

Cisco Performance Routing (PfR) Solution Guides

PfR - Enterprise Data dual-homed to two IP-VPN Service Providers Path Enforcement using PBR

Navigation

- [Go to PfR home page](#)
- [Go to PfR Solution Guides home page](#)

Contents

- [1 PfR Features that Enable Load Balancing](#)
 - ◆ [1.1 Link Utilization](#)
 - ◆ [1.2 Range](#)
 - ◆ [1.3 Traffic Class Performance](#)
- [2 Enterprise Needs](#)
- [3 PfR Solution Used](#)
- [4 PfR Network Topology Used](#)
- [5 PfR Components Configuration](#)
 - ◆ [5.1 Master Controller](#)
 - ◆ [5.2 Border Routers](#)
- [6 Flexible Netflow](#)
- [7 Checking Statistics and Flows](#)
- [8 Master Controller Verification](#)
 - ◆ [8.1 Master Controller and Traffic Classes](#)
 - ◆ [8.2 Display Learn-list](#)
 - ◆ [8.3 Policy Configuration](#)
- [9 Verify Traffic Classes Statistics](#)
 - ◆ [9.1 Traffic Classes](#)
 - ◆ [9.2 A closer look at the results](#)
 - ◇ [9.2.1 Best Effort TC](#)
 - ◇ [9.2.2 Business TC](#)
 - ◇ [9.2.3 Active Probes](#)

- 10 Verify Enforcement
 - ◆ 10.1 Policy-Based Routing used
 - ◆ 10.2 BGP Routes by default
 - ◆ 10.3 Business Traffic Class
 - ◆ 10.4 Best Effort Traffic Class
 - ◆ 10.5 Dynamic PBR on Border Router R4
 - ◆ 10.6 Dynamic PBR on Border Router R5
- 11 Add Delay on SP1
 - ◆ 11.1 ?normal? delay across SP1
 - ◆ 11.2 Adding a delay of 300ms on SP1
 - ◆ 11.3 Move TC to new exit ? HOLDDOWN State
 - ◆ 11.4 New exit ? INPOLICY State
- 12 Back to Normal
- 13 Conclusion

PfR Features that Enable Load Balancing

Link Utilization

Usage of this policy sets an upper threshold on the amount of traffic a specific link can carry. For example, if the upper threshold for a link is 90 % of total bandwidth, and it is currently running at 95 % of bandwidth, the link is Out-of-Policy (OOP). Cisco PfR will attempt to bring the link back into policy by repeatedly moving prefixes from the over-used link onto other exit links.


Range

Usage of this policy keeps all WAN links within a certain utilization range, relative to each other in order to ensure fair load-sharing across all concerned links. The range functionality allows the network administrator to instruct Cisco PfR to keep the usage on a set of exit links within a certain percentage range of each other. If the difference between the links becomes too great, Cisco PfR will attempt to bring the link back in to policy by distributing data traffic among the available exit links.

Traffic Class Performance

Usage of this policy enables the customer to define multiple paths that a set of traffic (ie voice traffic) could use as long as all the paths maintain the performance SLA's that are needed for that set of traffic. Hence, a policy that determines voice traffic to have a delay threshold of less than 250 msec can utilize multiple paths in the network if available, as long as all the paths deliver the traffic within its performance bounds.

Enterprise Needs



This solution is for an Enterprise which is dual-attached to two IP-VPN Service Providers. The load-balancing used here is taking place between external interfaces and specific policies are to be applied depending on the traffic type. In most (all) cases, Qos is already in place and classification/marketing procedures were studied a while back. Therefore packets entering on the Border Router are already classified and marked directly on the access switch connecting the station, IP Phone or multimedia terminal. That means PfR can probably use the dscp field as the most efficient way for the traffic profiling phase.

- **BUSINESS:** this group encompass all the critical applications. In most cases, these are transactional applications and by nature are delay intolerant. One of the key goal is to protect these applications and to be able to track variations in SLA. Traffic in this group is marked with DSCP=AF31.
- **BEST-EFFORT (BE):** traffic that has a low priority but could have a high bandwidth. Traffic in this group is marked with DSCP=0.

Requirements:

- The central site is dual-homed, the primary path is WAN1 (considered as the main IP-VPN with strict SLAs) and the secondary path is WAN2.
- Simple Primary/Backup load balancing solution.
- Business traffic over the primary path SP1 and low priority (email, bulk, internet) traffic over the back-up path SP2.
- When traffic reaches x% utilization or delay cross a threshold, move traffic over to the alternate path.
- Traffic Classes based on DSCP values. To keep this example simple, we have defined two Traffic Classes.

PfR Solution Used

This solution describes a more advanced load sharing and will be based on automatic application classification based on DSCP, learning-list and specific policies per group of traffic. As there is a requirement to move some Traffic Classes over a specific path, a new feature called link-group will be used for that.

Traffic is divided in 2 majors groups which will have specific policies applied. The main PfR features used in this configuration are the following (there are others but these ones are the most important):

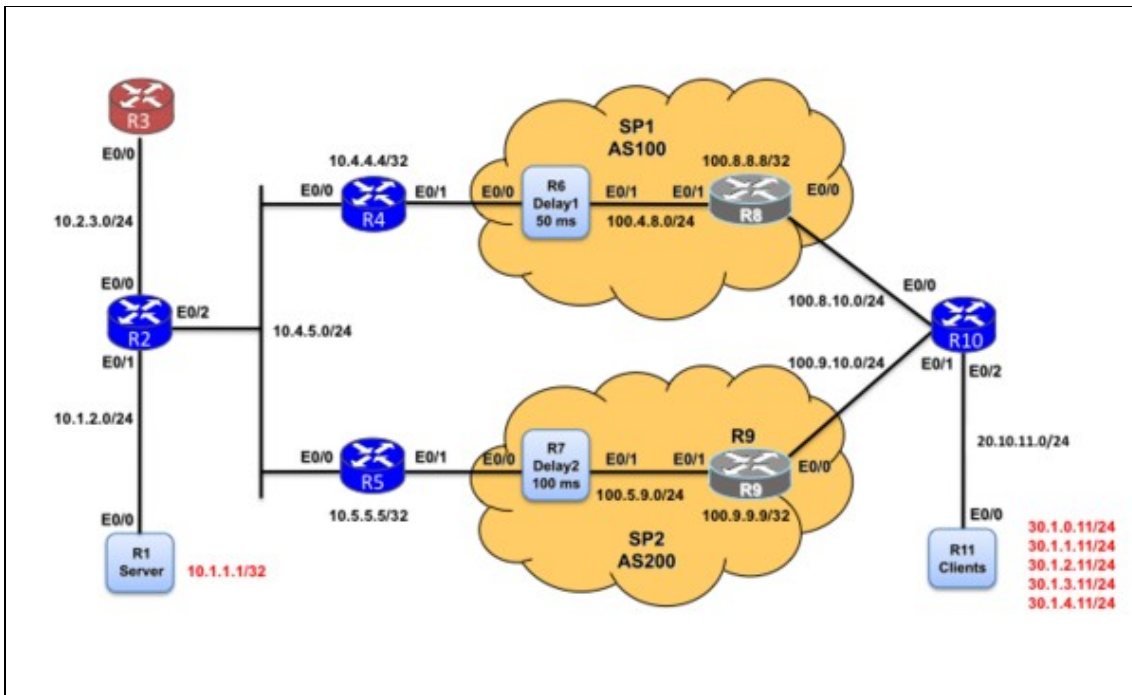
Traffic Profiling:

- **BUSINESS Policy:** traffic marked with DSCP AF31 from central site to branch considered as the critical business traffic. This one should take the path that the lowest delay, keep the link utilization below 90%. The primary path should be WAN1 and the backup path WAN2 if delay or utilization is beyond the defined boundaries. The PfR feature used to prefer one path over another path is link-group.
- **BE Policy:** low priority traffic. These flows should use the secondary path toward WAN2. We have only one policy that applies to this group: unreachable. The goal is to be able to move the low priority traffic over the primary path is something goes wrong with WAN2. Therefore the threshold for unreachable is set to 500000, which equal to 50% of the packets dropped as unreachable is expressed in flow per million (fpm).

PfR Network Topology Used

The central site has two Border Routers, connected to two separate Service Providers using eBGP. R2, R4 and R5 are iBGP peers.

- R3 is the Master Controller
- R4 and R5 the Border Routers
- Packet generation tool is used between R1 and R11 to emulates traffic
- R6 and R7 are delay generators that add delay/loss to the path through SP1 and SP2. By default, 100 ms through SP1 and 50 ms through SP2.
- R1 and R11 are packet generators that send/receive traffic (http, ssh, etc).



PfR Components Configuration

Master Controller

```

!
key chain pfr
  key 0
    key-string cisco
!
pfr master
!
! For Each TC, Associate a Policy
!
policy-rules MYMAP
!
! Following configuration is
! - to enable logging. This will print PfR related syslog messages on the console.
! ? to disable default policy of load balancing.
!
no max-range-utilization
logging
!
! =====
! Border Routers Definition
! =====
!
border 10.4.5.4 key-chain pfr
  interface Ethernet0/0 internal
  interface Ethernet0/1 external
  ! route excess traffic via Public WAN when the utilization in MPLS link reaches 80%
  max-xmit-utilization percentage 90
  link-group WAN1
!

```

PfR:Solutions:EnterpriseData

```
border 10.4.5.5 key-chain pfr
 interface Ethernet0/0 internal
 interface Ethernet0/1 external
   max-xmit-utilization percentage 90
   link-group WAN2
!
! =====
! Learning based on learn-list
! =====
!
learn
 throughput
 delay
!
! to enable continuous learn cycle, each of 1 minute duration.
! Sort the traffic-class based on ?throughput? at the end of
! each learning cycle.
!
 periodic-interval 0
 monitor-period 1
 list seq 10 refname BUSINESS
   traffic-class access-list BUSINESS
   throughput
 list seq 20 refname BE
   traffic-class access-list BE
   throughput
 holddown 90
 backoff 180 180
 mode route control
 periodic 180
 no resolve range
 no resolve utilization
!
! =====
! ACCESS-LIST TO MATCH DSCP
! =====
!
!
ip access-list extended BE
 permit ip any any dscp default
ip access-list extended BUSINESS
 permit ip any any dscp af31
!
! =====
! PFR POLICIES FOR BUSINESS
! =====
!
pfr-map MYMAP 10
 match pfr learn list BUSINESS
 set mode select-exit good
 set delay threshold 200
 set mode route control
 set mode monitor both
 set resolve delay priority 1 variance 20
 set resolve utilization priority 3 variance 20
 no set resolve range
 set probe frequency 30
 set link-group WAN1 fallback WAN2
!
!
! =====
! PFR POLICIES FOR BEST EFFORT
! =====
!
```

```
pfr-map MYMAP 20
match pfr learn list BE
set mode select-exit good
set mode route control
set mode monitor both
no set resolve delay
no set resolve range
no set resolve utilization
set unreachable threshold 500000
set probe frequency 30
set link-group WAN2 fallback WAN1
!
```

Notes:

- no max-range-utilization: unused here.
 - logging: enable PfR Syslogs (can be checked using show logging).
 - max-xmit-utilization percentage 90: If utilization on an external link exceeds 90% of the configured bandwidth, PfR will detect a Load OOP condition.
 - mode route control: PfR control the routes
 - periodic 180: policies are reevaluated every 180s (optional)
 - link-group: Define Link Group interfaces WAN1 and WAN2.
-
- Learn: Periodic-interval: set the time interval between prefix learning periods
 - Learn: Monitor period: Set the time period in which PfR learns traffic flows.
 - learn-list BUSINESS: learn traffic that match access-list BUSINESS (DSCP AF31)
 - learn-list BE: learn traffic that match access-list BUSINESS (DSCP 0)
-
- pfr-map MYMAP 10: Mode both, Delay set in absolute mode to 200 ms. Resolve on delay then utilization. Choose links from WAN1, fallback to WAN2 if OOP
 - pfr-map MYMAP 20: Mode both, resolve on utilization. Choose link from WAN2, fallback on WAN1 if unreachable above 50%

Border Routers

```
!
!=====
! Key-chain Definition
! Must match the one defined on MC
!=====
!
key chain pfr
  key 0
    key-string cisco
!
!
!=====
! Define the master controller address
!=====
```

```
!  
pfr border  
  local Ethernet0/0  
  master 10.2.3.3 key-chain pfr  
!  
!  
interface Ethernet0/0  
  description --INTERNAL--  
  ip address 10.4.5.4 255.255.255.0  
  ip ospf 100 area 0  
  load-interval 30  
!  
interface Ethernet0/1  
  description --WAN1--  
  bandwidth 500  
  ip address 100.4.8.4 255.255.255.0  
  load-interval 30  
!
```

Notes:

- The Master Controller needs to be configured to talk to the border routers. The IP address used is the one referred by the ?local <interface>? configuration on the BR.
- Netflow configuration is not needed for PfR to operate. Just configure here as a mean to verify flows that are crossing this BR.
- External interface: Don't forget to configure external interfaces with the contracted bandwidth so that PfR knows how to share load. Also set the load interval to 30 seconds for better accuracy.

Flexible Netflow

While configuring Netflow is not a mandatory task for PfR to work, it allows to have a good understanding of the traffic flows across the border routers.

Flow Record Definition

```
!  
flow record MYRECORD  
  match ipv4 protocol  
  match ipv4 source address  
  match ipv4 destination address  
  match transport source-port  
  match transport destination-port  
  match interface input  
  collect ipv4 dscp  
  collect interface output  
  collect counter bytes  
  collect counter packets  
!
```

Flow Monitor Definition

```
flow monitor MYMONITOR  
  record MYRECORD
```


!

And then apply the FNF Monitor on the interface

```
interface Ethernet0/0
 ip flow monitor MYMONITOR input
!
```

Checking Statistics and Flows

As explained before, explicitly enabling Netflow is not required for PfR to run but is a good practice to check active flows crossing the Border Routers, verify the ingress/egress interfaces used (must be internal to external or vice-versa).

In this example, we check if traffic is flowing through the border router R4 (by default R4 is the exit point according to BGP):

```
R4#sh flow monitor MYMONITOR cache format table
Cache type:                Normal
Cache size:                 4096
Current entries:           126
High Watermark:            287

Flows added:                2381
Flows aged:                2255
- Active timeout          ( 1800 secs)    0
- Inactive timeout        (   15 secs)    2255
- Event aged              0
- Watermark aged         0
- Emergency aged         0
```

IPV4 SRC ADDR	IPV4 DST ADDR	TRNS SRC PORT	TRNS DST PORT	INTF INPUT	IP PROT	int
10.2.3.3	10.4.5.4	3949	45158	Et0/0	6	NuL
10.4.5.2	224.0.0.5	0	0	Et0/0	89	NuL
10.4.5.5	224.0.0.5	0	0	Et0/0	89	NuL
10.1.1.1	30.1.0.11	232	10	Et0/0	6	EtC
10.1.1.1	30.1.2.11	234	10	Et0/0	6	EtC
10.1.1.1	30.1.3.11	21	3007	Et0/0	6	EtC
10.1.1.1	30.1.5.11	22	1075	Et0/0	6	EtC
10.1.1.1	30.1.1.11	22	1076	Et0/0	6	EtC
10.1.1.1	30.1.4.11	22	1077	Et0/0	6	EtC
10.1.1.1	30.1.5.11	22	1078	Et0/0	6	EtC
10.1.1.1	30.1.0.11	235	10	Et0/0	6	EtC
10.1.1.1	30.1.1.11	236	10	Et0/0	6	EtC
10.1.1.1	30.1.2.11	22	1079	Et0/0	6	EtC
10.1.1.1	30.1.5.11	22	1080	Et0/0	6	EtC
10.1.1.1	30.1.2.11	22	1081	Et0/0	6	EtC
10.1.1.1	30.1.3.11	80	3008	Et0/0	6	EtC
10.1.1.1	30.1.1.11	238	10	Et0/0	6	EtC

[snip]

We clearly see that traffic is crossing the BR between internal and external interfaces (E0/0 and E0/1). Next step is to check the Traffic Classes on the Master Controller.

Master Controller Verification

Before starting to look at the results, a few checks are needed to verify that the configuration is correct, that the learning is correctly defined and that policies are well defined and associated with learn-list.

Goal:

- Verify that the MC is active
- Verify that the learning process is enabled on the Master Controller
- Display the policy settings as well as the learning parameters and global timers.

Master Controller and Traffic Classes

The first step is to check the master controller configuration, verify the border routers, verify the parameters used (default and configured) and check the learn-list.

```
MC#sh pfr master
OER state: ENABLED and ACTIVE
  Conn Status: SUCCESS, PORT: 3949
  Version: 3.0
  Number of Border routers: 2
  Number of Exits: 2
  Number of monitored prefixes: 17 (max 5000)
  Max prefixes: total 5000 learn 2500
  Prefix count: total 17, learn 10, cfg 0
  PBR Requirements met
  Nbar Status: Inactive

Border          Status  UP/DOWN          AuthFail  Version
10.4.5.5        ACTIVE  UP               00:41:15  0 3.0
10.4.5.4        ACTIVE  UP               00:41:15  0 3.0

Global Settings:
max-range-utilization percent 0 recv 0
mode route metric bgp local-pref 5000
mode route metric static tag 5000
trace probe delay 1000
logging
exit holddown time 60 secs, time remaining 0

Default Policy Settings:
backoff 180 180 180
delay relative 50
holddown 90
periodic 180
probe frequency 56
number of jitter probe packets 100
```

```
mode route control
mode monitor both
mode select-exit good
loss relative 10
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
resolve delay priority 11 variance 20
```

Learn Settings:

```
current state : STARTED
time remaining in current state : 88 seconds
throughput
delay
no inside bgp
monitor-period 1
periodic-interval 0
aggregation-type prefix-length 24
prefixes 100 appls 100
expire after time 720
```

```
Learn-List seq 10 refname BUSINESS
Configuration:
Traffic-Class Access-list: BUSINESS
Aggregation-type: prefix-length 24
Learn type: throughput
Session count: 50 Max count: 100
Policies assigned: 10
Status: ACTIVE
Stats:
Traffic-Class Count: 4
```

```
Learn-List seq 20 refname BE
Configuration:
Traffic-Class Access-list: BE
Aggregation-type: prefix-length 24
Learn type: throughput
Session count: 50 Max count: 100
Policies assigned: 20
Status: ACTIVE
Stats:
Traffic-Class Count: 6
```

MC#

What to check:

- Both Border Routers are up and running
- Number of Monitored Prefixes: 9 are learned
- All default policy settings are displayed
- Learn is started (current state : STARTED)
- Then all learn-list are displayed
- For each learn-list, check the Traffic Class access-list, the policy number associated (which is the pfr-map it refers to).
- The Traffic Class count is also displayed which allows to check whether the learning process works well.

Display Learn-list

Because we use learn-list, there is a specific command to have a more detailed view:

```
MC#sh pfr master learn list

Learn-List seq 10 refname BUSINESS
Configuration:
  Traffic-Class Access-list: BUSINESS
  Aggregation-type: prefix-length 24
  Learn type: throughput
  Session count: 50 Max count: 100
  Policies assigned: 10
  Status: ACTIVE
Stats:
  Traffic-Class Count: 4
  Traffic-Class Learned:
    Appl Prefix 30.1.1.0/24 af31 256
    Appl Prefix 30.1.3.0/24 af31 256
    Appl Prefix 30.1.2.0/24 af31 256
    Appl Prefix 30.1.0.0/24 af31 256
Learn-List seq 20 refname BE
Configuration:
  Traffic-Class Access-list: BE
  Aggregation-type: prefix-length 24
  Learn type: throughput
  Session count: 50 Max count: 100
  Policies assigned: 20
  Status: ACTIVE
Stats:
  Traffic-Class Count: 5
  Traffic-Class Learned:
    Appl Prefix 30.1.2.0/24 defa 256
    Appl Prefix 30.1.3.0/24 defa 256
    Appl Prefix 30.1.1.0/24 defa 256
    Appl Prefix 30.1.5.0/24 defa 256
    Appl Prefix 30.1.4.0/24 defa 256
MC#
```

For each group, this command displays the prefixes that are dynamically learned as well as the DSCP. In this configuration, learning is application based (DSCP).

Policy Configuration

After the Master Controller and Traffic Classes verification, the second step is to check the policies associated with the Traffic Classes.

```
MC#sh pfr master policy
Default Policy Settings:
  backoff 180 180 180
  delay relative 50
  holddown 90
  periodic 180
  probe frequency 56
  number of jitter probe packets 100
```

```
mode route control
mode monitor both
mode select-exit good
loss relative 10
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
resolve delay priority 11 variance 20
oer-map MYMAP 10
  sequence no. 8444249301975040, provider id 1, provider priority 30
    host priority 0, policy priority 10, Session id 0
  match oer learn list BUSINESS
  backoff 180 180 180
  *delay threshold 200
  holddown 90
  periodic 180
  *probe frequency 30
  number of jitter probe packets 100
  *mode route control
  *mode monitor both
  *mode select-exit good
  loss relative 10
  jitter threshold 20
  mos threshold 3.60 percent 30
  unreachable relative 50
  next-hop not set
  forwarding interface not set
  *resolve delay priority 1 variance 20
  *resolve utilization priority 3 variance 20
  *link-group WAN1 fallback WAN2
oer-map MYMAP 20
  sequence no. 8444249302630400, provider id 1, provider priority 30
    host priority 0, policy priority 20, Session id 0
  match oer learn list BE
  backoff 180 180 180
  delay relative 50
  holddown 90
  periodic 180
  *probe frequency 30
  number of jitter probe packets 100
  *mode route control
  *mode monitor both
  *mode select-exit good
  loss relative 10
  jitter threshold 20
  mos threshold 3.60 percent 30
  *unreachable threshold 500000
  next-hop not set
  forwarding interface not set
  *link-group WAN2 fallback WAN1

* Overrides Default Policy Setting
MC#
```

Verify Traffic Classes Statistics

As soon as you see:

```
MC#
*Sep  3 16:01:22.924: %OER_MC-5-NOTICE: Prefix Learning WRITING DATA
```

PfR:Solutions:EnterpriseData

```
*Sep  3 16:01:22.966: %OER_MC-5-NOTICE: Prefix Learning STARTED
MC#
```

You should be able to see the traffic classes on the Master Controller. You will need a few cycles before having all prefixes in INPOLICY state. You will start getting syslog messages showing route changes.

```
MC>
*Nov  3 12:11:05.484: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.0.0/24 af31 256, BR 10.4.5.
*Nov  3 12:11:05.484: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.2.0/24 af31 256, BR 10.4.5.
*Nov  3 12:11:05.484: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.3.0/24 af31 256, BR 10.4.5.
*Nov  3 12:11:05.484: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.1.0/24 af31 256, BR 10.4.5.
*Nov  3 12:11:05.484: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.4.0/24 defa 256, BR 10.4.5.
*Nov  3 12:11:05.485: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.5.0/24 defa 256, BR 10.4.5.
*Nov  3 12:11:05.485: %OER_MC-5-NOTICE: Route changed Appl Prefix 20.10.11.0/24 defa 256, BR 10.4.
*Nov  3 12:11:33.202: %OER_MC-5-NOTICE: Discovered Exit for Appl Prefix 30.1.1.0/24 defa 256, BR 1
*Nov  3 12:11:33.202: %OER_MC-5-NOTICE: Discovered Exit for Appl Prefix 30.1.3.0/24 defa 256, BR 1
*Nov  3 12:11:33.202: %OER_MC-5-NOTICE: Discovered Exit for Appl Prefix 30.1.2.0/24 defa 256, BR 1
MC>
```

Traffic Classes

On the Master Controller, you have all the Traffic Classes (in this case based on DSCP values) learnt as well as statistics:

Using show pfr master traffic-class:

```
MC#sh 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 (packets-per-million), 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
-----
30.1.0.0/24          N af31 256          N          N 0.0.0.0/0
                    INPOLICY          35          10.4.5.4 Et0/1          PBR
                    52          59          0          0          0          0          65          7
                    49          49          0          0          N          N          N          N

30.1.1.0/24          N defa 256          N          N 0.0.0.0/0
                    INPOLICY          @41          10.4.5.5 Et0/1          PBR
                    103          98          0          0          0          0          32          3
                    104          104          0          0          N          N          N          N

30.1.1.0/24          N af31 256          N          N 0.0.0.0/0
                    INPOLICY          30          10.4.5.4 Et0/1          PBR
                    52          59          0          0          0          0          27          3
```

PfR:Solutions:EnterpriseData

	52	52	0	0	N	N	N	N
30.1.2.0/24			N defa 256		N	N	0.0.0.0/0	
			INPOLICY	@32		10.4.5.5	Et0/1	PBR
	103	97	0	0	0	0	28	3
	101	105	0	0	N	N	N	N
30.1.2.0/24			N af31 256		N	N	0.0.0.0/0	
			INPOLICY	30		10.4.5.4	Et0/1	PBR
	52	57	0	0	0	0	31	3
	52	50	0	0	N	N	N	N
30.1.3.0/24			N defa 256		N	N	0.0.0.0/0	
			INPOLICY	@34		10.4.5.5	Et0/1	PBR
	103	98	0	0	0	0	27	3
	U	102	0	0	N	N	N	N
30.1.3.0/24			N af31 256		N	N	0.0.0.0/0	
			INPOLICY	94		10.4.5.4	Et0/1	PBR
	52	52	0	0	0	0	32	3
	51	51	0	0	N	N	N	N
30.1.4.0/24			N defa 256		N	N	0.0.0.0/0	
			INPOLICY	@44		10.4.5.5	Et0/1	PBR
	104	100	0	0	0	0	54	6
	U	102	0	0	N	N	N	N
30.1.5.0/24			N defa 256		N	N	0.0.0.0/0	
			INPOLICY	@49		10.4.5.5	Et0/1	PBR
	104	100	0	0	0	0	55	6
	U	103	0	0	N	N	N	N

MC#

- Note the learnt Traffic Classes.
- Note the AF31 setting under the DSCP field for Traffic Classes belonging to BUSINESS group.
- Note the default setting under the DSCP field for Traffic Classes belonging to BE group.
- Note that BUSINESS Traffic Classes are going through WAN1 (R4) and BE Traffic Classes are going through WAN2 (R5).

A closer look at the results

Let's have a look at a few entries belonging to the Traffic Classes that we have defined: 30.1.1.0/24 (BE) and 30.1.1.0/24 (BUSINESS).

Best Effort TC

From the command `?show oer master traffic-class`. Let's focus on the prefix 30.1.1.0/24 to understand the interesting values reported here:

PfR:Solutions:EnterpriseData

```

MC#sh 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 (packets-per-million), 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
-----

[snip]

30.1.1.0/24          N defa  256          N          N 0.0.0.0/0
                   INPOLICY          @41          10.4.5.5 Et0/1          PBR
                   103          98          0          0          0          0          32          3
                   104          104          0          0          N          N          N          N

[snip]

MC#

```

- **Line1: prefix**

- Line2: State of INPOLICY?this prefix is controlled by PfR and is currently inpolicy. PBR is used to enforce the path. Forwarding decisions based on other packet fields (such as TCP port numbers, DSCP field) cannot be done via the traditional routing table. For this reason, PfR will create a dynamic Policy Based Routing (PBR) policy and apply it to the PfR internal interfaces of the border routers.
- Line2: The ?at sign? (@) on the Time Remaining value means the prefix is being actively probed. The numerical value is a countdown timer indicating when this state will expire.
- Line2: R5 is the exit point (WAN2)

- **Line3: Passive results**

- Line3 (PasSDly, PasLDly): short-term and long-term passive delay, measured from the TCP Syn/Ack. As explained previously, TCP Syn/Ack are used to check the reachability of the prefix and to collect the delay and loss information. The delay is around 100 ms, which is the delay through WAN2 (BR R5).
- Line3 (PasSUn, PasLUn): short-term and long-term statistics for Unreachable.
- Line3 (PasSLos, PasLLos): short-term and long-term statistics for Loss.
- Line3 (EBw/IBw): the bandwidth for this prefix in egress and ingress direction.

- **Line4: Active results**

- Line4 (ActSDly, ActLDly): short-term and long-term active delay, measured by active probes that are automatically learned.
- Line4 (ActSUn, ActLUn): short-term and long-term Unreachable results.
- Line4 (ActSLos, ActLLoss): short-term and long-term Loss results.

Business TC

```
MC#sh 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 (packets-per-million), 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	Flags	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Protocol
			State	Time		CurrBR	CurrI/F	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos

```
[snip]
```

30.1.1.0/24			N af31	256		N	N 0.0.0.0/0	
			INPOLICY		30		10.4.5.4 Et0/1	PBR
	52	59	0	0	0	0	27	3
	52	52	0	0	N	N	N	N

```
[snip]
```

• Line1: Prefix

- Line2: State of INPOLICY?this prefix is controlled by PfR and is currently inpolicy. PBR is used to enforce the path. Forwarding decisions based on other packet fields (such as TCP port numbers, DSCP field) cannot be done via the traditional routing table. For this reason, PfR will create a dynamic Policy Based Routing (PBR) policy and apply it to the PfR internal interfaces of the border routers.
- Line2: R4 is the exit point (WAN1)

• Line3: Passive results

- Line3 (PasSDly, PasLDly): short-term and long-term passive delay, measured from the TCP Syn/Ack. As explained previously, TCP Syn/Ack are used to check the reachability of the prefix and to collect the delay and loss information. The delay is around 50 ms, which is the delay through WAN1 (BR R4).
- Line3 (PasSUn, PasLUn): short-term and long-term statistics for Unreachable.
- Line3 (PasSLos, PasLLos): short-term and long-term statistics for Loss.
- Line3 (EBw/IBw): the bandwidth for this prefix in egress and ingress direction.

- **Line4: Active results**
- Line4 (ActSDly, ActLDly): short-term and long-term active delay, measured by active probes that are automatically learned.
- Line4 (ActSUn, ActLUn): short-term and long-term Unreachable results.
- Line4 (ActSLos, ActLLos): short-term and long-term Loss results.

Active Probes

Because the monitoring mode is ?both?, active probes are used to help choosing the exit interfaces. When there is a route change, PfR will probe all exits to gather information. Remember that because we use monitoring mode both, only passive events can trigger OOP messages.

```
MC#sh 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 (packets-per-million), 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
PasSDly PasLDly PasSUn PasLUn PasSLos PasLLos CurrBR CurrI/F Protocol
ActSDly ActLDly ActSUn ActLUn ActSJit ActPMOS ActSLos ActLLos
-----

[snip]

30.1.2.0/24          N defa 256          N          N 0.0.0.0/0
                    INPOLICY          @32          10.4.5.5 Et0/1          PBR
                    103          97          0          0          0          0          28          3
                    101          105          0          0          N          N          N          N

[snip]

MC#
```

- **Line2:** The ?at sign? (@) on the Time Remaining value means the prefix is being actively probed. The numerical value is a countdown timer indicating when this state will expire.

In this configuration active probes are automatically calculated and generated, but active probes can be manually defined per TC in the pfr-map section.

```
MC#sh pfr master active-probes
OER Master Controller active-probes
Border = Border Router running this Probe
State = Un/Assigned to a Prefix
Prefix = Probe is assigned to this Prefix
Type = Probe Type
```

PfR:Solutions:EnterpriseData

Target = Target Address
TPort = Target Port
How = Was the probe Learned or Configured
N - Not applicable

The following Probes exist:

State	Prefix	Type	Target	TPort	How	Codec
Assigned	20.10.11.0/24	echo	20.10.11.11	N	Lrnd	N
Assigned	30.1.0.0/24	echo	30.1.0.11	N	Lrnd	N
Assigned	30.1.3.0/24	echo	30.1.3.11	N	Lrnd	N
Assigned	30.1.2.0/24	echo	30.1.2.11	N	Lrnd	N
Assigned	30.1.1.0/24	echo	30.1.1.11	N	Lrnd	N
Assigned	30.1.5.0/24	echo	30.1.5.11	N	Lrnd	N
Assigned	30.1.4.0/24	echo	30.1.4.11	N	Lrnd	N

The following Probes are running:

Border	State	Prefix	Type	Target	TPort
10.4.5.4	ACTIVE	20.10.11.0/24	echo	20.10.11.11	N
10.4.5.5	ACTIVE	20.10.11.0/24	echo	20.10.11.11	N

MC#

Verify Enforcement

Policy-Based Routing used

Forwarding decisions based on other packet fields (in this case DSCP field) cannot be done via the traditional routing table. For this reason, PfR will create a dynamic Policy Based Routing (PBR) policy and apply it to the PfR internal interfaces of the border routers. But for PBR to be successful, BRs have to be adjacent either through a direct connection or via a GRE tunnel. This requirement can be verified with the "show pfr master" command and looking for: PBR Requirements met

```
MC#sh pfr master
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 3949
Version: 3.0
Number of Border routers: 2
Number of Exits: 2
Number of monitored prefixes: 15 (max 5000)
Max prefixes: total 5000 learn 2500
Prefix count: total 15, learn 9, cfg 0
PBR Requirements met
Nbar Status: Inactive
```

<----- Check

Border	Status	UP/DOWN	AuthFail	Version
10.4.5.4	ACTIVE	UP 00:20:13	0	3.0
10.4.5.5	ACTIVE	UP 00:20:17	0	3.0

Global Settings:

```
max-range-utilization percent 0 recv 0
mode route metric bgp local-pref 5000
mode route metric static tag 5000
trace probe delay 1000
logging
exit holddown time 60 secs, time remaining 0
```

Default Policy Settings:

[snip]

BGP Routes by default

PfR being not active, the parent routes are BGP based on R2, R4 and R5. R4 is the preferred exit point as seen below:

```
R2#sh ip bgp
BGP table version is 27, local router ID is 10.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
r  i10.4.5.0/24     10.5.5.5          0      100     0  i
r> i                10.4.4.4          0      100     0  i
*> i20.10.11.0/24  100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i20.11.11.11/32 100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.0.0/24    100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.1.0/24    100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.2.0/24    100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.3.0/24    100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.4.0/24    100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i30.1.5.0/24    100.4.8.8         0      100     0 100 20  i
   Network          Next Hop          Metric LocPrf Weight Path
*  i                100.5.9.9         0      100     0 200 20  i
r> i100.4.8.0/24    10.4.4.4          0      100     0  i
r> i100.5.9.0/24    10.5.5.5          0      100     0  i
*> i100.8.10.0/24  100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
*> i100.9.10.0/24  100.4.8.8         0      100     0 100 20  i
*  i                100.5.9.9         0      100     0 200 20  i
R2#
```

Business Traffic Class

Traffic Class entries for BUSINESS are under PfR control, enforcement is with PBR and exit point is primary link WAN1. PBR is used because this TC is based on the DSCP field (DSCP af31 here) and not on prefix network only. Forwarding decisions based L3+ information cannot be done via the traditional routing table. For this reason, PfR will create a dynamic Policy Based Routing (PBR) policy and apply it to the PfR internal interfaces of the border routers. The PBR route-maps are created dynamically and do NOT appear in the configuration files of the Border Routers.

The exit point enforced is belonging to link-group WAN1 which is the primary group defined in the configuration with a fallback on WAN2 if traffic classes are OOP.

Best Effort Traffic Class

Traffic Class entries for BE are under PfR control, enforcement is with PBR and exit point is secondary link WAN2. For the same reason as explained before for TC BUSINESS, PBR is used because this TC is based on the DSCP field (DSCP 0 here) and not on prefix network only. The PBR route-maps are created dynamically and do NOT appear in the configuration files of the Border Routers.

The exit point enforced is belonging to link-group WAN2 which is the secondary group defined in the configuration.

Dynamic PBR on Border Router R4

PfR creates a dynamic route-map to enforce the path for BUSINESS and BE traffic. Depending on two dynamic ACLs, PBR enforce a next-hop to R5 (WAN2) or directly to R8 (WAN1).

```
R4#sh route-map dynamic
route-map OER_INTERNAL_RMAP, permit, sequence 0, identifier 3388997633
  Match clauses:
    ip address (access-lists): oer#1
  Set clauses:
    ip next-hop 10.4.5.5
    interface Ethernet0/0
  Policy routing matches: 51425 packets, 35300582 bytes
route-map OER_INTERNAL_RMAP, permit, sequence 1, identifier 1979711490
  Match clauses:
    ip address (access-lists): oer#2
  Set clauses:
    ip next-hop 100.4.8.8
    interface Ethernet0/1
  Policy routing matches: 30042 packets, 20665961 bytes
Current active dynamic routemaps = 1
R4#
```

This route-map is applied on the ingress interface of R4:

```
R4#sh ip policy
Interface      Route map
Ethernet0/0    OER_INTERNAL_RMAP (Dynamic)
R4#
```

And the corresponding access-list are dynamically created by PfR.

We clearly see that:

- packets matching ACL oer#1 belong to TC BE (dscp default) and PBR will set ip next-hop 10.4.5.5 (R5, and therefore WAN2).
- packets matching ACL oer#2 belong to TC BUSINESS (dscp af31) and PBR will set ip next-hop 100.4.8.8 (R8, and therefore WAN1).

```
R4#sh ip access-list dynamic
Extended IP access list oer#1
 1048575 permit ip any 30.1.5.0 0.0.0.255 dscp default (11192 matches)
 2097151 permit ip any 30.1.4.0 0.0.0.255 dscp default (11193 matches)
 4194303 permit ip any 30.1.3.0 0.0.0.255 dscp default (7308 matches)
 8388607 permit ip any 30.1.1.0 0.0.0.255 dscp default (7215 matches)
 16777215 permit ip any 30.1.2.0 0.0.0.255 dscp default (6919 matches)
Extended IP access list oer#2
 134217727 permit ip any 30.1.0.0 0.0.0.255 dscp af31 (13334 matches)
 268435455 permit ip any 30.1.1.0 0.0.0.255 dscp af31 (5796 matches)
 536870911 permit ip any 30.1.2.0 0.0.0.255 dscp af31 (6209 matches)
 1073741823 permit ip any 30.1.3.0 0.0.0.255 dscp af31 (6468 matches)
R4#
```

Dynamic PBR on Border Router R5

PfR creates a dynamic route-map to enforce the path for BUSINESS and BE traffic. Depending on two dynamic ACLs, PBR enforce a next-hop to R9 (100.5.9.9, WAN2) or to R4 (10.4.5.4, WAN1).

```
R5#sh route-map dynamic
route-map OER_INTERNAL_RMAP, permit, sequence 0, identifier 788529153
  Match clauses:
    ip address (access-lists): oer#1
  Set clauses:
    ip next-hop 100.5.9.9
    interface Ethernet0/1
  Policy routing matches: 59776 packets, 41051578 bytes
route-map OER_INTERNAL_RMAP, permit, sequence 1, identifier 4194304002
  Match clauses:
    ip address (access-lists): oer#2
  Set clauses:
    ip next-hop 10.4.5.4
    interface Ethernet0/0
  Policy routing matches: 0 packets, 0 bytes
Current active dynamic routemaps = 1
R5#
```

This route-map is applied on the ingress interface of R5:

```
R5#sh ip policy
Interface      Route map
Ethernet0/0    OER_INTERNAL_RMAP (Dynamic)
R5#
```

And the corresponding access-list are dynamically created by PfR.

We clearly see that:

- packets matching ACL oer#1 belong to TC BE (dscp default) and PBR will set ip next-hop 100.5.9.9 (R9, and therefore WAN2).
- packets matching ACL oer#2 belong to TC BUSINESS (dscp af31) and PBR will set ip next-hop 10.4.5.4 (R4, and therefore WAN1).

```
R5#sh ip access-lists dynamic
Extended IP access list oer#1
 1048575 permit ip any 30.1.5.0 0.0.0.255 dscp default (13182 matches)
 2097151 permit ip any 30.1.4.0 0.0.0.255 dscp default (13183 matches)
 4194303 permit ip any 30.1.3.0 0.0.0.255 dscp default (8328 matches)
 8388607 permit ip any 30.1.1.0 0.0.0.255 dscp default (8320 matches)
 16777215 permit ip any 30.1.2.0 0.0.0.255 dscp default (8035 matches)
Extended IP access list oer#2
 134217727 permit ip any 30.1.0.0 0.0.0.255 dscp af31
 268435455 permit ip any 30.1.1.0 0.0.0.255 dscp af31
 536870911 permit ip any 30.1.2.0 0.0.0.255 dscp af31
 1073741823 permit ip any 30.1.3.0 0.0.0.255 dscp af31
R5#
```

Add Delay on SP1

?normal? delay across SP1

The Traffic Class BRANCH_BUSINESS has a delay resolver configured with an absolute threshold of 200 ms. In this section, we'll take the WAN1 path (which is the primary one) and add delay to it and observe what PfR does.

For TCP sessions, PfR is able to detect the delay along the path.

PfR maintain delay statistics for the short-term delay (5 min) and the long-term delay (60 min).

Remember that there are two ways to specify what is flagged as an OOP condition in PfR:

- Relative: Relative implies the use of short/long term statistics. You specify a percentage and that is used to gauge if a TC is OOP.
- Absolute: You now have a threshold that is compared only to the short-term delay.

Before adding delay, let's verify the RTT from R4 across the WAN1 link:

```
R4#ping 100.4.8.8

Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 100.4.8.8, timeout is 2 seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 51/53/54 ms
R4#
```

Adding a delay of 300ms on SP1

Now, we are adding 300 ms delay on WAN1. Again verify on R4:

```
R4#ping 100.4.8.8
```

```
Type escape sequence to abort.
```

```
Sending 5, 100-byte ICMP Echos to 100.4.8.8, timeout is 2 seconds:
```

```
!!!!!
```

```
Success rate is 100 percent (5/5), round-trip min/avg/max = 301/303/306 ms
```

```
R4#
```

The delay threshold being defined in an absolute mode, is therefore compared to the short-term delay. By issuing "sh pfr master traffic-class", you can notice that the short-term delay (PasSDly) is increasing for af31 sessions. The short-term delay is the average of the last 5 minutes of statistics.

```
MC#sh pfr master tra
```

```
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 (packets-per-million), 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	Flags	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Protocol
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos
30.1.0.0/24			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	82		10.4.5.4	Et0/1	PBR
	197	65	0	0	0	0	31	3
	302	86	0	0	N	N	N	N
30.1.1.0/24			N defa	256		N	N 0.0.0.0/0	
			INPOLICY	32		10.4.5.5	Et0/1	PBR
	103	102	0	0	0	0	20	2
	102	102	0	0	N	N	N	N
30.1.1.0/24			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	69		10.4.5.4	Et0/1	PBR
	213	68	0	0	0	0	16	2
	301	87	0	0	N	N	N	N
30.1.2.0/24			N defa	256		N	N 0.0.0.0/0	
			INPOLICY	18		10.4.5.5	Et0/1	PBR
	103	101	0	0	0	0	19	2
	104	102	0	0	N	N	N	N
30.1.2.0/24			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	91		10.4.5.4	Et0/1	PBR
	241	69	0	0	0	0	16	2
	305	87	0	0	N	N	N	N
30.1.3.0/24			N defa	256		N	N 0.0.0.0/0	
			INPOLICY	@47		10.4.5.5	Et0/1	PBR

PfR:Solutions:EnterpriseData

104	102	0	0	0	0	21	2
105	102	0	0	N	N	N	N
30.1.3.0/24		N af31 256		N		N 0.0.0.0/0	
		INPOLICY	90		10.4.5.4	Et0/1	PBR
222	67	0	0	0	0	16	2
302	87	0	0	N	N	N	N
30.1.4.0/24		N defa 256		N		N 0.0.0.0/0	
		INPOLICY	67		10.4.5.5	Et0/1	PBR
104	101	0	0	0	0	53	6
103	102	0	0	N	N	N	N
30.1.5.0/24		N defa 256		N		N 0.0.0.0/0	
		INPOLICY	80		10.4.5.5	Et0/1	PBR
104	101	0	0	0	0	53	6
102	102	0	0	N	N	N	N

MC#

Move TC to new exit ? HOLDDOWN State

After a few cycles we can notice that the Passive Short Delay is now above 200 ms. The absolute delay threshold is compared against the passive short-term delay value (PasSDly) and therefore trigger an OOP message for traffic class BRANCH_BUSINESS (packets with DSCP AF31) going over the WAN1 link. Because we are running in monitor mode both, only passive results can trigger OOP messages.

MC#

```
*Nov 3 14:12:12.205: %OER_MC-5-NOTICE: Passive ABS Delay OOP Appl Prefix 30.1.1.0/24 af31 256, de
*Nov 3 14:12:12.205: %OER_MC-5-NOTICE: Passive ABS Delay OOP Appl Prefix 30.1.2.0/24 af31 256, de
*Nov 3 14:12:12.205: %OER_MC-5-NOTICE: Passive ABS Delay OOP Appl Prefix 30.1.3.0/24 af31 256, de
*Nov 3 14:13:19.205: %OER_MC-5-NOTICE: Passive ABS Delay OOP Appl Prefix 30.1.0.0/24 af31 256, de
```

MC#

When OOP condition is detected, PfR tries to find an alternate path. When PfR makes a route change to a specific TC, this TC will move to the HOLDDOWN state to avoid erratic behavior. The timer associated is the holddown timer defined in the pfr master configuration. A syslog message indicates that a new exit is found. The primary path being OOP, PfR moves BUSINESS sessions to the fallback path.

MC#

```
*Nov 3 14:13:09.554: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.1.0/24 af31 256, BR 10.4.5.
*Nov 3 14:13:09.555: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.2.0/24 af31 256, BR 10.4.5.
*Nov 3 14:13:09.556: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.3.0/24 af31 256, BR 10.4.5.
*Nov 3 14:14:16.744: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.0.0/24 af31 256, BR 10.4.5.
```

MC#

BUSINESS Traffic Class sessions are now in HOLDDOWN State:

MC#sh pfr master tra

PfR:Solutions:EnterpriseData

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 (packets-per-million), 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		Protocol
	Flags	State					Time	CurrBR	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSJos	PasLJos	ActPMOS	ActSJos	ActLJos
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit				
30.1.0.0/24			N af31	256		N	N 0.0.0.0/0		
			HOLDDOWN		66		10.4.5.5 Et0/1		PBR
	U	U		0	0	0	0	20	2
	U	U		0	0	N	N	N	N
30.1.1.0/24			N defa	256		N	N 0.0.0.0/0		
			INPOLICY		62		10.4.5.5 Et0/1		PBR
	103	102		0	0	0	0	15	1
	101	102		0	0	N	N	N	N
30.1.1.0/24			N af31	256		N	N 0.0.0.0/0		
			HOLDDOWN		5		10.4.5.5 Et0/1		PBR
	104	104		0	0	0	0	13	1
	U	U		0	0	N	N	N	N
30.1.2.0/24			N defa	256		N	N 0.0.0.0/0		
			INPOLICY		49		10.4.5.5 Et0/1		PBR
	103	102		0	0	0	0	18	1
	101	101		0	0	N	N	N	N
30.1.2.0/24			N af31	256		N	N 0.0.0.0/0		
			HOLDDOWN		4		10.4.5.5 Et0/1		PBR
	104	104		0	0	0	0	16	1
	U	U		0	0	N	N	N	N
30.1.3.0/24			N defa	256		N	N 0.0.0.0/0		
			INPOLICY		31		10.4.5.5 Et0/1		PBR
	103	102		0	0	0	0	17	2
	106	102		0	0	N	N	N	N
30.1.3.0/24			N af31	256		N	N 0.0.0.0/0		
			HOLDDOWN		1		10.4.5.5 Et0/1		PBR
	160	160		0	0	0	0	12	2
	U	U		0	0	N	N	N	N
30.1.4.0/24			N defa	256		N	N 0.0.0.0/0		
			INPOLICY		100		10.4.5.5 Et0/1		PBR
	104	101		0	0	0	0	56	6
	105	102		0	0	N	N	N	N
30.1.5.0/24			N defa	256		N	N 0.0.0.0/0		
			INPOLICY		116		10.4.5.5 Et0/1		PBR
	104	102		0	0	0	0	53	6
	100	102		0	0	N	N	N	N

MC#

New exit ? INPOLICY State

After the holdown timer expires, the traffic classes go in the INPOLICY state again but on the fallback link (WAN2, through R5).

```
MC#sh pfr master tra
```

```
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 (packets-per-million), 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		Protocol
	Flags						CurrBR	CurrI/F	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw	
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos	
30.1.0.0/24		N	af31	256	N	N	0.0.0.0/0		
		INPOLICY			58	10.4.5.5	Et0/1		PBR
	104	104	0	0	0	0	34		3
	U	U	0	0	N	N	N		N
30.1.1.0/24		N	defa	256	N	N	0.0.0.0/0		
		INPOLICY			107	10.4.5.5	Et0/1		PBR
	104	102	0	0	0	0	22		2
	104	102	0	0	N	N	N		N
30.1.1.0/24		N	af31	256	N	N	0.0.0.0/0		
		INPOLICY			@50	10.4.5.5	Et0/1		PBR
	104	104	0	0	0	0	27		3
	U	U	0	0	N	N	N		N
30.1.2.0/24		N	defa	256	N	N	0.0.0.0/0		
		INPOLICY			93	10.4.5.5	Et0/1		PBR
	104	102	0	0	0	0	25		2
	102	102	0	0	N	N	N		N
30.1.2.0/24		N	af31	256	N	N	0.0.0.0/0		
		INPOLICY			2	10.4.5.5	Et0/1		PBR
	104	104	0	0	0	0	24		3
	U	U	0	0	N	N	N		N
30.1.3.0/24		N	defa	256	N	N	0.0.0.0/0		
		INPOLICY			76	10.4.5.5	Et0/1		PBR
	104	102	0	0	0	0	17		2
	103	102	0	0	N	N	N		N
30.1.3.0/24		N	af31	256	N	N	0.0.0.0/0		
		INPOLICY			@43	10.4.5.5	Et0/1		PBR
	137	137	0	0	0	0	31		2
	U	U	0	0	N	N	N		N
30.1.4.0/24		N	defa	256	N	N	0.0.0.0/0		
		INPOLICY			@19	10.4.5.5	Et0/1		PBR
	104	102	0	0	0	0	57		6

New exit ? INPOLICY State

PfR:Solutions:EnterpriseData

106	102	0	0	N	N	N	N
30.1.5.0/24		N defa 256		N		N 0.0.0.0/0	
		INPOLICY	@35		10.4.5.5	Et0/1	PBR
104	102	0	0	0	0	54	6
U	102	0	0	N	N	N	N

MC#

Back to Normal

We reset the delay on WAN1. Delay is now back to 50 ms. PfR will gradually move TCs belonging to BUSINESS group to the primary path WAN1.

```
*Nov 3 14:29:20.132: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.0.0/24 af31 256, BR 10.4.5.
*Nov 3 14:31:14.356: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.3.0/24 af31 256, BR 10.4.5.
```

Eventually, the system will settle into a state where the majority of the traffic is INPOLICY, BUSINESS Traffic Classes going over the primary path WAN1, and BE Traffic Classes going over the secondary path WAN2.

Conclusion

Performance Routing is an advanced technology that allows multiple deployments, one of them being the one described in this document. The configuration is straightforward and very efficient. Based on that, a customer can evolve to a more complex Performance Routing solution or can fine-tune the Traffic Class definition.