabled for all interfaces associated with that VRF. This includes the default VRF. When the firewall is enabled for local traffic only, there is no effect on wildcards within the VRF.
QoS	When an inbound QoS map is applied to an interface, flow bundling does not wildcard any fields on the interface matched by any match statements in the map. When an outbound QoS map is applied to an interface, flow bundling does not wildcard any fields matched by any match statements for outbound traffic in the map on any potential egress interfaces within the same VRF.
Access Groups	When an inbound access group is enabled on an interface, flow bundling will not wildcard any fields matched by the ACL on the interface. When an outbound access group is enabled on an interface, flow bundling will not wildcard any fields matched by the ACL on any potential ingress interfaces within the same VRF. If the **fragments** keyword is used on an ACL entry, wildcards are disabled for the specified IP protocol.
Route Maps	When route maps are configured and applied to an interface, they disable all wildcards on the interface to which they are applied.

Table 1. Flow Bundling Behavior and Its Interaction with Other Features *(Continued)*

Feature	Flow Bundling Behavior
Tunnels	Each VRF that is being used as any tunnel interface's source VRF cannot wildcard the source address or the Layer 4 protocol type for all interfaces within the VRF, or wildcard the GRE tunnel key for any interfaces within the VRF if the tunnel key is configured on that tunnel.
IPsec	When IPsec is enabled on a particular VRF, flow bundling is disabled for all interfaces associated with that VRF. When IPsec is enabled globally (**ip crypto** command), the AOS firewall is enabled and disables all wildcards on that VRF.
ITM	ITM disables all wildcards for ingress traffic on the interface on which it is enabled. When RapidRoute flows are reported by ITM, the ECN bits will always be **0** in the reported ToS byte.
Packet Capture	When a packet capture is enabled on an interface, flow bundling is disabled for all interfaces associated with that VRF. This includes the default VRF.
Media Anchoring	When media anchoring is enabled, and a media gateway command is configured on an interface, flow bundling is disabled on that interface.

Viewing and Managing RapidRoute Flow Bundling

RapidRoute flow bundling is enabled by default, so there is no configuration necessary for it to be operational. However, as of AOS firmware release R11.10.1, you can choose to disable flow bundling and all associated wildcards for all interfaces in an IP address family (IPv4 or IPv6) if needed using the **no [ip | ipv6] ffe wildcard** command from the Global Configuration mode prompt. This command disables all wildcards for all interfaces in either the IPv4 address family (**ip** keyword) or the IPv6 address family (**ipv6** keyword). Entering the command without the **no** keyword re-enables wildcards for all interfaces in the address family. To disable flow bundling and wildcards for all interfaces in the IPv4 address family, enter the command as follows:

```
(config)#no ip ffe wildcard
(config)#
```

You can view the wildcards used by flow bundling using several **show** commands. These commands are used to view which fields are wildcarded in a RapidRoute flow and which wildcards are used for each IP interface. Their output can be useful for understanding how flow bundling functions with other features. Configurations for other features, and the resulting flow bundling operation (displayed in **show** command output), can be viewed in *RapidRoute Flow Bundling Behavior with Other Features on page 16*.

Use the **show [ip | ipv6] ffe** command to display the fields that have been wildcarded in a RapidRoute flow. Wildcarded fields appear with the value **any**. If the Layer 4 Internet protocol type is being wildcarded, or it is not TCP or UDP, the port numbers are not displayed. Use the **ip** or **ipv6** keywords to specify whether IPv4 or IPv6 information is being displayed. Enter the command from the Enable mode prompt as follows to display IPv6 RapidRoute information:


```
>enable
#show ipv6 ffe
 Timeout          TCP        UDP       ICMP         AH        ESP        GRE      Other
 Age:          30m 0s     30m 0s     30m 0s     30m 0s     30m 0s     30m 0s     30m 0s
 Inactive:        15s        15s        30s        15s        15s        15s        15s

 Exceptions: 0/217/0 (current/max/drops)

 Type:  * valid, ! ineligible, - rejected
 Flags: F firewall, T altered TC
        H hardware assist, i ingress-filter, e egress-filter

 --------------------------------
 Ingress: system-management-evc
     0 hits, 101489 misses, 0 drops

 T Proto  TC         Age        Used       Drops Flags
   Source
   Destination
   ---------------
 ! udp    dscp:0     16s           0           0
   [2620:106:A00F:DE1::3]:123
   [2620:106:A00F:2A9D::2]:123
 Number of entries: 1 of 1  (4096 maximum)

 --------------------------------
 Total number of entries: 1 of 1  (32768 maximum)
```

Use the **show [ip | ipv6] ffe wildcard [interface** <*interface*>**]** command to display the wildcards being used for each IP interface. Use the **ip** or **ipv6** keywords to specify whether IPv4 or IPv6 information is being displayed. The optional **interface** <*interface*> parameter limits the output to a single interface. Interfaces are specified in the format <*interface type [slot/port | slot/port.subinterface id | interface id | interface id.subinterface id | group id]*>. For example, for an Ethernet interface, use **eth 0/1**; for an Ethernet subinterface, use **eth 0/1.1**; for a PPP interface, use **ppp 1**. Type **show ip ffe wildcard interface ?** to display a complete list of valid interfaces. In the output of this command, any **No** indicates the field cannot be wildcarded, and each variation will create a new RapidRoute flow. Any **Yes** indicates the field is wildcarded and RapidRoute flows will be bundled based on that field.

Enter the command from the Enable mode prompt as follows to view all IPv4 IP interfaces and their wildcards:

```
>enable
#show ip ffe wildcard
Field                                Wildcarded
=============================        ============================
eth 0/1
   Source IP Address                 : No
   Dest IP Address                   : No (always)
   IP Precedence                     : No
   IP DSCP                           : Yes
   IP Protocol (L4)                  : Yes
   TCP Source Port                   : Yes
   TCP Destination Port              : Yes
   UDP Source Port                   : Yes
   UDP Destination Port              : Yes
```



```
   ICMP Type, Code and ID              : Yes
   ESP SPI                             : Yes
   GRE Tunnel Key                      : Yes

eth 0/2
   Source IP Address                   : Yes
   Dest IP Address                     : No (always)
   IP Precedence                       : Yes
   IP DSCP                             : Yes
   IP Protocol (L4)                    : Yes
   TCP Source Port                     : Yes
   TCP Destination Port                : Yes
   UDP Source Port                     : Yes
   UDP Destination Port                : Yes
   ICMP Type, Code and ID              : Yes
   ESP SPI                             : Yes
   GRE Tunnel Key                      : Yes
```

RapidRoute flow bundling debug messaging (wildcard debugging) can be enabled or disabled using the **debug [ip | ipv6] ffe wildcard** command. This command can be useful when determining which of several subsystem configurations caused a specific RapidRoute wildcard to be disabled. Specify **ip** or **ipv6** to specify whether IPv4 or IPv6 debug messaging is enabled. When RapidRoute wildcard debug messages are enabled, two wildcard events are displayed. A **calculate** event, generated when an interface is called to recalculate its inbound or outbound wildcards, displays the results for each subsystem in the order in which wildcard processing is completed. A **finalize set** event is generated when the complete wildcards for an interface are pushed to either hardware or software processing and the new wildcards are used.

The wildcard bits in a wildcard debug message are displayed in the opposite order they are displayed in the **show [ip | ipv6] ffe wildcard** command, and exclude the destination IP address (which cannot be wildcarded). The last bit displayed at the end of the wildcard string is the least significant bit and represents the source IP address.

The following example enables RapidRoute wildcard debug messaging and provides sample event output. Also included in the example are the configuration of the IPv4 ACL named **ACL** and its application to the gigabit Ethernet subinterface **0/5.1**:

```
>enable
#debug ip ffe wildcard
#config t
(config)#ip access-list extended ACL
(config-ext-nacl)#permit icmp any any echo
(config-ext-nacl)#exit
(config)#interface gigabit ethernet 0/5.1
(config-giga-eth 0/5.1)#ip access-group ACL out
2015.01.01 12:00:34 FFEWILDCARDV4.giga-eth 0/5.1 calculate outbound:
    QoS:                    11111111111
    AccessGroup:            11011110111
2015.01.01 12:00:34 FFEWILDCARDV4.Loopback finalize set: 11011110111
2015.01.01 12:00:34 FFEWILDCARDV4.null 0 finalize set: 11011110111
```


2015.01.01 12:00:34 FFEWILDCARDV4.giga-eth 0/2.1 finalize set: 11011110111
2015.01.01 12:00:34 FFEWILDCARDV4.giga-eth 0/5.1 finalize set: 11011110111

> NOTE *Turning on a large amount of debug information can adversely affect the performance of your unit.You can view wildcard status on a per-interface basis using the **show [ip | ipv6] ffe** command.*

RapidRoute Flow Bundling Behavior with Other Features

The following configuration examples include the configuration of a feature that interacts with RapidRoute flow bundling and the wildcarding process. Each example gives a brief scenario, the AOS device configuration, and then the resulting flow bundling behavior (what fields have wildcards applied) as displayed by using the **show [ip | ipv6] ffe wildcard** command. Use these examples to become familiar with how different subsystems and matching criteria affect wildcard fields and their availability.

Simple Route Forwarding

In simple route forwarding, all wildcards are enabled. When simple route forwarding is configured as in the example below, the **show ip ffe wildcard** command will demonstrate that all wildcards are enabled.

```
no ip firewall
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When simple route forwarding is configured as in the previous example, the **show ip ffe wildcard** command shows that all wildcards are enabled.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
============================       ============================
eth 0/1
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : Yes
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
```



```
   ESP SPI                              : Yes
   GRE Tunnel Key                       : Yes

eth 0/2
   Source IP Address                    : Yes
   Dest IP Address                      : No (always)
   IP Precedence                        : Yes
   IP DSCP                              : Yes
   IP Protocol (L4)                     : Yes
   TCP Source Port                      : Yes
   TCP Destination Port                 : Yes
   UDP Source Port                      : Yes
   UDP Destination Port                 : Yes
   ICMP Type, Code and ID               : Yes
   ESP SPI                              : Yes
   GRE Tunnel Key                       : Yes
```

Firewall and IPsec

When the firewall is enabled (using the **ip firewall** or **ipv6 firewall** commands), all wildcards are disabled on all interfaces for that VRF. Enabling IPsec (using the **ip crypto** command) also disables all wildcards on all interfaces. When the firewall is enabled and configured as in the example below, the **show ip ffe wildcard** command shows that all wildcards are disabled.

```
ip firewall
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When the firewall is configured as in the previous example, the **show ip ffe wildcard** command shows that all wildcards are disabled.

```
>enable
#show ip ffe wildcard
Field                               Wildcarded
============================        ============================
eth 0/1
   Source IP Address                : No
   Dest IP Address                  : No (always)
   IP Precedence                    : No
   IP DSCP                          : No
   IP Protocol (L4)                 : No
   TCP Source Port                  : No
   TCP Destination Port             : No
```



```
   UDP Source Port                    : No
   UDP Destination Port               : No
   ICMP Type, Code and ID             : No
   ESP SPI                            : No
   GRE Tunnel Key                     : No

eth 0/2
   Source IP Address                  : No
   Dest IP Address                    : No (always)
   IP Precedence                      : No
   IP DSCP                            : No
   IP Protocol (L4)                   : No
   TCP Source Port                    : No
   TCP Destination Port               : No
   UDP Source Port                    : No
   UDP Destination Port               : No
   ICMP Type, Code and ID             : No
   ESP SPI                            : No
   GRE Tunnel Key                     : No
```

QoS Maps Example 1

Using QoS disables all wildcards for any flow bundling criteria also used to match traffic in the QoS map. In the following example, the QoS map **STATS** is configured to match on IP precedence, and the QoS map **DSCP** is configured to match on IP DSCP. When the AOS device is configured with the two QoS maps as follows, the **show ip ffe wildcard** command will show the wildcards used.

```
no ip firewall
!
qos map STATS 11
    match precedence 1
qos map STATS 12
    match precedence 2
qos map STATS 13
    match precedence 3
qos map STATS 14
    match any
!
qos map DSCP 11
    match dscp af11
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    qos-policy in STATS
    no shutdown
!
```



```
interface ethernet 0/2
    ip address 2.2.2.2 /24
    qos-policy in DSCP
    no shutdown
```

When QoS is configured as in the previous example, the **show ip ffe wildcard** command shows that the **IP Precedence** field in both interfaces, and the **IP DSCP** field in the **eth 0/2** interface, are not wildcarded due to the QoS maps applied to the interfaces.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      =============================
eth 0/1
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : Yes
   IP Protocol (L4)                : Yes
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes

eth 0/2
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : No
   IP Protocol (L4)                : Yes
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes
```

QoS Maps Example 2

As stated previously, QoS will disable wildcards for any fields used for matching in a QoS map. In the following example, the QoS map **VOICE** is configured to match on the Layer 4 IP protocol. When the AOS device is configured with the QoS map as follows, the **show ip ffe wildcard** will show which wildcards are used.


```
no ip firewall
!
qos map VOICE 10
    match ip rtp 5000 6000
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    qos-policy in VOICE
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When QoS is configured as in the previous example, the **show ip ffe wildcard** command shows that the **IP Protocol (L4)** and the **UDP Destination Port** fields in the **eth 0/1** interface are not wildcarded, due to the QoS map applied to the interface.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
==============================     ==============================
eth 0/1
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : No
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes

eth 0/2
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : Yes
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes
```


Access Groups

When configured and applied, ACLs and access groups disable wildcards on any **permit** or **deny** criteria. If the **fragments** keyword is used by any line in the ACL, wildcards are disabled for the specified IP protocol.

In the following example, the ACL **LIST** is configured to match on TCP, UDP, and ICMP. The ACL is then applied to the **eth 0/1** interface. When the AOS device is configured with the ACL as follows, the **show ip ffe wildcard** command will display which wildcards are used.

```
ip access-list extended LIST
    permit tcp 10.10.10.0 0.0.0.255 eq ssh any eq ssh
    permit udp 10.10.10.0 0.0.0.255 eq snmp any eq snmp
    permit icmp any any echo
    deny ip any any
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    ip access-group LIST in
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When ACLs and access groups are configured as in the previous example, the **show ip ffe wildcard** command shows that the **IP Protocol (L4)**, **TCP Source** and **Destination Port**, **UDP Source** and **Destination Port**, and **ICMP** fields in the **eth 0/1** interface are not wildcarded, due to the ACL applied to the interface.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      =============================
eth 0/1
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : No
   TCP Destination Port            : No
   UDP Source Port                 : No
   UDP Destination Port            : No
   ICMP Type, Code and ID          : No
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes
```



```
eth 0/2
   Source IP Address                  : Yes
   Dest IP Address                    : No (always)
   IP Precedence                      : Yes
   IP DSCP                            : Yes
   IP Protocol (L4)                   : Yes
   TCP Source Port                    : Yes
   TCP Destination Port               : Yes
   UDP Source Port                    : Yes
   UDP Destination Port               : Yes
   ICMP Type, Code and ID             : Yes
   ESP SPI                            : Yes
   GRE Tunnel Key                     : Yes
```

Route Maps

When route maps are configured and applied to an interface, they disable all wildcards on the interface to which they are applied. In the following example, route map **RM** is configured to match and act on IP precedence. When the AOS device is configured with the route map as follows, the **show ip ffe wildcard** command displays which wildcards are used.

```
route-map RM permit 100
    match ip precedence 3 2
    set ip precedence 1
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    ip policy route-map RM
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When route maps are configured as in the previous example, the **show ip ffe wildcard** command shows that all wildcards in the **eth 0/1** interface have been disabled, due to the route map applied to the interface.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      ============================
eth 0/1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : No
   IP Protocol (L4)                : No
   TCP Source Port                 : No
   TCP Destination Port            : No
```



```
   UDP Source Port                    : No
   UDP Destination Port               : No
   ICMP Type, Code and ID             : No
   ESP SPI                            : No
   GRE Tunnel Key                     : No

eth 0/2
   Source IP Address                  : Yes
   Dest IP Address                    : No (always)
   IP Precedence                      : Yes
   IP DSCP                            : Yes
   IP Protocol (L4)                   : Yes
   TCP Source Port                    : Yes
   TCP Destination Port               : Yes
   UDP Source Port                    : Yes
   UDP Destination Port               : Yes
   ICMP Type, Code and ID             : Yes
   ESP SPI                            : Yes
   GRE Tunnel Key                     : Yes
```

Tunnel Interfaces and VRFs Example 1

Tunnel interfaces with a tunnel key configured prevent wildcards on the following criteria: source address, Layer 4 IP protocol, and GRE tunnel key. These wildcards are prevented on all interfaces within the tunnel's source VRF, including the tunnel itself (regardless of the tunnel's unencapsulated VRF). If the tunnel type is multipoint GRE (MGRE), the source address can still be wildcarded, but otherwise the same rules apply. When the AOS device is configured with the tunnel interface as follows, the **show ip ffe wildcard** command will display which wildcards are used.

```
interface tunnel 1 gre ip
    ip address 3.3.3.3 /24
    tunnel source eth 0/2
    tunnel key 1000
    no shutdown
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When the tunnel interface is configured as in the previous example, the **show ip ffe wildcard** command shows that the **Source IP Address**, **IP Protocol (L4)** and the **GRE Tunnel Key** fields in all interfaces are not wildcarded, due to the tunnel key configuration on the tunnel interface and the fact that all interfaces are in the same VRF as the tunnel interface.


```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      ============================
eth 0/1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : No

eth 0/2
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : No

tunnel 1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : No
```


Tunnel Interfaces and VRFs Example 2

When a tunnel interface is part of two VRFs, the tunnel's overall wildcard set is the most restrictive intersection of the wildcards calculated for the interface in both VRFs. This occurs because there is only one set of wildcards per interface, but the tunnel interface effectively spans both VRFs. For example, if a tunnel is sourced on VRF B, and the logical side of the interface is sourced in VRF A, the most restrictive wildcard rules will apply to the tunnel interface, no matter which VRF they come from. In the example below, the tunnel interface cannot wildcard IP precedence even in VRF A since the tunnel also participates in VRF B, which is configured with a QoS map (**PREC**) matching on IP precedence.

When the AOS device is configured with the tunnel interface and VRFs A and B as follows, the **show ip ffe wildcard** command will display which wildcards are used.

```
interface tunnel 1 gre ip
    vrf forwarding A
    ip address 3.3.3.3 /24
    tunnel source eth 0/2
    tunnel key 1000
    tunnel vrf B
    no shutdown
!
qos map PREC 11
    match precedence 1
qos map PREC 12
    match any
!
interface ethernet 0/1
    vrf forwarding A
    ip address 1.1.1.1 /24
    no shutdown
!
interface ethernet 0/2
    vrf forwarding B
    ip address 2.2.2.2 /24
    qos-policy outbound PREC
    no shutdown
```

When the AOS device is configured with the tunnel interface and VRFs A and B as shown in the previous example, the **show ip ffe wildcard** command shows that the **Source IP Address**, **IP Precedence**, **IP Protocol (L4)**, and the **GRE Tunnel Key** fields in both the **eth 0/2** and **tunnel 1** interfaces are not wildcarded, due to the tunnel key configuration on the tunnel interface and the fact that the **eth 0/2** interface has QoS map **PREC** applied.


```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      =============================
eth 0/1
   Source IP Address               : Yes
   Dest IP Address                 : No (always)
   IP Precedence                   : Yes
   IP DSCP                         : Yes
   IP Protocol (L4)                : Yes
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : Yes

eth 0/2
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : No

tunnel 1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : Yes
   IP Protocol (L4)                : No
   TCP Source Port                 : Yes
   TCP Destination Port            : Yes
   UDP Source Port                 : Yes
   UDP Destination Port            : Yes
   ICMP Type, Code and ID          : Yes
   ESP SPI                         : Yes
   GRE Tunnel Key                  : No
```


ITM

When ITM is enabled in the ingress direction, all wildcards are disabled for the ingress interface. If ITM is enabled in the egress direction, wildcards for all interfaces within the same VRF are disabled. The following example configures the **eth 0/1** interface with egress flow monitoring enabled.

```
interface ethernet 0/1
    ip address 1.1.1.1 /24
    ip flow egress
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When ITM is configured as in the previous example, the **show ip ffe wildcard** command shows that all wildcards for both interfaces have been disabled, due to the use of egress ITM on one interface within the same VRF.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
==============================     ==============================
eth 0/1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : No
   IP Protocol (L4)                : No
   TCP Source Port                 : No
   TCP Destination Port            : No
   UDP Source Port                 : No
   UDP Destination Port            : No
   ICMP Type, Code and ID          : No
   ESP SPI                         : No
   GRE Tunnel Key                  : No

eth 0/2
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : No
   IP Protocol (L4)                : No
   TCP Source Port                 : No
   TCP Destination Port            : No
   UDP Source Port                 : No
   UDP Destination Port            : No
   ICMP Type, Code and ID          : No
   ESP SPI                         : No
   GRE Tunnel Key                  : No
```


Packet Capture

If packet capture is enabled on an interface, all wildcards are disabled for all interfaces within the same VRF. For example, if the AOS device is configured with packet capture as follows, the **show ip ffe wildcard** command will show that all wildcards for both interfaces have been disabled.

```
ip access-list extended PCAP_ACL
    deny udp any host 10.10.10.1
    deny udp host 10.10.10.1 any
    permit udp any
!
packet-capture PCAP standard
    export tftp 10.10.10.1
    match list PCAP_ACL
    no shutdown
!
interface ethernet 0/1
    ip address 1.1.1.1 /24
    ip packet-capture PCAP
    no shutdown
!
interface ethernet 0/2
    ip address 2.2.2.2 /24
    no shutdown
```

When packet capture is configured as shown in the previous example, the **show ip ffe wildcard** command shows that all wildcards for both interfaces have been disabled, due to the use of packet capture on one interface within the same VRF.

```
>enable
#show ip ffe wildcard
Field                              Wildcarded
=============================      =============================
eth 0/1
   Source IP Address               : No
   Dest IP Address                 : No (always)
   IP Precedence                   : No
   IP DSCP                         : No
   IP Protocol (L4)                : No
   TCP Source Port                 : No
   TCP Destination Port            : No
   UDP Source Port                 : No
   UDP Destination Port            : No
   ICMP Type, Code and ID          : No
   ESP SPI                         : No
   GRE Tunnel Key                  : No
```



```
eth 0/2
   Source IP Address                    : No
   Dest IP Address                      : No (always)
   IP Precedence                        : No
   IP DSCP                              : No
   IP Protocol (L4)                     : No
   TCP Source Port                      : No
   TCP Destination Port                 : No
   UDP Source Port                      : No
   UDP Destination Port                 : No
   ICMP Type, Code and ID               : No
   ESP SPI                              : No
   GRE Tunnel Key                       : No
```

Using RapidRoute Service Assurance

RapidRoute service assurance is a feature that allows better visibility into the functioning of RapidRoute itself. It can be used to diagnose performance issues that may be otherwise difficult to determine by providing CLI commands that display the current peak count and peak history of RapidRoute flows. In addition, SNMP alarms can be used to monitor the peak RapidRoute flow counts through the use of SNMP traps based on network monitoring configurations.

Viewing and Clearing RapidRoute Peak Flow Statistics

Like RapidRoute flow bundling, RapidRoute service assurance is a feature that is enabled by default. Therefore, **show** commands are used to view the peak flow statistics, and **clear** commands are used to clear gathered statistics.

Use the **show [ip | ipv6] ffe peak** command to display the current and peak count of RapidRoute sessions on the AOS device. Maximum entries, current entries, peak entries, and the last time the counters were cleared are displayed for each eligible interface, as well as the global values. Use the **ip** and **ipv6** keywords to specify whether IPv4 or IPv6 data is displayed. The current number of entries (**Entries**) can be used to determine if the current state of the system is similar to the state of the system when it reached the peak entry count. The maximum number of entries (**MaxEntries**) can be useful because if the peak entry count is approaching this value, it indicates that an upgrade to a device with more capacity or the deployment of additional devices for load sharing may be necessary.

Enter the command from the Enable mode prompt as follows to view RapidRoute service assurance statistics:

```
>enable
#show ip ffe peak
```

Ingress	MaxEntries	Entries	Peak	Last Cleared
giga-eth 0/4.1	64000	1000	16000	16 Sep 2015 18:00:00
giga-eth 0/5.3158	64000	8000	8000	Never
Global	64000	9000	20000	Never

Use the **show [ip | ipv6] ffe peak history** command to display a graphical presentation of the peak and average global RapidRoute count per second for the last 60 seconds. This command also displays the peak and average global RapidRoute count per minute for the last 60 minutes, as well as the peak and average global RapidRoute count per hour for the last 72 hours. Use the **ip** and **ipv6** keywords to specify whether IPv4 or IPv6 data is displayed.

> NOTE: *Data is presented as a percentage of the value configured using the **ip ffe max-entries** <value> or **ipv6 ffe max-entries** <value> commands. If these values are changed, the output for the **show [ip | ipv6] ffe peak history** command is cleared.*

Enter the command from the Enable mode prompt as follows to view the historical data for RapidRoute service assurance statistics:

>enable
#show ip ffe peak history

```
#: Average FFE entry % of configured FFE max-entries per interval
@: Maximum FFE entry % of configured FFE max-entries per interval
Most current interval starts on the left.

Previous 1 minute peak:
100
 90
 80                                          @@@@
 70                                      @@@@@@@@
 60                                     @@@@@@@@@
 50                                   @@@@@@@@@@@
 40                               @@@@@@@@@@@@@@@@     @@@@@
 30    @@    @@   @@@@@@@@@@     @@@@@@@@@@@@@@@@@@@@ @@@@@@@@@@
 20  @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
 10  @@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@
  0 +++++++++1+++++++++2+++++++++3+++++++++4+++++++++5+++++++++6
             0         0         0         0         0         0
                              seconds
Previous 1 hour peak:
100
 90
 80
 70
 60
 50               @        @@@                            @
 40          @@@@@@@@@@@@@@@@@@@@        @      @@@      @ @ @@
 30 @@@@@@@@@##########@#######@@ @@@@@@@@ @@@@@@@@ @  @@@ @ @@
 20 ############################################### @@@#########
 10 ############################################################
  0 +++++++++1+++++++++2+++++++++3+++++++++4+++++++++5+++++++++6
             0         0         0         0         0         0
                              minutes

Previous 72 hours peak:
100
 90
 80
 70
 60
 50                                                                 @@@
 40          @@             @@@                                   @@@##@
 30 @@@@@@@@@@@@@@@@@@@@@@@@@@@@ @@@@@@@@@@@@@@@@@@@         @@@@@@####@   @
 20 ########@@#######@@@######@@@#########@@#######@@@@@@@@@@@@########@@@@@
 10 ########################################################################
  0 +++++++++1+++++++++2+++++++++3+++++++++4+++++++++5+++++++++6+++++++++7++
             0         0         0         0         0         0         0
                                     hours
```


Use the **clear [ip | ipv6] ffe [**<*interface*>**] peak** command to clear the RapidRoute peak entry count. Use the **ip** and **ipv6** keywords to specify whether IPv4 or IPv6 statistics are cleared. The optional <*interface*> parameter specifies that statistics are cleared only on the specified interface. Interfaces are specified in the format <*interface type [slot/port | slot/port.subinterface id | interface id | interface id.subinterface id | group id]*>. For example, for an Ethernet interface, use **eth 0/1**; for an Ethernet subinterface, use **eth 0/1.1**; for a PPP interface, use **ppp 1**. Type **clear [ip | ipv6] ffe ?** to display a complete list of valid interfaces. If an interface is not specified, the global peak entry count and all interface peaks are cleared.

To clear the RapidRoute peak entry count for all IPv4 interfaces, and the global IPv4 peak entry count, enter the command from the Enable mode prompt as follows:

```
>enable
#clear ip ffe peak
```

RapidRoute Service Assurance and Network Monitoring

RapidRoute service assurance can be used in conjunction with network monitoring. SNMP traps can be configured to signal alarms based on track state changes, which can be configured to change state based on RapidRoute peak count limits. This functionality is another method to diagnose performance issues by viewing RapidRoute statistics. To configure network monitoring to monitor RapidRoute peak statistics, follow these steps:

1. Configure the network monitor track to test that the number of RapidRoute flow entries is less than the specified number. You can configure a track test (using the command **test** from the Network Monitoring Track Configuration mode), a track test list (using the commands **test list and** or **test list or** from the Network Monitoring Track Configuration mode), or a weighted track test (using the **test list weighted** command from the Network Monitoring Track Configuration mode) for this feature. Once the test type is configured, use the **if [not] [ip | ipv6] ffe [**<*interface*>**] entries less-than** <*number*> command to specify that the network monitor track tests that the specified number of RapidRoute flow entries is not exceeded. Use the optional **not** parameter to indicate that the individual track state will negate the result of the object test (the RapidRoute flow entry value). Use the **ip** and **ipv6** keywords to specify whether the track is monitoring IPv4 or IPv6 RapidRoute. The optional <*interface*> parameter specifies an RapidRoute ingress interface to test. Interfaces are specified in the format <*interface type [slot/port | slot/port.subinterface id | interface id | interface id.subinterface id | group id]*>. For example, for an Ethernet interface, use **eth 0/1**; for an Ethernet subinterface, use **eth 0/1.1**; for a PPP interface, use **ppp 1**. If no interface is specified, the global RapidRoute entry count is tested. The <*number*> parameter specifies the maximum number of RapidRoute flow entries for the test. Valid range is **1** to **500000** entries. Use the **no** version of this command to remove the test. The following example configures a track to monitor the RapidRoute flow entries on an Ethernet interface and will cause the track to change state if the maximum number of flows (**35000**) is exceeded:

   ```
   (config)#track TRACK1
   (config-track)#test if ip ffe ethernet 0/1 entries less-than 35000
   ```

2. Next, you can configure the track to send an SNMP trap when the track changes state. If RapidRoute service assurance is being used on the track, the track will change state if the maximum number of RapidRoute flows is exceeded, and an SNMP trap will be generated. To achieve this, enter the **snmp trap state-change** command from the Network Monitoring Track Configuration mode. The following example configures **TRACK1** (which is configured to test RapidRoute flows) to send an SNMP trap upon track state change:


```
(config)#track TRACK1
(config-track)#snmp trap state-change
```

3. You can use the **show track** command from the Enable mode to display what tracks are being used and their respective RapidRoute maximum flow entry values. Enter the command as follows:

```
>enable
#show track
Track TRACK1:
    Current State: PASS          (Admin: UP)
    Testing:
        ip ffe eth 0/1 entries less-than 35000          (PASS)
    Dampening Interval: 1 seconds
    Time in current state: 0 days, 0 hours, 0 minutes, 30 seconds
    Track State Changes: 1
    Tracking:
```

RapidRoute Command Summary

The following tables summarize the commands used in conjunction with RapidRoute Global and interface configuration, as well as RapidRoute flow bundling and RapidRoute service assurance.

Table 2. RapidRoute Global Configuration Commands

Prompt	Command	Description
(config)#	**[no] [ip | ipv6] ffe max-entries** <*value*>	Specifies the global maximum of FFE entries allowed at a given time. Valid range is **1** to **500000** entries, with a default value of **16384** entries.
(config)#	**[no] [ip | ipv6] ffe timeout [ah | esp | gre | icmp | other | tcp | udp]** <*max timeout*> **[**<*inactive timeout*>**]**	Configures the timeout values for FFE entries in the traffic flow table. You can optionally specify the protocol using the protocol keywords. The <*max timeout*> value is the maximum time (in seconds) that an entry stays in the table before being removed. Valid range is **60** to **86400** seconds, with a default value of **1800** seconds. The optional <*inactive timeout*> parameter specifies the maximum time (in seconds) that an FFE entry stays in the table if no traffic matches it. Valid range is **10** to **86400** seconds, with a default value of **15** seconds.
(config)#	**[no] [ip | ipv6] ffe limit exceptions** <*number*>	Limits the number of unhandled FFE exception packets allowed at any given time. Valid range is **1** to **1024**, with a default value of **128** packets.
(config)#	**[no] ip crypto ffe [max-entries** <*value*>**]**	Enables RapidRoute for IPv4 IPsec security associations. The optional **max-entries** parameter specifies the maximum number of entries per inbound IPsec security association. Valid range is **1** to **8192** entries, with a default value of **4096** entries.

Table 3. RapidRoute Interface Configuration Commands

Prompt	Command	Description
(config-int)#	**[no] [ip | ipv6] ffe [max-entries** <*value*>**]**	Enables or disables RapidRoute on the interface. The optional **max-entries** parameter specifies the maximum number of FFE entries supported by the interface. Valid range is **1** to **500000** entries, with a default value of **4096** entries. RapidRoute is enabled by default on all IP interfaces.

Table 4. RapidRoute Flow Bundling Commands

Prompt	Command	Description
(config)#	**[no] [ip | ipv6] ffe wildcard**	Using the **no** version of this command disables the RapidRoute flow bundling feature for all interfaces in the IPv4 (**ip** keyword) or IPv6 (**ipv6** keyword) address family. To re-enable flow bundling, enter the command without the **no** keyword. Flow bundling is enabled by default.
#	**show [ip | ipv6] ffe**	Displays IPv4 or IPv6 RapidRoute flows and statistics. Wildcarded fields appear with the value **any**.
#	**show [ip | ipv6] ffe wildcard [interface** <*interface*>**]**	Displays the wildcards being used for each IP interface. The optional **interface** <*interface*> parameter limits output to a single interface.
#	**[no] debug [ip | ipv6] ffe wildcard**	Enables debug messaging for IPv4 or IPv6 RapidRoute flow bundling. Use the **no** form of this command to disable debug messaging for RapidRoute flow bundling.

Table 5. RapidRoute Service Assurance Commands

Prompt	Command	Description
#	**show [ip | ipv6] ffe peak**	Displays the current and peak count of RapidRoute sessions on the AOS device.
#	**show [ip | ipv6] ffe peak history**	Displays a graphical presentation of the peak and average global RapidRoute count per second for the last 60 seconds, the peak and average global RapidRoute count per minute for the last 60 minutes, and the peak and average global RapidRoute count per hour for the last 72 hours.
#	**clear [ip | ipv6] ffe [**<*interface*>**] peak**	Clears the RapidRoute peak entry count. The optional <*interface*> parameter specifies that statistics are only cleared on that interface. If no interface is specified, all RapidRoute peak entry counts are cleared.

Table 5. RapidRoute Service Assurance Commands *(Continued)*

Prompt	Command	Description
(config-track)# -OR- (config-track-test)#	**[no] test if [not] [ip | ipv6] ffe [**<*interface*>**] entries less-than** <*number*> -OR- **[no] if [not] [ip | ipv6] ffe [**<*interface*>**] entries less-than** <*number*>	Specifies that network monitor track test that the specified number of RapidRoute flow entries is not exceeded. The optional **not** parameter indicates that the individual track state will negate the result of the object test. The optional <*interface*> parameter specifies a RapidRoute ingress interface to test. Use the <*number*> parameter to specify the maximum number of RapidRoute flow entries for the test. Valid range is **1** to **500000**. The **no** form of this command removes the test from the track's configuration.
(config-track)#	**[no] snmp trap state-change**	Configures the network monitor track to send an SNMP trap when the track changes state.The **no** form of this command disables the SNMP trap on the track.
#	**show track**	Displays any configured tracks and their associated RapidRoute flow maximum entry values.

Additional Resources

The following table lists several features, and their associated documentation, that may interact with RapidRoute itself, RapidRoute flow bundling, or RapidRoute service assurance. All documents are available online at https://supportforums.adtran.com.

Table 6. Documentation of Features that may Interact with RapidRoute

Feature	Document
IPv4 Firewall	*Configuring the Firewall (IPv4) in AOS*
IPv6 Firewall	*IPv6 Firewall Protection in AOS*
IPv6	*Configuring IPv6 in AOS*
ITM	*Configuring ITM in AOS*
SNMP	*Configuring SNMP in AOS*
Network Monitoring	*Configuring Network Monitor in AOS*
Media Anchoring	*Configuring Media Anchoring in AOS*
Quality of Service (QoS)	*Configuring QoS in AOS*
Carrier Ethernet QoS	*Carrier Ethernet Services QoS*