

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Configure](#)

[Network Diagram](#)

[Configurations](#)

[Verify](#)

[Case 1: Parent route via BGP](#)

[Case 2: Parent route is via EIGRP](#)

[Related Cisco Support Community Discussions](#)

Introduction

This document describes how Performance Routing version 2 (PfRv2) controls traffic depending on the PfRv2 policy decision. The method and criteria used to control traffic depends on the underlying protocol via which parent route is learnt. In this document, PfRv2 traffic control action is demonstrated when parent route is learnt via BGP and EIGRP.

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of Performance Routing (PfR).

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

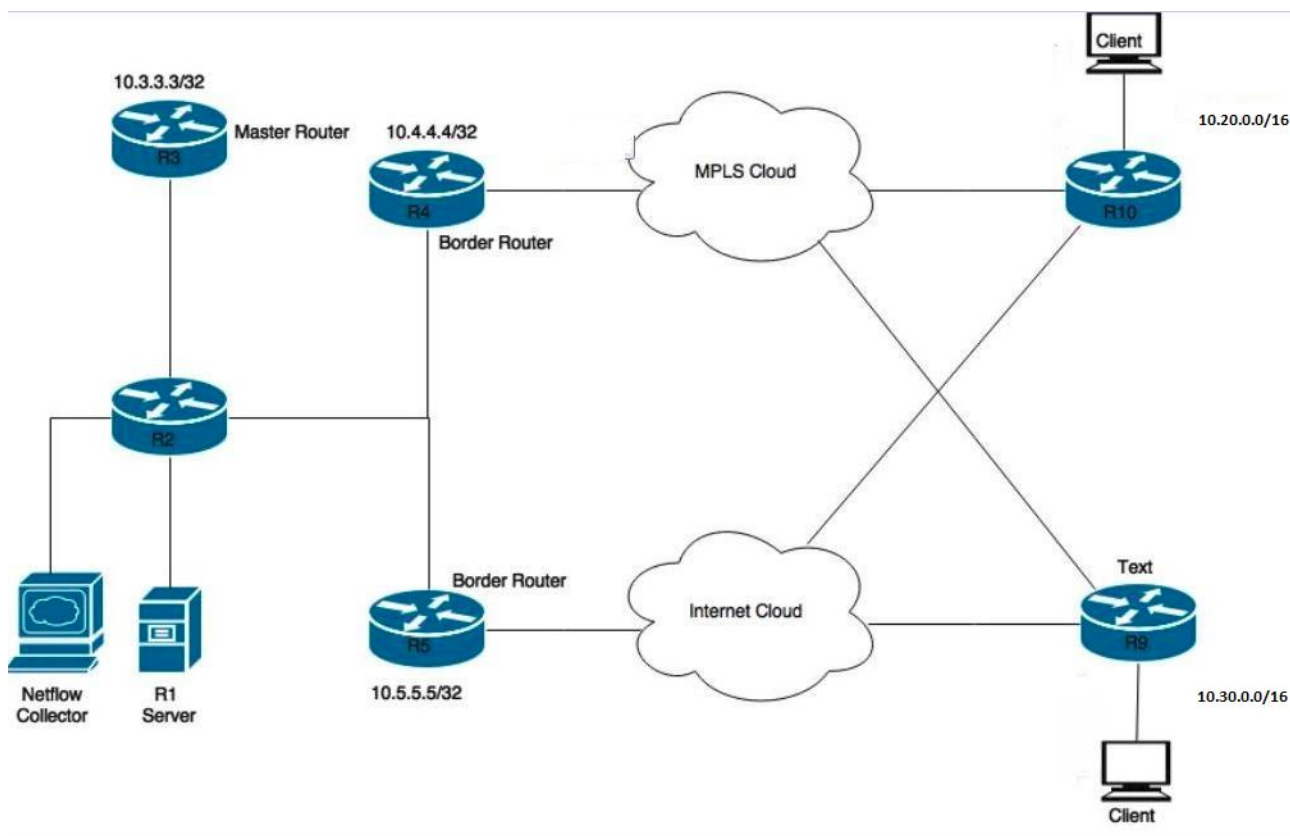
Configure

PfRv2 allows network administrator to configure learn-list to group traffic, apply configured policy and choose best border router (BR) that meets certain set of parameters like delay, jitter, utilization etc defined in policy. There are various modes in which PfRv2 controls traffic and it depends on the protocol via which the parent route for destination prefix is learnt. PfRv2 is capable of changing routing information base (RIB) by manipulating routing protocols, injecting static routes or via dynamic policy based routing. Below is a table that highlights route control method for various protocols.

Parent route	Prefix control method
BGP	BGP via modifying local preference
EIGRP	EIGRP via injecting more specific route
Static	Static via injecting more specific route
RIP,OSPF,ISIS	Dynamic policy based routing

Network Diagram

This document would refer following image as a sample topology for rest of the document.



Devices shown in the diagram:

R1- Server, Initiating traffic.

R3- PfR Master Router.

R4&R5- PfR Border Router.

Clients connected to R9 & R10 are devices receiving the traffic from the R1 server.

Configurations

Verify

Case 1: Parent route via BGP

In this case parent route for both prefixes i.e. 10.20.0.0/16 and 10.30.0.0/16 are learnt via BGP. Below is an output for parent route from both border routers (R4 and R5).

```
R4#show ip route
--output suppressed--
B    10.20.0.0/16 [20/0] via 10.0.46.6, 01:26:58
B    10.30.0.0/16 [20/0] via 10.0.46.6, 01:26:58
```

```
R5#show ip route
--output suppressed--
B    10.20.0.0/16 [20/0] via 10.0.57.7, 00:42:37
B    10.30.0.0/16 [20/0] via 10.0.57.7, 00:42:37
```

There is active flow of traffic for both traffic classes and both could be seen in INPOLICY state in below outputs. R4 can be seen below to be selected for prefix 10.20.0.0/24 and R5 has been selected for prefix 10.30.0.0/24. This is as per configured link-group preference for each learn-

list.

R3#show pfr master traffic-class

OER Prefix Statistics:

Pas - Passive, Act - Active, S - Short term, L - Long term, Dly - Delay (ms),
P - Percentage below threshold, Jit - Jitter (ms),
MOS - Mean Opinion Score
Los - Packet Loss (percent/10000), Un - Unreachable (flows-per-million),
E - Egress, I - Ingress, Bw - Bandwidth (kbps), N - Not applicable
U - unknown, * - uncontrolled, + - control more specific, @ - active probe all
- Prefix monitor mode is Special, & - Blackholed Prefix
% - Force Next-Hop, ^ - Prefix is denied

DstPrefix	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Flags	State	Time	CurrBR	CurrI/F	Protocol		
							PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw
							ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos

10.20.20.0/24		N	N	N		N						N	N	
			INPOLICY	56		10.4.4.4				Et1/0				BGP
		N	N	N		N				N		N		N
		1	2	0		0				N		N		N

10.30.30.0/24		N	N	N		N						N	N	
			INPOLICY	59		10.5.5.5				Et1/0				BGP
		N	N	N		N				N		N		N
		3	2	0		0				N		N		N

Since R4 has been selected by PfRv2 as the exit router for 10.20.20.0/24 ,R4 injects a route with higher local preference for 10.20.20.0/24 as shown below. Properties of injected route are inherited by parent route.

R4#show ip bgp 10.20.20.0/24

BGP routing table entry for 10.20.20.0/24, version 60
Paths: (1 available, best #1, table default, not advertised to EBGp peer)
Advertised to update-groups:
10
Refresh Epoch 1
200, (injected path from 10.20.0.0/16)
10.0.46.6 from 10.0.46.6 (10.6.6.6)
Origin incomplete, metric 0, localpref 100, valid, external, best
Community: no-export
rx pathid: 0, tx pathid: 0x0

Higher local preference is not seen on router that injects the route. Instead it is visible on other BRs that receive this route via iBGP. Below is an example route seen on R5 for prefix 10.20.20.0/24.

R5#show ip bgp 10.20.20.0/24

BGP routing table entry for 10.20.20.0/24, version 17
Paths: (1 available, best #1, table default)
Advertised to update-groups:
6
Refresh Epoch 1
200
10.0.45.4 from 10.0.45.4 (10.4.4.4)
Origin incomplete, metric 0, localpref 5000, valid, internal, best
rx pathid: 0, tx pathid: 0x0

Therefore, any traffic that is received by R5 for prefix 10.20.20.0/24 is routed to R4 back so as traffic could exit the BR selected by PfRv2.

R4#show pfr border routes bgp

BGP table version is 60, local router ID is 10.4.4.4

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
 x best-external, a additional-path, c RIB-compressed,
 Origin codes: i - IGP, e - EGP, ? - incomplete
 RPKI validation codes: V valid, I invalid, N Not found
 OER Flags: C - Controlled, X - Excluded, E - Exact, N - Non-exact, I - Injected

Network	Next Hop	OER	LocPrf	Weight	Path
*> 10.20.20.0/24	10.0.46.6	CEI	5000	0	200 ?
*>i10.30.30.0/24	10.0.45.5	XN	5000	0	300 ?

For prefix 10.20.20.0/24 three flags can be seen. 'C' (controlled) means that route was locally controlled and injected. 'E' (exact) means that this route is exact and is present in BGP table and there is no more specific route present than this. 'I' (injected) says that this route has been locally injected on this router.

Likewise for prefix 10.30.30.0/24, two flags can be seen. 'X' (excluded) shows that this route has not been locally injected and instead originated at some other BR, R5 in our case. And with 'X' flag set, 'N' flag can be ignored.

An important thing to note is by default the route injected carries a local preference value of 5000. Therefore, if your BGP policy is already using some value that is higher than 5000, there might be a problem and results seen could not be expected. You can tune default local preference value by following command.

R4#show pfr border routes bgp

BGP table version is 60, local router ID is 10.4.4.4
 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
 r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
 x best-external, a additional-path, c RIB-compressed,
 Origin codes: i - IGP, e - EGP, ? - incomplete
 RPKI validation codes: V valid, I invalid, N Not found
 OER Flags: C - Controlled, X - Excluded, E - Exact, N - Non-exact, I - Injected

Network	Next Hop	OER	LocPrf	Weight	Path
*> 10.20.20.0/24	10.0.46.6	CEI	5000	0	200 ?
*>i10.30.30.0/24	10.0.45.5	XN	5000	0	300 ?

Case 2: Parent route is via EIGRP

Consider this case where parent route for both prefixes i.e. 10.20.0.0/16 and 10.30.0.0/16 are learnt via EIGRP. Below is an output for parent route from both border routers (R4 and R5). In this case these routes are external however these might be internal eigrp parent routes as well depending on network design.

R4#show ip route

```
--output suppressed--
D EX    10.20.0.0/16 [170/25651200] via 10.0.46.6, 00:04:25, Ethernet1/0
D EX    10.30.0.0/16 [170/25651200] via 10.0.46.6, 00:04:25, Ethernet1/0
```

R5#show ip route

```
--output suppressed--
D EX    10.20.0.0/16 [170/25651200] via 10.0.57.7, 00:05:46, Ethernet1/0
D EX    10.30.0.0/16 [170/25651200] via 10.0.57.7, 00:05:46, Ethernet1/0
```

As shown in previous case, there is active flow of traffic for both traffic classes and both could be seen in INPOLICY state in below output. R4 has been selected for prefix 10.20.20.0/24 and R5 has been selected for prefix 10.30.30.0/24. This is as per configured link-group preference for each learn-list.

R3#show pfr master traffic-class

OER Prefix Statistics:

Pas - Passive, Act - Active, S - Short term, L - Long term, Dly - Delay (ms),
 P - Percentage below threshold, Jit - Jitter (ms),
 MOS - Mean Opinion Score

Los - Packet Loss (percent/10000), Un - Unreachable (flows-per-million),
 E - Egress, I - Ingress, Bw - Bandwidth (kbps), N - Not applicable
 U - unknown, * - uncontrolled, + - control more specific, @ - active probe all
 # - Prefix monitor mode is Special, & - Blackholed Prefix
 % - Force Next-Hop, ^ - Prefix is denied

DstPrefix	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix		
Flags	State		Time	CurrBR	CurrI/F	Protocol		
PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw	
ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos	

10.20.20.0/24		N	N	N		N	N	
		INPOLICY		31	10.4.4.4	Et1/0		EIGRP
	N	N	N	N	N	N	N	N
	1	2	0	0	N	N	N	N
10.30.30.0/24		N	N	N		N	N	
		INPOLICY		24	10.5.5.5	Et1/0		EIGRP
	N	N	N	N	N	N	N	N
	2	2	0	0	N	N	N	N

Since R4 has been selected by PfRv2 as the best exit router for 10.20.20.0/24 ,R4 injects a more specific route with tag 5000 as shown below. This injected route is always an EIGRP internal route even if parent route is external. Also if parent route carries a tag value, that is not inherited by injected route.

Note: Not all properties of injected route are inherited by parent route.

R4#show ip route 10.20.20.0 255.255.255.0

Routing entry for 10.20.20.0/24

Known via "eigrp 100", distance 90, metric 25651200

Tag 5000, type **internal**

Redistributing via eigrp 100

Last update from 10.0.46.6 on Ethernet1/0, 00:17:04 ago

Routing Descriptor Blocks:

* 10.0.46.6, from 0.0.0.0, 00:17:04 ago, via Ethernet1/0

Route metric is 25651200, traffic share count is 1

Total delay is 2000 microseconds, minimum bandwidth is 100 Kbit

Reliability 255/255, minimum MTU 1500 bytes

Loading 12/255, Hops 0

Route tag 5000

R4#show ip eigrp topology 10.20.20.0/24

EIGRP-IPv4 Topology Entry for AS(100)/ID(10.4.4.4) for 10.20.20.0/24

State is Passive, Query origin flag is 1, 1 Successor(s), FD is 25651200

Descriptor Blocks:

10.0.46.6 (Ethernet1/0), from 0.0.0.0, Send flag is 0x0

Composite metric is (25651200/0), route is Internal

Vector metric:

Minimum bandwidth is 100 Kbit

Total delay is 2000 microseconds

Reliability is 255/255

Load is 12/255

Minimum MTU is 1500

Hop count is 0

Originating router is 10.4.4.4

Internal tag is **5000**

R4#show pfr border routes eigrp

Flags: C - Controlled by oer, X - Path is excluded from control,

E - The control is exact, N - The control is non-exact

Flags	Network	Parent	Tag
CE	10.20.20.0/24	10.20.0.0/16	5000
XN	10.30.30.0/24		

Above case has parent route that was less specific i.e. 10.20.0.0/16 and injecting a more specific route 10.20.20.0/24 provided desired results. Any traffic received on R5 would be redirected to R4 using below route and hence traffic would flow as per PfRv2 selected best exit BR.

R5#show ip route 10.20.20.0

```

Routing entry for 10.20.20.0/24
  Known via "eigrp 100", distance 90, metric 26931200
  Tag 5000, type internal
  Redistributing via eigrp 100
  Last update from 10.0.45.4 on Tunnel10, 00:25:34 ago
  Routing Descriptor Blocks:
  * 10.0.45.4, from 10.0.45.4, 00:25:34 ago, via Tunnel10    // 10.0.45.4 is R4 IP.
    Route metric is 26931200, traffic share count is 1
    Total delay is 52000 microseconds, minimum bandwidth is 100 Kbit
    Reliability 255/255, minimum MTU 1476 bytes
    Loading 28/255, Hops 1
    Route tag 5000

```

In case parent route is also /24 route, R4 injects a /24 route in a way that makes the injected route more preferred than the parent route.

R4#show ip eigrp topology 10.20.20.0/24

```

EIGRP-IPv4 Topology Entry for AS(100)/ID(10.4.4.4) for 10.20.20.0/24
  State is Passive, Query origin flag is 1, 1 Successor(s), FD is 25600000
  Descriptor Blocks:
  10.0.46.6 (Ethernet1/0), from 0.0.0.0, Send flag is 0x0
    Composite metric is (25600000/0), route is Internal
    Vector metric:
      Minimum bandwidth is 100 Kbit
      Total delay is 1 microseconds           // Injected route with a delay of 1.
      Reliability is 255/255
      Load is 102/255
      Minimum MTU is 1500
      Hop count is 0
      Originating router is 10.4.4.4
      Internal tag is 5000
  10.0.45.5 (Tunnel10), from 10.0.45.5, Send flag is 0x0
    Composite metric is (26931200/25651200), route is External
    Vector metric:
      Minimum bandwidth is 100 Kbit
      Total delay is 52000 microseconds
      Reliability is 255/255
      Load is 99/255
      Minimum MTU is 1476
      Hop count is 2
      Originating router is 10.0.78.7
  External data:
    AS number of route is 0
    External protocol is Static, external metric is 0
    Administrator tag is 0 (0x00000000)
  10.0.46.6 (Ethernet1/0), from 10.0.46.6, Send flag is 0x0    //Parent route
    Composite metric is (25651200/281600), route is External
    Vector metric:
      Minimum bandwidth is 100 Kbit
      Total delay is 2000 microseconds
      Reliability is 255/255
      Load is 102/255
      Minimum MTU is 1500
      Hop count is 1
      Originating router is 10.0.68.6

```

```
External data&colon;  
AS number of route is 0  
External protocol is Static, external metric is 0  
Administrator tag is 0 (0x00000000)
```

As shown above, when parent route and injected prefix are of same subnet mask, injected route inherits minimum bandwidth, load, reliability, MTU etc from parent route but delay of injected route is set less and hence this becomes a preferred route. So, when traffic is received on other BR i.e. R5, R5 can send traffic via this route with better metric to R4 and R4 would then send it out of its exit interface in agreement with PfRv2.