

Cisco Performance Routing (PfR) Solution Guides

PfR Internet Presence - Inbound Load Balancing

Navigation

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

Contents

- [1 Internet Presence and Inbound Load Balancing](#)
- [2 PfR Solution Used](#)
 - ◆ [2.1 Overview](#)
 - ◆ [2.2 Learning phase](#)
 - ◆ [2.3 Measuring phase](#)
 - ◆ [2.4 Optimization Policies](#)
 - ◆ [2.5 Enforcement](#)
- [3 PfR Network Topology Used](#)
- [4 Flexible Netflow](#)
- [5 Checking Statistics and Flows](#)
- [6 Display Routing Table \(Central Site\)](#)
- [7 Display Routing Table \(Remote Sites\)](#)
- [8 Dynamic Configuration](#)
 - ◆ [8.1 PfR Configuration](#)
 - ◆ [8.2 Master Controller Verification](#)
 - ◆ [8.3 Display Outside Prefixes](#)
 - ◆ [8.4 Display Inside Prefixes](#)
- [9 Static Configuration and PfR Map](#)
 - ◆ [9.1 PfR Configuration](#)
 - ◆ [9.2 Master Controller Verification](#)
 - ◆ [9.3 Display Inside Traffic Classes](#)
- [10 Dynamic Configuration and PfR Map](#)
 - ◆ [10.1 PfR Configuration](#)

- ◆ 10.2 Master Controller Verification
- ◆ 10.3 Display Inside Prefixes
- 11 Verify Load balancing and Enforcement
 - ◆ 11.1 Bandwidth used on exit links
 - ◆ 11.2 Verifying Enforcement
- 12 Conclusion

Internet Presence and Inbound Load Balancing

This use case is a common deployment scenario; optimization of large numbers of client devices sourced from several ISP connections. In terms of megabits per second, the bulk of the user traffic is from servers to clients. PfR, therefore, is configured and addresses the path selection from servers to clients over two or more links to typically multiple ISPs.

But there is sometimes a need to influence the ingress traffic using BGP. This solution guide will address the need of inbound optimization, ie from clients to servers.

The PfR BGP inbound optimization feature introduced the ability to influence inbound traffic and therefore the ability to support inside prefixes.

A network advertises reachability of its inside prefixes to the Internet using eBGP advertisements to its ISPs. If the same prefix is advertised to more than one ISP, then the network is multihoming. PfR BGP inbound optimization works best with multihomed networks, but it can also be used with a network that has multiple connections to the same ISP. Using BGP, PfR can select inside prefixes to support best entrance selection for traffic that originates from prefixes outside an autonomous system destined for prefixes inside the autonomous system. Company networks advertise the inside prefixes over the Internet using an Internet service provider (ISP) and receive advertisements for outside prefixes from an ISP.

PfR Solution Used

Overview

PfR Inbound optimization only supports passive monitoring. The load-balancing used here is taking place between external interfaces for inbound traffic in addition to outbound traffic. This deployment does not use active probes. For PfR to verify reachability for a destination network prefix, TCP traffic must be observed on more than one exit interface so PfR has more than one exit with validated reachability to the target network prefix. For Internet Edge based designs where HTTP traffic is predominant, this is usually not an issue.

BGP is used between the site and all upstream ISPs.

- For outbound traffic optimization, Internet Service Providers can send:

Default routing or
 Partial routes or
 Full routes

(In all the above 3 cases, the PfR configuration will remain the same).

Note: the outbound traffic optimization is described in another solution guide, we will therefore focus on the PfR configuration for inbound optimization.

See [Internet - Outbound Load Balancing](#)

- For inbound traffic optimization, the site has to send the exact inside prefixes that have to be controlled by PfR.

Learning phase

BGP inbound optimization provides the ability to manually configure (based on prefix-list) or automatically learn inside prefixes.

the inside prefix learning that happens for inbound optimization is by looking at the BGP tables on the border routers. The Master Controller asks BGP on the BRs for all networks that were advertised to eBGP peers over PfR external interfaces. Once the MC receives all the prefixes from the BRs, it just enters them into the monitoring database and start optimizing them. Note that PfR only looks at the prefixes that are originated in the local AS, therefore transit prefixes that were advertised by the BRs to eBGP peers, if any, will not be learnt.

The maximum number of inside prefixes that can be learned in a monitoring period is 30.

Measuring phase

Passive			Active		
Reachability	Delay	Loss	Reachability	Delay	Loss
Egress BW	Ingress BW		Jitter	MOS	
<ul style="list-style-type: none"> ▪ PfR Netflow Monitoring ▪ Flows Need not be symmetrical 			<ul style="list-style-type: none"> ▪ PfR enables IP SLA feature ▪ Probes sourced from BR ▪ ICMP probes learned or configured ▪ TCP, UDP, JITTER need ip sla responder 		

Monitor mode passive is the only mode supported for inbound optimization, which means PfR only supports passive data collection based on Netflow monitoring to optimize inbound traffic. Metrics available are unreachable, loss, delay, ingress and egress bandwidth.

Optimization Policies

Link policies defining traffic load or range performance characteristics can be applied against PfR-managed entrance links.

range: traffic should be load-balanced so that the average bandwidth range between external interfaces should be maintain within X%.

Configuration sample:

```
max range receive percent 5
resolve range priority 3
no resolve delay
no resolve utilization
```

Enforcement

To enforce an entrance link selection, PfR offers the following methods:

- BGP Autonomous System Number Prepend

When an entrance link goes out-of-policy (OOP) due to delay, or in images prior to Cisco IOS Releases 15.2(1)T1 and 15.1(2)S, and PfR selects a best entrance for an inside prefix, extra autonomous system hops are prepended one at a time (up to a maximum of six) to the inside prefix BGP advertisement over the other entrances.

In Cisco IOS Releases 15.2(1)T1, 15.1(2)S, and later releases, when an entrance link goes out-of-policy (OOP) due to unreachable or loss reasons, and PfR selects a best entrance for an inside prefix, six extra autonomous system hops are prepended immediately to the inside prefix BGP advertisement over the other entrances. The extra autonomous system hops on the other entrances increase the probability that the best entrance will be used for the inside prefix and allows PfR to quickly move the traffic away from the old entrance link.

This is the default method PfR uses to control an inside prefix, and no user configuration is required.

- BGP Autonomous System Number Community Prepend

When an entrance link goes out-of-policy (OOP) due to delay, or in images prior to Cisco IOS Releases 15.2(1)T1 and 15.1(2)S, and PfR selects a best entrance for an inside prefix, a BGP prepend community is attached one at a time (up to a maximum of six) to the inside prefix BGP advertisement from the network to another autonomous system such as an ISP.

In Cisco IOS Releases 15.2(1)T1, 15.1(2)S, and later releases, when an entrance link goes out-of-policy (OOP) due to unreachable or loss reasons, and PfR selects a best entrance for an inside prefix, six BGP prepend communities are attached to the inside prefix BGP advertisement. The BGP prepend community will increase the number of autonomous system hops in the advertisement of the inside prefix from the ISP to its peers.

PfR:Solutions:InternetInboundLoadBalancing

Autonomous system prepend BGP community is the preferred method to be used for PfR BGP inbound optimization because there is no risk of the local ISP filtering the extra autonomous system hops. There are some issues, for example, not all ISPs support the BGP prepend community, ISP policies may ignore or modify the autonomous system hops, and a transit ISP may filter the autonomous system path. If you use this method of inbound optimization and a change is made to an autonomous system, you must issue an outbound reconfiguration using the `clear ip bgp` command.

The prepend community has to be configured per Border Router and per external interfaces because the community value is specific per Service Provider.

Configuration Sample:

```
!  
pfr master  
border 10.4.5.6 key-chain pfr  
  interface Ethernet0/1 external  
    downgrade bgp community aa:nn (community number in aa:nn format)  
  interface Ethernet0/0 internal  
!
```

In this solution guide, PfR will enforce the path by using BGP AS-PATH prepend for controlled prefixes. PfR will not try to control an inside prefix unless there is an exact match in the BGP routing information base (RIB) because PfR does not advertise a new prefix to the Internet.

PfR Network Topology Used

The central site has three Border Routers, connected to three separate Service Providers using eBGP. R2, R4, R5 and R6 are iBGP peers. For an Internet Presence solution, it may be recommended to have a dedicated Master Controller given the possible high number of prefixes that have to be optimized and managed.

- R2, R4, R5 and R6 are iBGP peers in AS 100
- R3 is the Master Controller
- R4, R5 and R6 are the Border Routers
- Traffic Simulator tool is used between R1 and R11, R12 to emulate traffic
- R1, R11 and R12 are traffic generators (to send/receive http, ssh, etc.).

And then apply the Flexible NetFlow Monitor on the Border Routers as well as R2:

```
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).

Here is the output on R2 which sees all flows:

```
R2#sh flow monitor MYMONITOR cache format table
```

```
R2#shflow
Cache type:                Normal
Cache size:                4096
Current entries:          208
High Watermark:           208

Flows added:              208
Flows aged:               0
- Active timeout          ( 1800 secs) 0
- Inactive timeout        (   15 secs) 0
- 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
10.10.3.1	30.30.16.11	80	2037	Et0/1	6	EtC
10.10.1.1	30.30.20.11	7000	7001	Et0/1	6	EtC
10.10.1.1	30.30.8.11	7000	7005	Et0/1	6	EtC
10.10.2.1	30.30.9.11	25	1028	Et0/1	6	EtC
10.10.1.1	30.30.17.11	7000	7004	Et0/1	6	EtC
10.10.2.1	20.20.8.12	25	1016	Et0/1	6	EtC
10.10.3.1	30.30.18.11	80	2038	Et0/1	6	EtC
10.10.4.1	20.20.20.12	80	2045	Et0/1	6	EtC
10.10.1.1	20.20.12.12	7000	7005	Et0/1	6	EtC
10.10.4.1	20.20.6.12	80	2001	Et0/1	6	EtC
10.10.1.1	30.30.5.11	7000	7008	Et0/1	6	EtC
10.10.2.1	20.20.5.12	25	1029	Et0/1	6	EtC
10.10.1.1	30.30.11.11	7000	7006	Et0/1	6	EtC
10.10.2.1	30.30.4.11	25	1027	Et0/1	6	EtC
10.10.1.1	20.20.4.12	7000	7003	Et0/1	6	EtC
10.10.1.1	20.20.15.12	7000	7009	Et0/1	6	EtC
10.10.4.1	20.20.19.12	80	2048	Et0/1	6	EtC

[SNIP]

Note: 10.10.0.0/16 is the prefix allocated to the servers in our example above.

Display Routing Table (Central Site)

Let's have a look at the routing table before applying a PfR configuration. Central site has prefixes 10.10.0.0/16 in Autonomous System 100, remote sites have prefixes 20.20.0.0/16 in Autonomous System 200 and 30.30.0.0/16 in Autonomous System 300. For clarity, only interesting part matching the destination prefixes of the routing tables are displayed. The servers subnets (inside prefixes) are 10.10.1.0/24, 10.10.2.0/24, 10.10.3.0/24 and 10.10.4.0/24. These prefixes must be in the BGP table for inbound optimization to work.

On the Border Router R4:

```
R4#sh bgp
BGP table version is 1476, 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

   Network          Next Hop          Metric LocPrf Weight Path
* i 10.10.1.0/24    10.4.5.2           21     100      0 i <-- INSIDE
* i                  10.4.5.2           21     100      0 i
*>                  10.4.5.2           21             32768 i
* i 10.10.2.0/24    10.4.5.2           21     100      0 i <-- INSIDE
* i                  10.4.5.2           21     100      0 i
*>                  10.4.5.2           21             32768 i
* i 10.10.3.0/24    10.4.5.2           21     100      0 i <-- INSIDE
* i                  10.4.5.2           21     100      0 i
*>                  10.4.5.2           21             32768 i
* i 10.10.4.0/24    10.4.5.2           21     100      0 i <-- INSIDE
* i                  10.4.5.2           21     100      0 i
*>                  10.4.5.2           21             32768 i
*>i 20.20.0.0/16    100.6.83.1         0       200      0 300 20 i <-- REMOTE AS200
*                    100.4.81.1         50       0 100 20 i
*>i 30.30.0.0/16    100.6.83.1         0       200      0 300 30 i <-- REMOTE AS300
*                    100.4.81.1         50       0 100 30 i
```

[SNIP]

R4#

On the Border Router R5:

```
R5#sh bgp
BGP table version is 1442, local router ID is 10.5.5.5
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
```


PfR:Solutions:InternetInboundLoadBalancing

```

      Network          Next Hop          Metric LocPrf Weight Path
* i 10.10.1.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.2.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.3.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.4.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
*>i 20.20.0.0/16      100.6.83.1           0       200      0 300 20 i <-- REMOTE AS200
*                   100.5.82.1           100     0 200 20 i
*>i 30.30.0.0/16      100.6.83.1           0       200      0 300 30 i <-- REMOTE AS300
*                   100.5.82.1           100     0 200 30 i

```

[SNIP]

R5#

On the Border Router R6:

R6#sh bgp

```

BGP table version is 1436, local router ID is 10.6.6.6
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

```

```

      Network          Next Hop          Metric LocPrf Weight Path
* i 10.10.1.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.2.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.3.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
* i 10.10.4.0/24      10.4.5.2             21     100      0 i <-- INSIDE
* i                   10.4.5.2             21     100      0 i
*>                   10.4.5.2             21                   32768 i
*> 20.20.0.0/16      100.6.83.1           200     0 300 20 i <----- HERE
*> 30.30.0.0/16      100.6.83.1           200     0 300 30 i <----- HERE

```

[SNIP]

R6#

Display Routing Table (Remote Sites)

Let's look at the AS PATH for the server subnets on the remote sites.

On the Remote Router R9:

```
R9#sh bgp
BGP table version is 248, local router ID is 20.9.9.9
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
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.10.1.0/24	20.9.81.1			0	100 10 i
*	20.9.83.1		100	0	300 10 i
*	20.9.82.1			0	200 10 i
*> 10.10.2.0/24	20.9.81.1			0	100 10 i
*	20.9.83.1		100	0	300 10 i
*	20.9.82.1			0	200 10 i
*> 10.10.3.0/24	20.9.81.1			0	100 10 i
*	20.9.83.1		100	0	300 10 i
*	20.9.82.1			0	200 10 i
*> 10.10.4.0/24	20.9.81.1			0	100 10 i
*	20.9.83.1		100	0	300 10 i
*	20.9.82.1			0	200 10 i

[SNIP]

R9#

On the Remote Router R10:

```
R10#sh bgp
BGP table version is 256, local router ID is 30.10.10.10
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
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 10.10.1.0/24	30.10.81.1			0	100 10 i
*	30.10.83.1		100	0	300 10 i
*	30.10.82.1			0	200 10 i
*> 10.10.2.0/24	30.10.81.1			0	100 10 i
*	30.10.83.1		100	0	300 10 i
*	30.10.82.1			0	200 10 i
*> 10.10.3.0/24	30.10.81.1			0	100 10 i
*	30.10.83.1		100	0	300 10 i
*	30.10.82.1			0	200 10 i
*> 10.10.4.0/24	30.10.81.1			0	100 10 i
*	30.10.83.1		100	0	300 10 i
*	30.10.82.1			0	200 10 i

[SNIP]

R10#

Dynamic Configuration

PfR Configuration

This configuration is the easiest one where dynamic learning is enabled for both outside and inside prefixes.

The inside prefix learning that happens for inbound optimization is by looking at the BGP tables on the border routers. It is not done through NetFlow top talkers. We just ask BGP on the BRs for all networks that were advertised to eBGP peers over PfR external interfaces. Once we receive all the prefixes from the BRs, we just enter them into the monitoring database and start optimizing them (note that we ask for prefixes that originated in our AS only, so transit prefixes that were advertised by our BRs to eBGP peers, if any, will not be learnt).

```
pfr master
max-range-utilization percent 8
logging
!
border 10.4.5.6 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.4.5.5 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.4.5.4 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
learn
inside bgp
traffic-class filter access-list ALLOW_REMOTE
expire after time 300
!
max prefix total 10000 learn 10000
max range receive percent 5
mode monitor passive
resolve range priority 3
no resolve delay
no resolve utilization
!
```

Note: the traffic class filter is optional and just there to ease the reading of the outputs.

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.

Goal:

- Verify that the MC is active
- Verify that the learning process is enabled on the Master Controller for the outgoing traffic
- Verify that the inside prefixes are correctly learnt from the BGP tables on the Border Routers
- Display the policy settings.

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

You may have to wait a few cycles for the MC to be able to learn a few prefixes and then issue the command:

```
MC#sh pfr master
OER state: ENABLED and ACTIVE
  Conn Status: SUCCESS, PORT: 3949
  Version: 3.3
  Number of Border routers: 3
  Number of Exits: 3
  Number of monitored prefixes: 52 (max 10000)
  Max prefixes: total 10000 learn 10000
  Prefix count: total 52, learn 52, cfg 0
  PBR Requirements met
  Nbar Status: Inactive
  Auto Tunnel Mode: On

Border          Status   UP/DOWN          AuthFail  Version  DOWN Reason
10.4.5.4        ACTIVE  UP               00:39:36    0   3.3
10.4.5.5        ACTIVE  UP               00:39:34    0   3.3
10.4.5.6        ACTIVE  UP               00:39:34    0   3.3

Global Settings:
  max-range-utilization percent 8 recv 5
  rsvp post-dial-delay 0 signaling-retries 1
  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 20

Default Policy Settings:
  backoff 90 900 90
  delay relative 50
  holddown 90
  periodic 0
  probe frequency 56
  number of jitter probe packets 100
  mode route control
  mode monitor passive
  loss relative 10
  jitter threshold 20
  mos threshold 3.60 percent 30
  unreachable relative 50
```

PfR:Solutions:InternetInboundLoadBalancing

```
resolve range priority 3 variance 0
```

Learn Settings:

```
current state : STARTED
time remaining in current state : 92 seconds
throughput
no delay
inside bgp
traffic-class filter access-list ALLOW_REMOTE
monitor-period 1
periodic-interval 0
aggregation-type prefix-length 24
prefixes 100 appls 100
expire after time 300
```

MC#

What to check:

- All Border Routers are up and running
- Number of monitored prefixes: 52 (max 10000) - 52 prefixes. if there is more than 10000 prefixes, the rest will be under the routing control.
- Prefixes: total 10000 learn 10000 - The max number of monitored prefixes, the max number of learnt prefixes.
- Global Settings: max-range-utilization is 8 (percentage of bandwidth difference between all exit interfaces), used for outgoing traffic
- Learn Settings: prefixes 100 appls 100 - 100 prefixes max automatically learnt per cycle (this could be changed via configuration)
- All default policy settings are displayed
- Learn is started (current state : STARTED)

Display Outside Prefixes

On the Master Controller, you have all the Traffic Classes (in this case prefixes) learnt as well as statistics. We are not going into details as this is explained in the Internet Load Balancing Solution Guide.

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

Display Outside Prefixes

PfR:Solutions:InternetInboundLoadBalancing

```

-----
20.20.0.0/24          N    N    N          N          N N
                    INPOLICY          0          10.4.5.6 Et0/1          BGP
                    152    153          0          0          278    362          36          4
                    N    N    N          N          N          N
20.20.8.0/24         N    N    N          N          N N
                    INPOLICY          0          10.4.5.6 Et0/1          BGP
                    151    154          0          0          469    352          38          4
                    N    N    N          N          N          N
20.20.16.0/24        N    N    N          N          N N
                    INPOLICY          0          10.4.5.6 Et0/1          BGP
                    152    155          0          0          242    353          35          4
                    N    N    N          N          N          N
30.30.0.0/24         N    N    N          N          N N
                    INPOLICY          0          10.4.5.4 Et0/1          BGP
                    52    52          0          0          394    273          58          7
                    N    N    N          N          N          N
30.30.8.0/24         N    N    N          N          N N
                    INPOLICY          0          10.4.5.5 Et0/1          BGP
                    104    105          0          0          809    474          48          5
                    N    N    N          N          N          N
30.30.16.0/24        N    N    N          N          N N
                    INPOLICY          0          10.4.5.5 Et0/1          BGP
                    104    104          0          0          171    467          51          5
                    N    N    N          N          N          N
20.20.1.0/24         N    N    N          N          N N
                    INPOLICY          0          10.4.5.6 Et0/1          BGP
                    151    152          0          0          259    336          36          3
                    N    N    N          N          N          N
20.20.9.0/24         N    N    N          N          N N
                    INPOLICY          0          10.4.5.6 Et0/1          BGP
                    152    153          0          0          719    380          35          4
                    N    N    N          N          N          N
20.20.17.0/24        N    N    N          N          N N
                    INPOLICY          0          10.4.5.4 Et0/1          BGP
                    52    52          0          0          0    258          62          7
                    N    N    N          N          N          N

```

[SNIP]

MC#

Display Inside Prefixes

On the Master Controller, you have all the inside Traffic Classes. Because the inside prefixes are dynamically learnt in this example, we have all inside prefixes that are in the Border Routers BGP tables (and locally originated), and not only those we want to optimize.

Using show pfr master traffic-class inside:

Display Inside Prefixes

PfR:Solutions:InternetInboundLoadBalancing

MC#sh pfr master traffic-class inside

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 (inside)	Appl_ID	Dscp Prot			SrcPort	DstPort	SrcPrefix			
		State					Time	CurrBR	CurrI/F	Protocol
		PasSDly	PasLDly	PasSUn						
ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos			
10.10.1.0/24		N	N	N		N	N			
		INPOLICY			0	10.4.5.6	Et0/1			BGP
	U	U	0	0	0	0	1106		125	
	N	N	N	N	N	N				
100.4.81.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U
10.1.2.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U
10.10.2.0/24		N	N	N		N	N			
		INPOLICY			0	10.4.5.5	Et0/1			BGP
	U	U	0	0	0	0	19		2	
	N	N	N	N	N	N				
100.5.82.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U
10.2.3.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U
10.10.3.0/24		N	N	N		N	N			
		INPOLICY			0	10.4.5.5	Et0/1			BGP
	U	U	0	0	0	0	913		110	
	N	N	N	N	N	N				
100.6.83.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U
10.10.4.0/24		N	N	N		N	N			
		INPOLICY			0	10.4.5.5	Et0/1			BGP
	81	81	0	0	0	0	22		2	
	N	N	N	N	N	N				
10.4.5.0/24		N	N	N		N	N			
		DEFAULT*			0		U			U

MC#

What to check:

Display Inside Prefixes

PfR:Solutions:InternetInboundLoadBalancing

- We are only interested by inside prefixes under 10.10.0./16, but because we use the dynamic learning mode we have all inside prefixes that are in the BGP routing table. Apart from 10.10.x.y/24 prefixes, the rest are inter routers prefixes and there is very few traffic going there.
- Line1: prefix
- Line2: State of INPOLICY - Inside prefix in controlled by PfR and is in policy
- Line2: State of DEFAULT* - Inside prefix is not controlled by PfR and are under the routing protocol control
- Line2: Display the BR and external interface used for entrance
- Line3: Passive results

(EBw/IBw): the bandwidth for this prefix in egress and ingress direction.

- Line4: Active results

Not applicable for inside prefixes. Only Passive Monitoring is available for inbound optimization.

A new command available from 15.2(1)T gives exhaustive metrics for a specific Traffic class:

```
MC#sh pfr master traffic-class performance ip any 10.10.1.0/24
```

```
=====
Traffic-class: (inside)
Destination Prefix : 10.10.1.0/24          Source Prefix   : N/A
Destination Port   : N/A                   Source Port     : N/A
DSCP                : N                    Protocol        : N/A
Application Name   : N/A

General:
Control State      : Controlled using BGP
Traffic-class status : INPOLICY
Current Exit       : BR 10.4.5.6 interface Et0/1, Tie breaker was Non-OER
Time on current exit : 0d 0:0:11
Time remaining in current state : 0 seconds
Traffic-class type  : Configured
Improper config     : None

Last Out of Policy event:
Exit                : BR 10.4.5.6 interface Et0/1
Reason              : Range
Time since Out of Policy event : 0d 2:31:14
Link OOP; no prefix performance :

Average Passive Performance Current Exit: (Average for last 5 minutes)
Unreachable        : 208% -- Threshold: 50%
Delay               : 102 msec -- Threshold: 200 msec
Loss                : 84% -- Threshold: 10%
Egress BW          : 1025 kbps
Ingress BW         : 131 kbps
Time since last update : 0d 0:0:11
=====
```


MC#

Static Configuration and PfR Map

PfR Configuration

In the following Master Controller configuration, the usual outbound optimization commands are in place and we add a specific policy for inside prefixes. These inside prefixes are statically defined using a prefix-list.

The goal here is twofold:

- to statically defined the inside prefixes that we want PfR to optimize, ie prefixes under 10.10.0.0/16.
- to be able to define specific policies for inside prefixes by using a pfr-map.

```

!
key chain pfr
  key 0
    key-string cisco
!
pfr master
!-----
! Specific Policies for inside prefixes
!-----
!
policy-rules MAP-TEST2
max-range-utilization percent 10
logging
!
border 10.4.5.4 key-chain pfr
  interface Ethernet0/0 internal
  interface Ethernet0/1 external
!
border 10.4.5.5 key-chain pfr
  interface Ethernet0/0 internal
  interface Ethernet0/1 external
!
border 10.4.5.6 key-chain pfr
  interface Ethernet0/0 internal
  interface Ethernet0/1 external
!
learn
  traffic-class filter access-list ALLOW_REMOTE
  expire after time 300
!
!-----
! Global Policies for outbound traffic
!-----
!
max prefix total 10000 learn 10000
max range receive percent 10
mode monitor passive

```

```

resolve range priority 3
no resolve delay
no resolve utilization
!
!-----
! Inside Prefixes - Policy Definitions
!-----
!
pfr-map MAP-TEST2 10
match traffic-class prefix-list HQ_PREFIX inside
set holddown 90
set mode route control
set resolve range priority 3
no set resolve delay
no set resolve utilization
!
!
ip prefix-list HQ_PREFIX seq 5 permit 10.10.1.0/24
ip prefix-list HQ_PREFIX seq 10 permit 10.10.2.0/24
ip prefix-list HQ_PREFIX seq 15 permit 10.10.3.0/24
ip prefix-list HQ_PREFIX seq 20 permit 10.10.4.0/24
!

```

Master Controller Verification

Before starting to look at the results, a few checks are needed to verify that the configuration is correct. In this example, automatic learning is enabled for outside prefixes and static configuration is used for inside prefixes. By using static configuration, we can narrow down the inside prefixes to the only ones that are really interested.

Goal:

- Verify that the MC is active
- Verify that the learning process is enabled on the Master Controller for the outgoing traffic
- Display the policy settings.

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

You may have to wait a few cycles for the MC to be able to learn a few prefixes and then issue the command:

```

MC#sh pfr master
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 3949
Version: 3.3
Number of Border routers: 3
Number of Exits: 3
Number of monitored prefixes: 46 (max 10000)
Max prefixes: total 10000 learn 10000
Prefix count: total 46, learn 42, cfg 4
PBR Requirements met
Nbar Status: Inactive

```

PfR:Solutions:InternetInboundLoadBalancing

Auto Tunnel Mode: On

Border	Status	UP/DOWN		AuthFail	Version	DOWN Reason
10.4.5.6	ACTIVE	UP	02:37:28	0	3.3	
10.4.5.5	ACTIVE	UP	02:37:28	0	3.3	
10.4.5.4	ACTIVE	UP	02:37:28	0	3.3	

Global Settings:

```
max-range-utilization percent 8 recv 5
rsvp post-dial-delay 0 signaling-retries 1
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 37
```

Default Policy Settings:

```
backoff 90 900 90
delay relative 50
holddown 90
periodic 0
probe frequency 56
number of jitter probe packets 100
mode route control
mode monitor passive
loss relative 10
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
resolve range priority 3 variance 0
```

Learn Settings:

```
current state : STARTED
time remaining in current state : 77 seconds
throughput
no delay
no inside bgp
traffic-class filter access-list ALLOW_REMOTE
monitor-period 1
periodic-interval 0
aggregation-type prefix-length 24
prefixes 100 appls 100
expire after time 300
```

MC#

What to check:

- All Border Routers are up and running
- Number of Monitored Prefixes: 46 (max 10000) - we have 46 prefixes compared to 52 in the dynamic configuration. This is because we have statically configured the only interesting inside prefixes (and not all inside prefixes available).
- Prefix count: total 46, learn 42, cfg 4 - 42 prefixes learnt (remote prefixes) and 4 statically defined (the 4 inside prefixes)
- Prefixes: total 10000 learn 10000 - The max number of monitored prefixes, the max number of learnt prefixes.
- Global Settings: max-range-utilization is 8 (percentage of bandwidth difference between all exit interfaces), used for outgoing traffic

- Learn Settings: prefixes 100 appls 100 - 100 prefixes max automatically learnt per cycle (this could be changed via configuration)
- All default policy settings are displayed
- Learn is started (current state : STARTED)

Display Inside Traffic Classes

On the Master Controller, you have all the inside Traffic Classes previously configured as well as statistics. Because the inside prefixes are statically defined in this example, there is no learning and the associated TCs are directly inserted in the database. We also only have the interesting inside prefixes, those we want to optimize.

Using show pfr master traffic-class inside:

```
MC#sh pfr master traffic-class inside
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 (inside)	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Flags	State	Time	CurrBR	CurrI/F	Protocol
							PasSDly	PasLUn	PasSLos	PasLLos	EBw	IBw
							ActSDly	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos
10.10.1.0/24		N	N	N		N						
		INPOLICY		0		10.4.5.6						BGP
	103	101	133	31	50	119	1025	129				
	N	N	N	N	N	N						
10.10.2.0/24		N	N	N		N						
		INPOLICY		0		10.4.5.5						BGP
	U	U	0	0	0	0	24	2				
	N	N	N	N	N	N						
10.10.3.0/24		N	N	N		N						
		INPOLICY		0		10.4.5.5						BGP
	U	U	0	0	0	0	962	109				
	N	N	N	N	N	N						
10.10.4.0/24		N	N	N		N						
		INPOLICY*		0		10.4.5.4						U
	U	U	0	0	0	0	22	2				
	N	N	N	N	N	N						

MC#

What to check:

- Line1: prefix
- Line2: State of INPOLICY - Inside prefix in controlled by PfR and is in policy
- Line2: State of INPOLICY* - An asteric (*) indicates this prefix is uncontrolled by PfR (the parent route controls routing), but is currently inpolicy. The last column display U instead of BGP.
- Line2: Display the BR and external interface used for entrance
- Line3: Passive results

(EBw/IBw): the bandwidth for this prefix in egress and ingress direction.

- Line4: Active results

Not applicable for inside prefixes. Only Passive Monitoring is available for inbound optimization.

Dynamic Configuration and PfR Map

PfR Configuration

With this configuration, dynamic learning is enabled for both outside and inside prefixes. The inside prefix learning that happens for inbound optimization is by looking at the BGP tables on the border routers. We just ask BGP on the BRs for all networks that were advertised to eBGP peers over PfR external interfaces. Once we receive all the prefixes from the BRs, we just enter them into the monitoring database and start optimizing them (note that we ask for prefixes that originated in our AS only, so transit prefixes that were advertised by our BRs to eBGP peers, if any, will not be learnt). In addition to that we define a pfr-map for these inside prefixes. We can therefore define policies for inside prefixes with the "set" commands.

```
pfr master
policy-rules MAP-TEST1
max-range-utilization percent 8
logging
!
border 10.4.5.6 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.4.5.5 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.4.5.4 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
learn
inside bgp
traffic-class filter access-list ALLOW_REMOTE
expire after time 300
max prefix total 10000 learn 10000
max range receive percent 5
mode monitor passive
```

```

resolve range priority 3
no resolve delay
no resolve utilization
!
pfr-map MAP-TEST1 10
match pfr learn inside
set holddown 90
set mode route control
set resolve range priority 3
no set resolve delay
no set resolve utilization
!

```

Master Controller Verification

Again the first step is to check the Master Controller configuration.

Goal:

- Verify that the MC is active
- Verify that the BRs are connected
- Verify that the learning process is enabled on the Master Controller for the outgoing traffic
- Verify that the inside prefixes are correctly learnt from the BGP tables on the Border Routers
- Display the policy settings.

You may have to wait a few cycles for the MC to be able to learn a few prefixes and then issue the command:

```

MC#sh pfr master
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 3949
Version: 3.3
Number of Border routers: 3
Number of Exits: 3
Number of monitored prefixes: 52 (max 10000)
Max prefixes: total 10000 learn 10000
Prefix count: total 52, learn 52, cfg 0
PBR Requirements met
Nbar Status: Inactive
Auto Tunnel Mode: On

```

Border	Status	UP/DOWN		AuthFail	Version	DOWN Reason
10.4.5.4	ACTIVE	UP	00:03:17	0	3.3	
10.4.5.5	ACTIVE	UP	00:03:17	0	3.3	
10.4.5.6	ACTIVE	UP	00:03:15	0	3.3	

```

Global Settings:
max-range-utilization percent 8 recv 5
rsvp post-dial-delay 0 signaling-retries 1
mode route metric bgp local-pref 5000
mode route metric static tag 5000
trace probe delay 1000
logging

```

PfR:Solutions:InternetInboundLoadBalancing

```
exit holddown time 60 secs, time remaining 57
```

Default Policy Settings:

```
backoff 90 900 90
delay relative 50
holddown 90
periodic 0
probe frequency 56
number of jitter probe packets 100
mode route control
mode monitor passive
loss relative 10
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
resolve range priority 3 variance 0
```

Learn Settings:

```
current state : STARTED
time remaining in current state : 82 seconds
throughput
no delay
inside bgp
traffic-class filter access-list ALLOW_REMOTE
monitor-period 1
periodic-interval 0
aggregation-type prefix-length 24
prefixes 100 appls 100
expire after time 300
```

MC#

What to check:

- All Border Routers are up and running
- Number of monitored prefixes: 52 (max 10000) - 52 prefixes. if there is more than 10000 prefixes, the rest will be under the routing control.
- Prefixes: total 10000 learn 10000 - The max number of monitored prefixes, the max number of learnt prefixes.
- Global Settings: max-range-utilization is 8 (percentage of bandwidth difference between all exit interfaces), used for outgoing traffic
- Learn Settings: prefixes 100 appls 100 - 100 prefixes max automatically learnt per cycle (this could be changed via configuration)
- All default policy settings are displayed
- Learn is started (current state : STARTED)

Display Inside Prefixes

On the Master Controller, you have all the inside Traffic Classes. Because the inside prefixes are dynamically learnt in this example, we have all inside prefixes that are in the Border Routers BGP tables (and locally originated), and not only those we want to optimize. But because we have defined a specific pfr map for inside prefixes, we can define specific policies for inside prefixes that are different from the ones defined for outside prefixes.

Using show pfr master traffic-class inside:

Master Controller Verification

PfR:Solutions:InternetInboundLoadBalancing

MC#sh pfr master traffic-class inside

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 (inside)	Appl_ID	Dscp Prot			SrcPort	DstPort	SrcPrefix								
		State					Time	CurrBR	CurrI/F	Protocol					
		PasSDly	PasLDly	PasSUn							PasLUn	PasSLos	PasLLos	EBw	IBw
		ActSDly	ActLDly	ActSUn							ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos
10.10.1.0/24		N	N	N		N	N								
		INPOLICY			0	10.4.5.5	Et0/1	BGP							
	U	U	0	0	0	0	1094	148							
	N	N	N	N	N	N									
100.4.81.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							
10.1.2.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							
10.10.2.0/24		N	N	N		N	N								
		HOLDDOWN			59	10.4.5.4	Et0/1	BGP							
	U	U	0	0	0	0	20	2							
	N	N	N	N	N	N									
100.5.82.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							
10.2.3.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							
10.10.3.0/24		N	N	N		N	N								
		HOLDDOWN			58	10.4.5.4	Et0/1	BGP							
	U	U	0	0	0	0	874	107							
	N	N	N	N	N	N									
100.6.83.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							
10.10.4.0/24		N	N	N		N	N								
		INPOLICY*			0	10.4.5.4	Et0/1	U							
	U	U	0	0	0	0	17	2							
	N	N	N	N	N	N									
10.4.5.0/24		N	N	N		N	N								
		DEFAULT*			0		U	U							

MC#

What to check:

- We are only interested by inside prefixes under 10.10.0/16, but because we use the dynamic learning mode we have all inside prefixes that are in the BGP routing table. Apart from 10.10.x.y/24 prefixes,

the rest are inter routers prefixes and there is very few traffic going there.

- Line1: prefix
- Line2: State of INPOLICY - Inside prefix in controlled by PfR and is in policy
- Line2: State of INPOLICY* - An asteric (*) indicates this prefix is uncontrolled by PfR (the parent route controls routing), but is currently inpolicy. The last column display U instead of BGP.
- Line2: State of DEFAULT* - Inside prefix is not controlled by PfR and are under the routing protocol control
- Line2: Display the BR and external interface used for entrance
- Line3: Passive results

(EBw/IBw): the bandwidth for this prefix in egress and ingress direction.

- Line4: Active results

Not applicable for inside prefixes. Only Passive Monitoring is available for inbound optimization.

Verify Load balancing and Enforcement

Bandwidth used on exit links

To verify the accuracy of the load-balancing scheme on external interfaces, there is one command available on the Master Controller:

```
MC#sh pfr master border detail
Border          Status  UP/DOWN      AuthFail  Version  DOWN Reason
10.4.5.6        ACTIVE  UP           02:36:02    0   3.3
  Tu0           TUNNEL  UP
  Et0/1         EXTERNAL UP
  Et0/0         INTERNAL UP

External        Capacity  Max BW  BW Used  Load Status  Exit Id
Interface       (kbps)   (kbps)  (kbps)  (%)
-----
Et0/1           Tx        2000    1800    653    32 UP        3
                Rx        2000    148     7

-----
Border          Status  UP/DOWN      AuthFail  Version  DOWN Reason
10.4.5.5        ACTIVE  UP           02:36:02    0   3.3
  Tu0           TUNNEL  UP
  Et0/1         EXTERNAL UP
  Et0/0         INTERNAL UP

External        Capacity  Max BW  BW Used  Load Status  Exit Id
Interface       (kbps)   (kbps)  (kbps)  (%)
-----
Et0/1           Tx        2000    1800    680    34 UP        2
                Rx        2000    134     6
```

PfR:Solutions:InternetInboundLoadBalancing

```

Border      Status  UP/DOWN      AuthFail  Version  DOWN Reason
10.4.5.4    ACTIVE  UP           02:36:02    0   3.3
  Tu0       TUNNEL  UP
  Et0/1     EXTERNAL UP
  Et0/0     INTERNAL UP
  
```

```

External      Capacity      Max BW      BW Used      Load Status      Exit Id
Interface      (kbps)        (kbps)      (kbps)        (%)
-----
Et0/1         Tx           2000         1800          773           38 UP           1
              Rx           2000         2000           3              0
  
```

MC#

Verifying Enforcement

To implement BGP inbound optimization, PfR manipulates eBGP advertisements to influence the best entrance selection for traffic bound for inside prefixes.

Two methods can be used:

- **BGP Autonomous System Number Prepend** (this is the default method): when PfR selects a best entrance for an inside prefix, extra autonomous system hops are prepended one at a time (up to a maximum of six) to the inside prefix BGP advertisement over the other entrances.
- **BGP Autonomous System Number Community Prepend**: when PfR selects a best entrance for an inside prefix, a BGP prepend community is attached to the inside prefix BGP advertisement from the network to another autonomous system such as an ISP. This community value has a special meaning for the upstream ISP and is therefore configured under the external interface definition (under Border Router definition).

In this example, we simply use the AS-PATH prepend enforcement method. After a few cycles, PfR controls the prefixes and use AS-PATH prepend to enforce the ingress path to the appropriate Border Routers. If we look at R9 or R10 BGP tables, we will see the new BGP table with prefixes that have the AS prepended.

On the Master Controller, you can check whether entrance for inside prefixes is controlled by BGP, and which Border Router (and its external interface) reports the biggest bandwidth for a specific inside prefix.

Using show pfr master traffic-class inside:

```

MC#sh pfr master traffic-class inside
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 (inside)  Appl_ID Dscp Prot      SrcPort      DstPort SrcPrefix
  
```

Bandwidth used on exit links

PfR:Solutions:InternetInboundLoadBalancing

	Flags		State		Time		CurrBR	CurrI/F	Protocol
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos		EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS		ActSLos	ActLLos
10.10.1.0/24			N	N	N		N	N	
			INPOLICY		0		10.4.5.6	Et0/1	BGP
	103	101		133	31	50	119	1025	129 <--- HERE
	N	N	N	N	N	N	N		
10.10.2.0/24			N	N	N		N	N	
			INPOLICY		0		10.4.5.5	Et0/1	BGP
	U	U		0	0	0	0	24	2 <--- HERE
	N	N	N	N	N	N	N		
10.10.3.0/24			N	N	N		N	N	
			INPOLICY		0		10.4.5.5	Et0/1	BGP
	U	U		0	0	0	0	962	109 <--- HERE
	N	N	N	N	N	N	N		
10.10.4.0/24			N	N	N		N	N	
			INPOLICY*		0		10.4.5.4	Et0/1	U
	U	U		0	0	0	0	22	2 <--- HERE
	N	N	N	N	N	N	N		

MC#

What to check:

Line3 (marked with: <--- HERE):

- Gives the ingress and egress bandwidth for each inside prefix (EBw/IBw) as reported by the BR.
- As an example, inside prefix 10.10.2.0/24 has an egress Bandwidth of 24 Kbps and an Ingress Bandwidth of 2 Kbps as reported by Border Router R5 (10.4.5.5).

Now, you can also have all the details for this inside prefix 10.10.2.0/24 and check what is the bandwidth reported by the other 2 Border Routers:

```
MC#sh pfr master traffic-class inside detail
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 (inside)	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]

PfR:Solutions:InternetInboundLoadBalancing

```

10.10.2.0/24      INPOLICY      0 10.4.5.6      Et0/1          BGP      0
                  105          105            0              0          0          0
                  N            N              N              N          7          1
                  10.4.5.4      Et0/1          BGP          1
                  0            105            0              0          0          0
                  N            N              N              N          5          0
                  10.4.5.5      Et0/1          BGP          1
                  0            101            0              0          0          348
                  N            N              N              N          4          0
  
```

[SNIP]

It appears clearly in this example that Border Router R5 (10.4.5.5) reports the highest bandwidth for the inside prefix 10.10.2.0/24.

Let's have a look at the remote site router to check that AS-PATH prepend has been used to influence the entrance:

```

R9#sh bgp
BGP table version is 66, local router ID is 20.9.9.9
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

   Network          Next Hop           Metric LocPrf Weight Path
*  10.10.1.0/24     20.9.82.1
*>                   20.9.81.1
*                   20.9.83.1          100
*  10.10.2.0/24     20.9.82.1
*                   20.9.81.1          100
*>                   20.9.83.1          100
*  10.10.3.0/24     20.9.82.1
*                   20.9.81.1          100
*>                   20.9.83.1          100
*  10.10.4.0/24     20.9.82.1
*>                   20.9.81.1          100
*                   20.9.83.1          100
  
```

[SNIP]

R9#

Conclusion

One of the basic solution provided by PfR is to be able to load-balance traffic among multiple exit interface. The configuration is straightforward and very efficient for outgoing traffic. Inbound traffic entrance can be influenced by using the AS Path Prepend or Community but is still dependant upon the various BGP policies implemented by the Service Providers over the Internet.