APPLICATION NOTE

FILTER-BASED FORWARDING

Using Filter-Based Forwarding to Control Next-Hop Selection

Table of Contents

Table of Figures

Introduction

Filter-based forwarding enables you to configure packet filters that classify packets based on header information, such as IP source address, IP destination address, IP protocol field, and source and destination TCP/UDP port numbers. If a packet matches the conditions of the filter, then traditional destination-based forwarding occurs using the routing table that is specified in the accept action of the filter definition language. Filter-based forwarding provides a very simple yet powerful tool: a policy-based routing table selector.

Scope

This document describes the purpose and mechanics of filter-based forwarding and then discusses some key applications. The examples mainly focus around IPv4, but are applicable to IPv6 as well.

The features described in this application note are supported on Juniper Networks® J Series Services Routers, M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers.

Description and Deployment Scenario

Filter-based forwarding allows you to control the next-hop selection for customer traffic by defining input packet filters that examine the fields in a packet's header. If a packet satisfies the match conditions of the filter, the packet is forwarded using the routing table instance specified in the filter action statement. This application note describes the use of FBF deployment in the following scenarios:

- Open access
- BGP/MPLS VPNs (RFC 2547bis)
- Traffic engineering without MPLS

Packet Classification

The packet filter can classify packets based on any of the fields that can be examined by the Juniper Networks Junos® operating system filter definition language. These fields include the following:

- Source and/or destination IP addresses
- Protocol number
- Source and/or destination port numbers
- IP precedence value
- DSCP value
- IP options
- TCP flags
- Packet length
- ICMP type
- Incoming and/or outgoing logical or physical interface

Filter Actions

If a packet satisfies the conditions of the filter, you can specify the filter action known as a *<routing-instance>*. This filter action allows you to specify the routing table instance that is used to forward traffic that matches the filter's conditions. Once the routing table is identified, traditional destination-based routing occurs. In addition to the routing-instance action, you can also specify the following action modifiers in the filter:

- Alert
- Count
- Log
- Output-queue
- PLP
- Police
- Sample

Filter-Based Forwarding Example

Figure 1 illustrates a sample packet flow where an input packet filter is used to classify packets, and each packet is forwarded to a different next hop using different routing tables based on the result of the packet classification process.

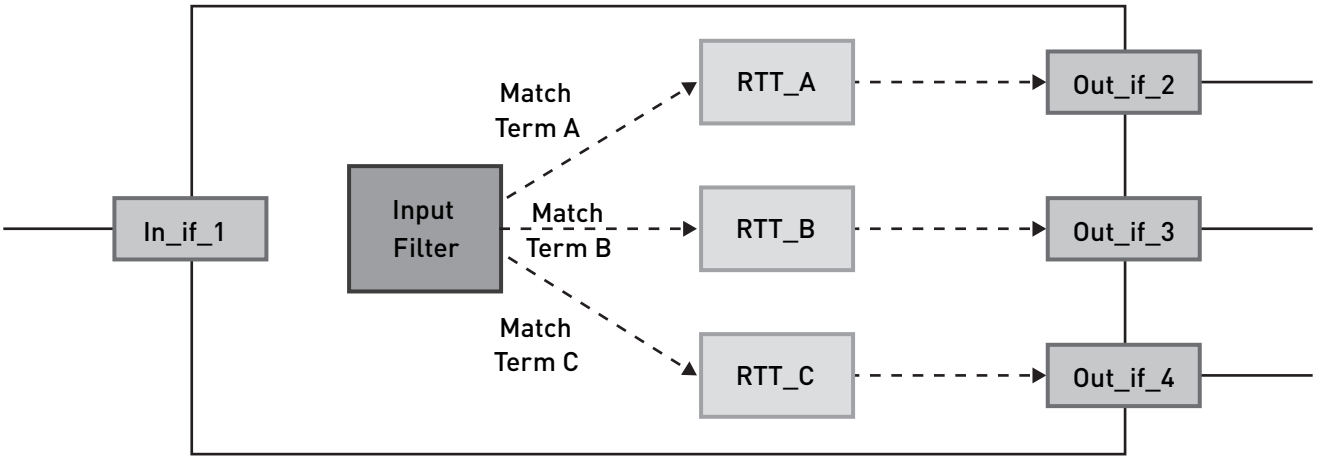

Figure 1: Sample packet flow

In this example, incoming packets arrive on interface In_if_1. The Packet Forwarding Engine (PFE), using an input packet filter, examines each packet. If the packet matches Term 1 of the filter, then destination-based forwarding occurs using RTT_A. If the packet matches Term 2 of the filter, then destination-based forwarding occurs using RTT_B. If the packet matches Term 3 of the filter, then destination-based forwarding occurs using RTT_C.

Sample Syntax

This section provides sample syntax that illustrates how you can configure filter-based forwarding on a Juniper Networks router. The first configuration fragment defines a packet filter that directs customer traffic to a next-hop router in SP 1 or SP 2 based on the packet's source address. Figure 2 shows the network topology for this example.

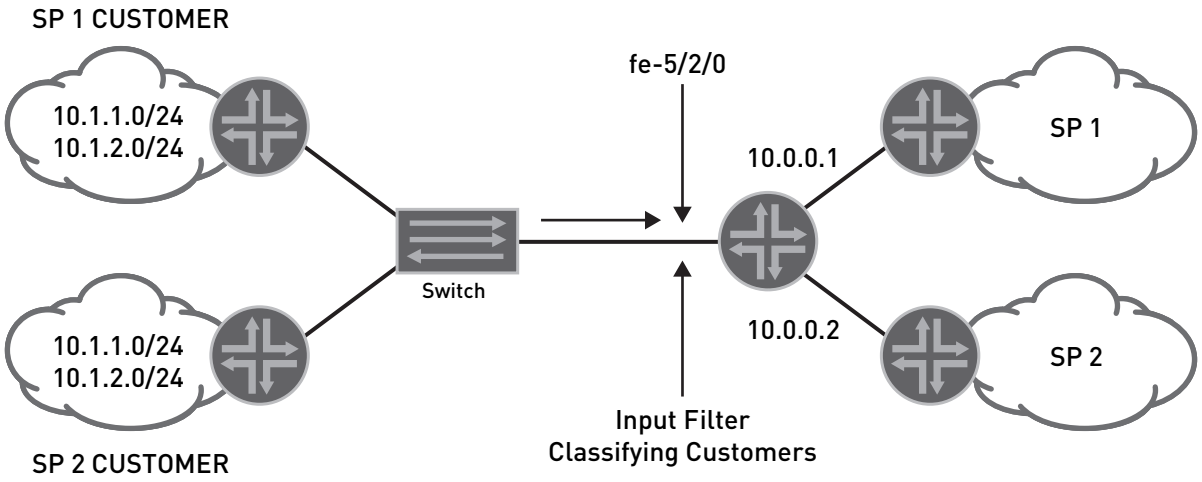

Figure 2: Network topology for sample syntax example

If the packet has a source address assigned to an SP 1 customer, then destination-based forwarding occurs using the sp1-route-table. If the packet has a source address assigned to an SP 2 customer, then destination-based forwarding occurs using the sp2-route-table. If a packet does not match either of these conditions, then the filter accepts the packet, and then destination-based forwarding occurs using the standard inet.0 routing table.

```
filter classify-customers {
      term sp1-customers {
            from {
                  source-address { /* SP 1 customer prefixes */
                        10.1.1.0/24;
                        10.1.2.0/24;
                  }
            }
            then {
                  routing-instance sp1-route-table;
            }
       }
       term sp2-customers {
           from {
                source-address { /* SP 2 customer prefixes */
                      10.2.1.0/24;
                      10.2.2.0/24;
                }
           }
           then {
                 routing-instance sp2-route-table;
           }
      }
      term default { /* Accept all other traffic */
           then {
                accept; /* Forward using inet.0 */
           }
      }
}
```

The following configuration segment defines the routing-instances referenced in the filter classify-customers:

```
routing-instance {
      sp1-route-table {
            instance-type forwarding;
            static
                  route 0.0.0.0/0 nexthop 10.0.0.1; /* Static default route */
            }
       }
       sp2-route-table {
             instance-type forwarding;
             static {
                 route 0.0.0.0/0 nexthop 10.0.0.2; /* Static default route */
             }
       }
```

The following configuration fragment resolves the interface routes installed in the routing instances to directly connected next hops on that interface:

```
routing-options {
      interface-routes {
            rib-group inet fbf-group;
      }
      rib-groups {
            fbf-group {
```

```
                    import-rib [inet.0 sp1-route-table.inet.0
                                sp2-route-table.inet.0];
                }
        }
}
```

The following configuration segment assigns the filter classify-customers to router interface fe-5/2/0 as an input packet filter:

```
interfaces fe-5/2/0 {
        unit 0 {
                family inet {
                        filter {
                                input classify-customers;
                        }
                }
        }
}
```

Filter-Based Forwarding Applications

As mentioned above, filter-based forwarding can be used to support the following service provider applications, including:

- Open access
- BGP/MPLS VPNs (RFC 2547bis)
- Traffic engineering without MPLS

Open Access

One of the largest applications for filter-based forwarding is to support open-access requirements or Carrier supporting Carrier (CsC) applications. This is the data networking equivalent of allowing you to pick your own long-distance telephone carrier independently of your local provider that delivers the local loop. Similar to the telephone network, service providers are required to give competitors access to their infrastructure to allow customers to select their own transit provider. Figure 3 illustrates one way to use filter-based forwarding to support a service provider's open-access requirements.

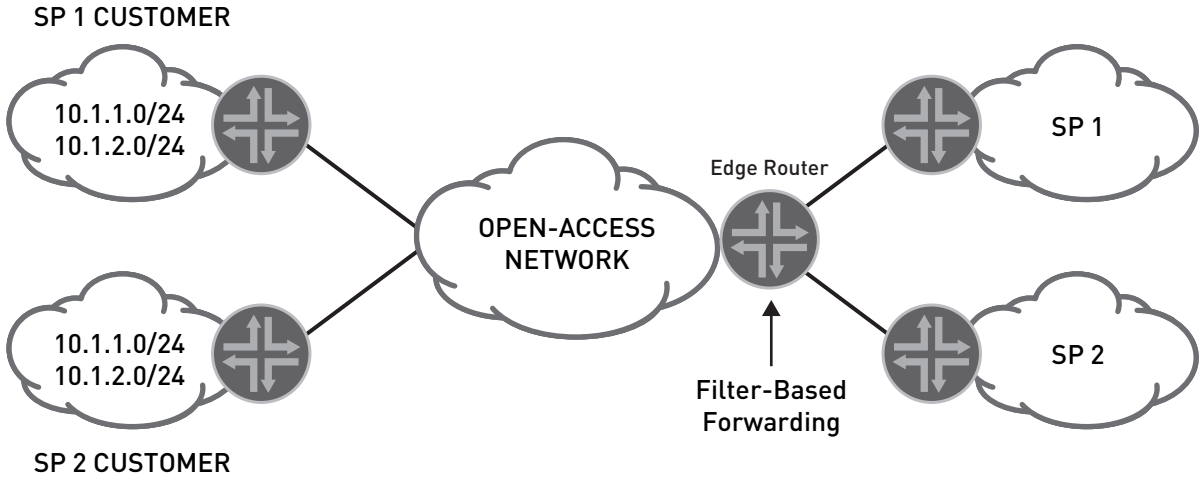

Figure 3: Support for provider open-access requirements

The open-access network provider (with a shared cable modem, DSL or Ethernet infrastructure) has a pool of IP addresses from each of the different service providers for which it is providing customer access. The open-access provider assigns these addresses to the customers of each specific service provider. Since traffic from these subscribers must be forwarded to the appropriate service provider, the forwarding decision for traffic leaving the open-access network cannot be based solely on the destination address carried in each packet header.

In this solution, all customer traffic from open-access network elements is routed using the default route to the access provider's edge router. When a customer packet arrives at the ingress port of the access provider's edge router, filter-based forwarding analysis is performed based on each packet's source address, and the packet is forwarded to the appropriate SP transit provider determined by the results of the packet filter.

BGP/MPLS VPNs (RFC 2547bis)

Another application for filter-based forwarding is to support customer access to a Layer 3 BGP/MPLS VPN across an open-access network. In a classic RFC 2547bis VPN, each PE router forwards customer traffic using the VPN Routing and Forwarding (VRF) table that is associated with the packet's incoming interface. Typically, one or more router interfaces are associated with, or bound to, a VPN by including the interfaces in the configuration of the VPN routing instance. By binding an interface to a VPN, the VPN's VRF table is used to make forwarding decisions for any packets arriving on that interface.

When different CE routes belonging to different VPNs gain access to a common PE router using an open-access network, a mechanism is needed to map each incoming packet to its VRF table so that it can be forwarded correctly. Filter-based forwarding provides a non-classic solution that allows a PE router to determine which traffic is assigned to a VPN based on attributes other than the physical or logical interface binding (Figure 4).

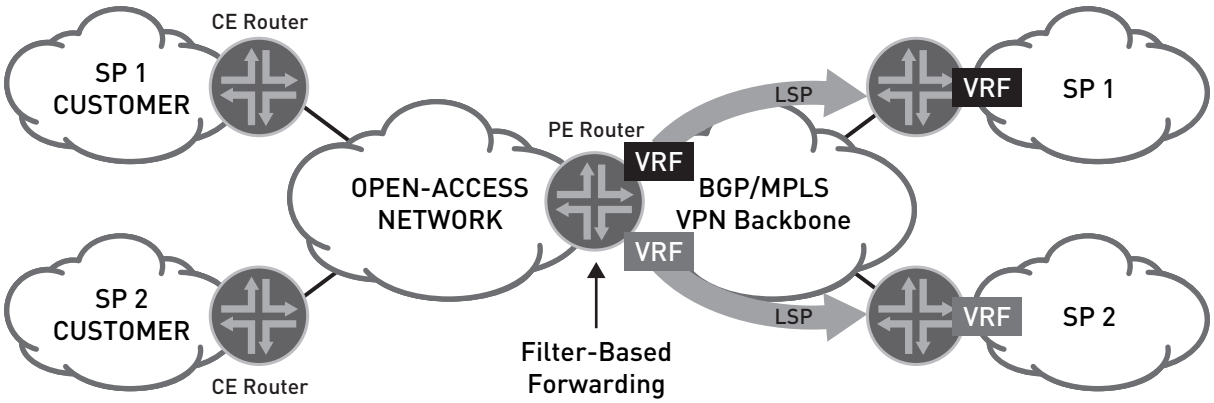

Figure 4: Support for BGP/MPLS VPNs

For this example, assume that VPN_Red is associated with SP 1 and VPN_Blue is associated with SP 2. The PE router attached to the open-access network is configured with an inbound packet filter that determines which traffic is forwarded using VRF_Red (SP 1's customers) and which traffic is forwarded according to VRF_Blue (SP 2's customers).

As a service provider, filter-based forwarding gives you tremendous flexibility when delivering BGP/MPLS VPN services:

- You can deliver a one-way VPN where outbound traffic from a customer site flows across the BGP/MPLS VPN infrastructure while return traffic follows the best-effort IP routed path.
- You can configure a PE router so that certain traffic arriving on the open-access network interface follows a BGP/MPLS VPN path while other traffic arriving on the same interface is forwarded using the best-effort routes contained in inet.0.
- Interfaces of other CE routers attached to other PE routers can be physically mapped to a VRF table if a CE router's open-access network interface is not physically mapped to a VRF table.

The alternate approach to supporting these connectivity options requires a two-router solution. The first router performs policy-based routing to classify packets and forward them out a specific physical or logical interface. The second router functions as the PE router and maps traffic to a VRF table based on its incoming physical or logical interface (Figure 5).

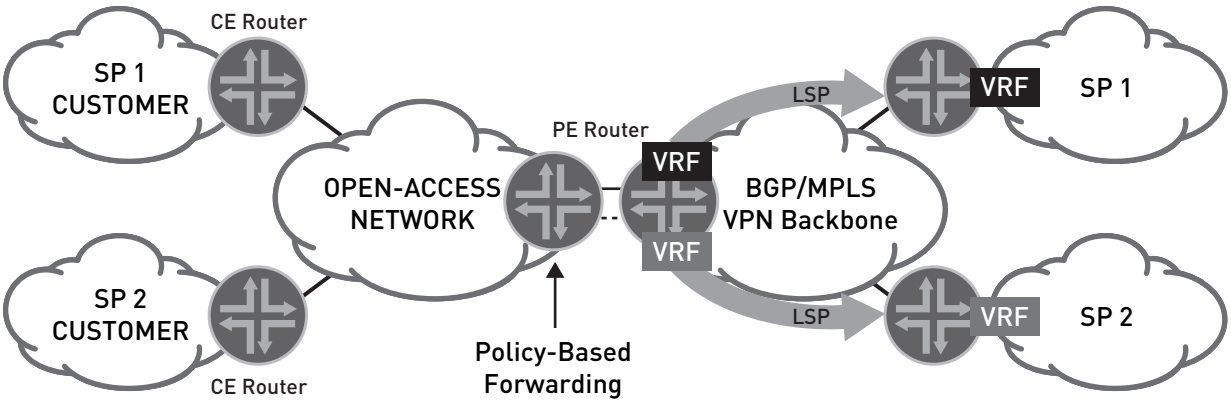

Figure 5: The alternative policy-based routing solution

Traffic Engineering Without MPLS

You can also use filter-based forwarding to support a rudimentary form of traffic engineering without deploying MPLS. This example also illustrates how you can use filter-based forwarding to support application-based forwarding. Consider the network topology illustrated in Figure 6.

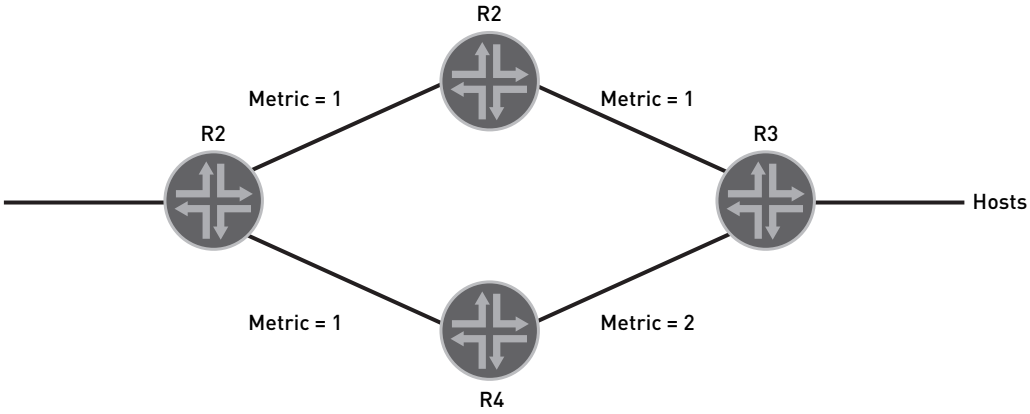

Figure 6: Support for rudimentary traffic engineering

Using conventional destination-based IP routing, all traffic arriving at R1 that is addressed to hosts downstream of R3 follow the least-cost IGP metric path from R1 to R2 to R3. However, to support specific customer requirements, assume that you might want to forward all voice over IP (VoIP) traffic along the less congested but higher cost IGP metric path from R1 to R4 to R3. You can accomplish this model by configuring filter-based forwarding on R1 to forward all VoIP traffic that is addressed to specific hosts downstream of R3 to a next hop of R4 rather than to R2.

Despite the performance advantage of executing filter-based forwarding hardware, it is important to note that the traffic engineering/application-based forwarding application has scalability limitations that materialize when the next hop disappears or when redirection is attempted across multiple hops. These limitations are not related to the Junos OS implementation. Rather, they are a natural consequence of deploying static routing.

Summary

Filter-based forwarding provides a policy-based routing table selection tool that you can use to support a number of critical applications, including open access, BGP/MPLS VPNs and traffic engineering. Filter-based forwarding is another mechanism for controlling the flow of subscriber traffic in service provider networks.

Acronyms

ASIC	Application-Specific Integrated Circuit
BGP	Border Gateway Protocol
CsC	Carrier Supporting Carrier
CE	Customer Edge
DSL	Digital Subscriber Line
DSCP	DiffServ Code Point
ICMP	Internet Control Message Protocol
IP	Internet Protocol
SP	Service Provider
MPLS	Multiprotocol Label Switching
PE	Provider Edge
PLP	Packet Loss Priority
RFC	Request for Comments
TCP	Transmission Control Protocol
TE	Traffic Engineering
UDP	User Datagram Protocol
VPN	Virtual Private Network
VRF	VPN Routing and Forwarding

About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

3500136-002-EN May 2010