

使用BGP或EIGRP配置PfRv2流量控制机制

目录

[简介](#)

[先决条件](#)

[要求](#)

[使用的组件](#)

[配置](#)

[网络图](#)

[配置](#)

[验证](#)

[第 1 种情况：通过BGP的父路由](#)

[案例2：父路由通过EIGRP](#)

[相关的思科支持社区讨论](#)

简介

本文档介绍性能路由版本2(PfRv2)如何根据PfRv2策略决策控制流量。用于控制流量的方法和标准取决于获取父路由所依据的底层协议。在本文档中，当通过BGP和EIGRP获知父路由时，将演示PfRv2流量控制操作。

先决条件

要求

思科建议您具备性能路由(PfR)的基本知识。

使用的组件

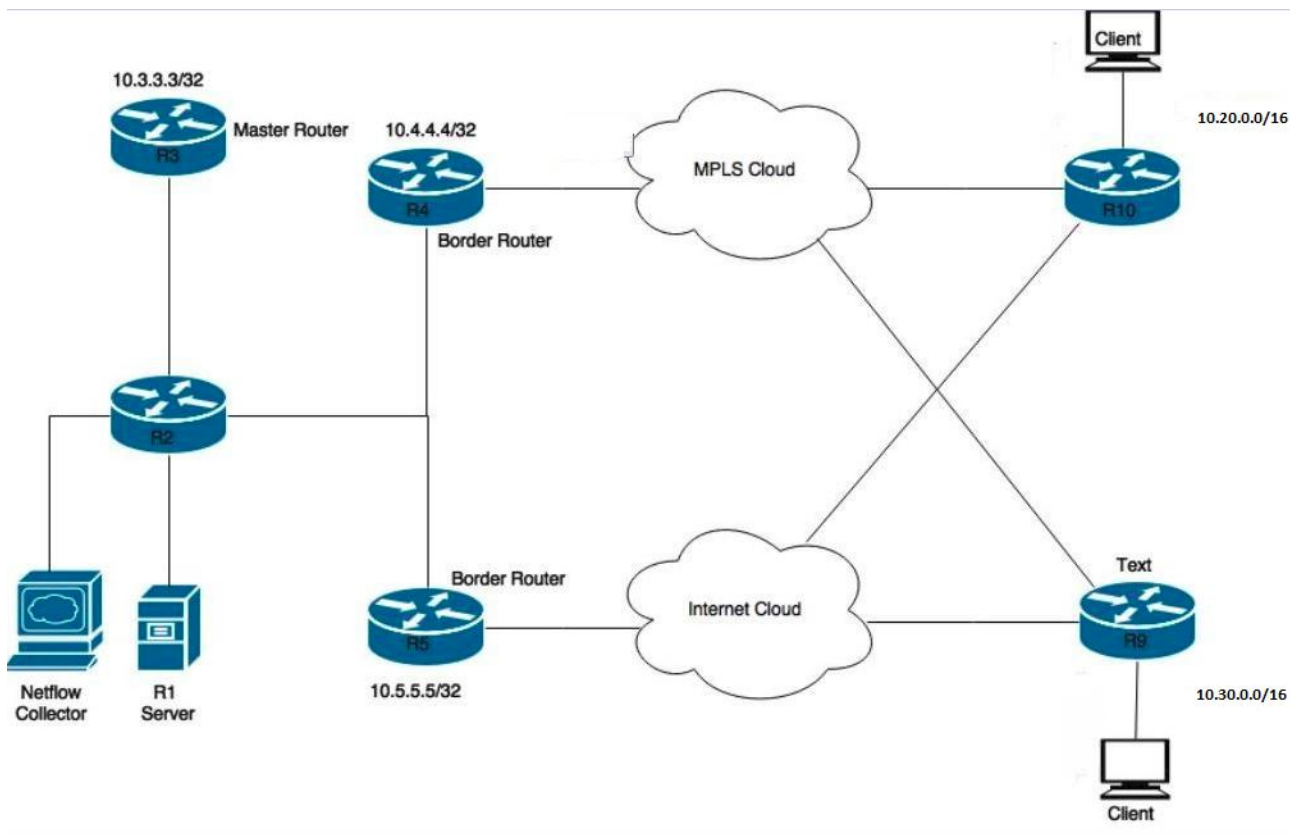
配置

PfRv2允许网络管理员配置学习列表以分组流量，应用已配置的策略并选择满足策略中定义的延迟、抖动、利用率等特定参数集的最佳边界路由器(BR)。PfRv2控制流量有多种模式，它取决于用于获取目的前缀的父路由的协议。PfRv2能够通过操作路由协议、注入静态路由或基于动态策略的路由来更改路由信息库(RIB)。下表重点介绍了各种协议的路由控制方法。

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

网络图

本文档将以下图像作为文档其余部分的示例拓扑。



R1 —
R3- PfR
R4&R5- PfR
R9R10R1

配置

```
!  
key chain pfr  
  key 0  
  key-string cisco  
pfr master  
  policy-rules PFR  
  !  
  border 10.4.4.4 key-chain pfr  
  interface Ethernet1/0 external  
  interface Ethernet1/2 internal  
  link-group MPLS  
  !  
  border 10.5.5.5 key-chain pfr  
  interface Ethernet1/3 internal  
  interface Ethernet1/0 external  
  link-group INET  
  !  
learn  
traffic-class filter access-list DENY-ALL  
  list seq 10 rename APPLICATION-LEARN-LIST  
  traffic-class prefix-list APPLICATION
```

```

    throughput
list seq 20 refname DATA-LEARN-LIST
    traffic-class prefix-list DATA
    throughput
!
pfr-map PFR 10
match pfr learn list APPLICATION-LEARN-LIST
set periodic 90
set delay threshold 25
set mode monitor active
set active-probe echo 10.20.21.1
set probe frequency 5
set link-group MPLS fallback INET
!
pfr-map PFR 20
match pfr learn list DATA-LEARN-LIST
set periodic 90
set delay threshold 25
set mode monitor active
set active-probe echo 10.30.31.1
set probe frequency 5
set link-group INET fallback MPLS
!
ip prefix-list APPLICATION: 1 entries
    seq 5 permit 10.20.0.0/16
!
ip prefix-list DATA: 1 entries
    seq 5 permit 10.30.0.0/16
!

```

验证

1 BGP

(10.20.0.0/16 10.30.0.0/16) BGPR4R5

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

```

INPOLICY 10.20.20.0/24 R4 10.30.30.0/24 R5

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	PasSJos	PasLJos	EBw	IBw	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSJos	ActLJos
10.20.20.0/24		N	N	N		N	N								
			INPOLICY	56	10.4.4.4	Et1/0									BGP
	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
	1	2	0	0	N	N	N	N	N	N	N	N	N	N	N
10.30.30.0/24		N	N	N		N	N								
			INPOLICY	59	10.5.5.5	Et1/0									BGP
	N	N	N	N	N	N	N	N	N	N	N	N	N	N	N
	3	2	0	0	N	N	N	N	N	N	N	N	N	N	N

由于Pfrv2已选择R4作为10.20.20.0/24的送出路由器，因此R4会注入本地优先级较高的10.20.20.0/24路由，如下所示。注入路由的属性由父路由继承。

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
```

iBGPBRR510.20.20.0/24

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
```

R510.20.20.0/24R4Pfrv2BR

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 ?

10.20.20.0/24“C”“E”BGP“1”

10.30.30.0/24“X”BRR5“X”“N”

5000BGP5000

R3(config-pfr-mc)#mode route metric bgp local-pref

2:EIGRP

EIGRP(10.20.0.0/1610.30.0.0/16)R4R5 eigrp

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

如上例所示，两个流量类都有活动的流量，在以下输出的INPOLICY状态中可以看到这两个流量。已为前缀10.20.20.0/24选择R4，为前缀10.30.30.0/24选择R5。这与每个学习列表的已配置链路组首选项相同。

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

Table with 8 columns: DstPrefix, Appl_ID, Dscp, Prot, SrcPort, DstPort, SrcPrefix, and Protocol. It shows statistics for two prefixes: 10.20.20.0/24 and 10.30.30.0/24, including flags like INPOLICY and various performance metrics.

由于Pfrv2已选择R4作为10.20.20.0/24的最佳送出路由器，因此R4使用标签5000注入更具体的路由，如下所示。即使父路由是外部路由，此注入路由也始终是EIGRP内部路由。此外，如果父路由带有标记值，则注入路由不会继承该标记值。

注意：并非注入路由的所有属性都由父路由继承。

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		

10.20.0.0/16:10.20.20.0/24R5R4Pfrv2BR

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
```

/24R4/24

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:
    AS number of route is 0
    External protocol is Static, external metric is 0
    Administrator tag is 0 (0x00000000)
```

MTUBRR5R5R4R4Pfrv2