

Cisco Performance Routing (PfR) Solution Guides

PfR NBAR Based Application Control

Navigation

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

Contents

- [1 Introduction](#)
- [2 PfR Network Topology Used](#)
- [3 Configuration](#)
 - ◆ [3.1 Master Controller Configuration - R3](#)
 - ◆ [3.2 Border Routers Configuration ? Routers R4 and R5](#)
- [4 PfR Configuration Verification](#)
 - ◆ [4.1 Master Controller](#)
 - ◆ [4.2 Border Routers](#)
- [5 Further Debugging Tips](#)
- [6 Conclusion](#)

Introduction

The Performance Routing with NBAR/CCE Application Recognition feature introduces the ability to profile an application-based traffic class using Network-Based Application Recognition (NBAR). NBAR is a classification engine that recognizes and classifies a wide variety of protocols and applications, including web-based and other difficult-to-classify applications and protocols that use dynamic TCP/UDP port assignments. PfR uses NBAR to recognize and classify a protocol or application, and the resulting traffic classes are added to the PfR application database to be passively and/or actively monitored.

As a consequence PfR is capable of optimizing applications identified by Network-Based Application Recognition (NBAR). Using the `?match traffic-class application nbar nbar-appl-name [nbar-appl-name ...]`

prefix-list prefix-listname? command or ?match pfr learn list <list-name>? , PFR will be able to use NBAR to recognize applications and optimize them based on the policies installed on the Master Controller. Internally, PFR uses classmaps to control and optimize NBAR based applications.

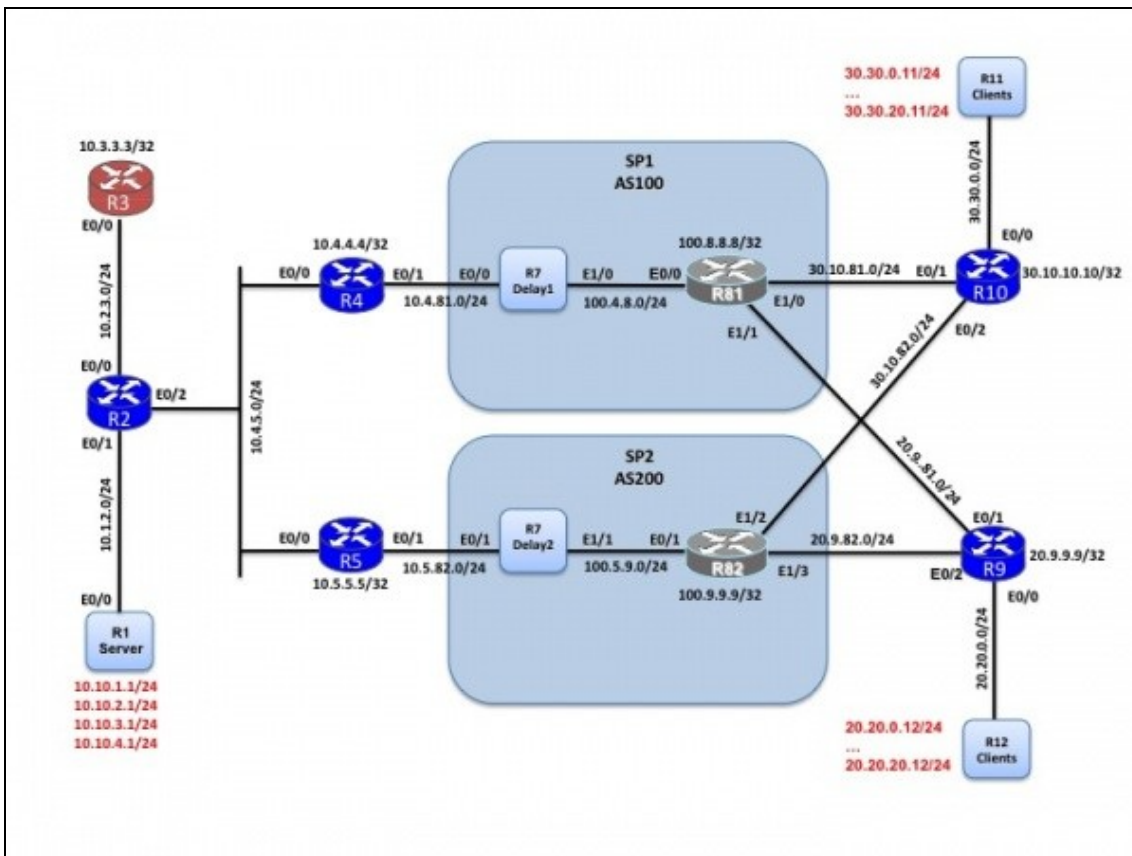
The current version of PfR (PfRv2) only supports the initial version of NBAR and therefore only support a small subset of all applications supported in the new DPI engine NBAR2. While it is possible to use Application Based policies with NBAR1 directly within PfR, it is recommended to mark the packets on ingress with NBAR2 and then use the remarked DSCP to apply PfR policies. For more information, check [here](#)

PfR Network Topology Used

The central site has two Border Routers(R4 & R5), connected to two separate Service Providers.

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
- Traffic Simulator tool is used between R1 and R11 to emulate traffic, both TCP and UDP.
- 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).



R2 is an iBGP peer for both border routers and as such has the BGP route table. PfR being not active, the parent routes are BGP based on R2, R4 and R5 and R5 is the preferred exit point (higher local preference applied on ingress on R5) as seen below:

Route to destination prefixes BRANCH1 (20.20.0.0/16):

```
R2#sh ip bgp 20.20.0.0/16
BGP routing table entry for 20.20.0.0/16, version 7
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  200 20
    100.5.82.1 (metric 11) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 200, valid, internal, best
R2#
```

Route to destination prefixes BRANCH2 (30.30.0.0/16):

```
R2#sh ip bgp 30.30.0.0/16
BGP routing table entry for 30.30.0.0/16, version 12
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  200 30
    100.5.82.1 (metric 11) from 10.5.5.5 (10.5.5.5)
      Origin IGP, metric 0, localpref 200, valid, internal, best
R2#
```

Configuration

As usual, most of the configuration is done on the Master Controller.

Master Controller Configuration - R3

Basic configuration to establish session between MC and BR

```
! This is the basic configuration required to establish session between MC and BR.
! With this basic configuration, learning is enabled, route control is in place and
! load-balancing happens on all configured external interfaces.
!
! It includes
! ? Key-chain configuration for authentication.
! ? Specification of BR?s IP Address and internal/external interface on the BR.
!   If there is a tunnel between the BRs, it should be configured as internal
! - Specification of the Carrier name. Not used in the policies described hereafter.
! - Disabling auto-tunnels as BRs are L2 adjacent.
```

PfR:Solutions:NBAR

```
!  
! Notes on the new defaults:  
! - Automatic learning is enabled  
! - mode route control is enabled  
! - Maximum transmit utilization is 90%  
! - Load-balancing is now used as the last resolver and cannot be disabled. Range is 20%.  
!  
!  
key chain pfr  
  key 0  
    key-string cisco  
!  
pfr master  
  logging  
!  
border 10.4.4.4 key-chain pfr  
  interface Ethernet0/0 internal  
  interface Ethernet0/1 external  
  link-group SP1  
!  
border 10.5.5.5 key-chain pfr  
  interface Ethernet0/0 internal  
  interface Ethernet0/1 external  
  link-group SP2  
!  
no mode auto-tunnels  
!
```

Enable NetFlow version9 Export (optional)

```
! Flow Exporter Definition  
! Where do I want my data sent?  
!  
flow exporter MYEXPORTER  
  destination 10.151.1.131  
  source Loopback0  
  transport udp 9991  
  option interface-table timeout 300  
  option sampler-table timeout 300  
  option application-table timeout 300  
!  
!  
! Add NetFlow export to PfR  
!  
pfr master  
  exporter MYEXPORTER  
!
```

Enable logging

```
! Following configuration is to enable logging.  
! This will print PfR related syslog messages on the console.  
!  
pfr master  
  logging
```

!

Define learning parameters, disable global learning

```

! By default now:
!
! - Automatic learning is enabled
!
! - Continuous learn cycle, each 1 minute duration
!   periodic-interval = 0 (forever)
!   monitor-period = 1 (minutes)
!
! - traffic-class sorting based on ?throughput? at the end of
!   each learning cycle.
!
! - Anything traffic that doesn?t match the learn list will be
!   learned under global learn and will be optimized using default policy.
!
! - In this example we decided to disable global learning
!   configure a filter using a named access-list.
!   The goal here is to learn only branch traffic using learn list (described in the next section)
!
pfr master
  learn
    ! Disable Global learn
    !
    traffic-class filter access-list DENY_GLOBAL_LEARN_LIST
    !
    !
    ! Access-list for disabling global learn.
    !
ip access-list extended DENY_GLOBAL_LEARN_LIST
  deny ip any any
!

```

Define learn-list for an NBAR based application

```

! Following is a learn list configuration for NBAR based application control.
! Learn-list is defined for a specific remote site BRANCH1.
! In the configuration below, http has been used as an example.
! You can tune the aggregation mask to match your need.
! Here we used the default value of /24
! Sorting is based on throughput (default)
!
pfr master
  learn
    list seq 10 refname HTTP
    traffic-class application nbar http filter BRANCH1_PREFIX
    aggregation-type prefix-length 24
    throughput
    !
    !
    ! Define the prefix list of the branch
    !

```

PfR:Solutions:NBAR

```
ip prefix-list BRANCH1_PREFIX seq 5 permit 20.20.0.0/16
!
```

Policy configuration for controlling NBAR applications and specific for a branch

```
! Following is a policy configuration to control NBAR applications
! specific to one branch.
!
! This should be repeated for each branch. It includes
!
! ? match command is to specify that this policy should be applied
!   to all the traffic-classes learned under list HTTP
!
! - Re-evaluate exit every 90 sec (periodic 90)
!
! ? delay threshold is configured as 60 msec. The delay measured by PfR is
!   Round-Trip-Time.
!
! ? Resolver setting is configured to set delay as the highest priority
!
! - Load-balancing is now used as the last resolver (cannot be disabled for now)
!
! - monitor mode is set to fast. This means probe all external interfaces
!   all the time. When Out-of-Policy condition is detected on the current
!   exit results on alternate exit is available for quick decision. In other
!   modes alternate exits are probed only when current link is determined to
!   be OOP. The fast mode helps in switching the path quickly when the
!   problem is detected.
!
! ? Probe frequency is set to 8 seconds and can be changed to a lesser value
!   if the application being controlled is critical.
!
!- A forced echo probe has been setup
!
! Default Values (new in 15.2(3)T and 3.6):
! - holddown timer set to 90 to shorten the time needed to move to INPOLICY state
!
!
pfr-map MYMAP 10
match pfr learn list HTTP
set periodic 90
set mode select-exit good
set delay threshold 60
set mode monitor fast
set resolve delay priority 1 variance 5
set active-probe echo 20.9.9.9
set probe frequency 8
!
```

Assign the policy to the PfR Master Controller

```
!
pfr master
policy-rules MYMAP
!
```

Note: You can choose to replace the learnt list based NBAR application based learning with the configured application/prefix-list method as well. This would require the user to replace the match statement in the configured pfr-map:

From:

```
pfr-map MYMAP 10
  match pfr learn list HTTP
  (?)
```

To:

```
pfr-map MYMAP 10
  match traffic-class application nbar http prefix-list myprefix
  (?)
!
! where myprefix is a configured ip prefix-list:
!
ip prefix-list myprefix seq 5 permit 100.2.5.2/32
```

Border Routers Configuration ? Routers R4 and R5

Following are common PfR configuration commands in the Border Router R4 and R5.

```
! This is the minimum configuration required on the BR?s.
! It includes
! ? Key-chain configuration for authentication.
! ? Specification of MC?s IP Address and Local interface. The IP address
!   of the local interface will be used as source IP address in communicating
!   with MC.
!
key chain pfr
  key 0
    key-string cisco
!
pfr border
  local Ethernet0/0
  master 10.3.3.3 key-chain pfr
!
```

PfR Configuration Verification

Master Controller

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.

```
R3#sh pfr master
```

```
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 3949
Version: 3.3
Number of Border routers: 2
Number of Exits: 2
Number of monitored prefixes: 4 (max 5000)
Max prefixes: total 5000 learn 2500
Prefix count: total 4, learn 2, cfg 0
PBR Requirements met
Nbar Status: Active
Auto Tunnel Mode: Off
```

Border	Status	UP/DOWN	AuthFail	Version	DOWN Reason
10.4.4.4	ACTIVE	UP	14:01:33	0	3.3
10.5.5.5	ACTIVE	UP	14:01:33	0	3.3

Global Settings:

```
max-range-utilization percent 20 recv 0
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 0
```

Default Policy Settings:

```
backoff 90 900 90
delay relative 50
holddown 90
periodic 0
probe frequency 56
number of jitter probe packets 0
mode route control
mode monitor both
loss relative 10
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
trigger-log percentage 30
```

Learn Settings:

```
current state : STARTED
time remaining in current state : 80 seconds
throughput
no delay
no inside bgp
traffic-class filter access-list DENY_GLOBAL_LEARN_LIST
monitor-period 1
```



```
periodic-interval 0
aggregation-type prefix-length 24
prefixes 100 appls 100
expire after time 720

Learn-List seq 10 refname HTTP
Configuration:
  Traffic-Class Application: http
  Filter: BRANCH1_PREFIX
  Aggregation-type: prefix-length 24
  Learn type: throughput
  Session count: 1000 Max count: 1000
  Policies assigned: 10
  Status: ACTIVE
Stats:
  Traffic-Class Count: 2
```

R3#

What to check:

- Both Border Routers are up and running
- Check to make sure that NBAR Status is Active.
- Check to make sure that PBR Requirements are met
- Number of automatically learned applications ? 2 in this case
- All default policy settings are displayed
- Learn is started (current state : STARTED)
- Learn-list status is ACTIVE and learn-type is throughput
- Make sure that all the configured learn lists show up
- For each learn-list, check the Traffic Class access-list (if configured), prefix-list (if configured) and the policy number associated (which is the pfr-map it refers to).
- The Traffic Class count is also displayed which allows you to check whether the learning process worked well (here 2 applications is being learnt under HTTP learn list).

Check if the application is automatically learnt under each learn-list

R3#show pfr master learn list

```
Learn-List seq 10 refname HTTP
Configuration:
  Traffic-Class Application: http
  Filter: BRANCH1_PREFIX
  Aggregation-type: prefix-length 24
  Learn type: throughput
  Session count: 1000 Max count: 1000
  Policies assigned: 10
  Status: ACTIVE
Stats:
  Traffic-Class Count: 2
  Traffic-Class Learned:
    Appl Prefix 20.20.1.0/24 http
    Appl Prefix 20.20.2.0/24 http
```

R3#

Check the status of NBAR application being controlled on the Master

This command is used to display information about the status of an application identified using NBAR for each PfR border router. The following partial output shows information about the status of applications identified using NBAR at PfR border routers R4 (10.4.4.4) and R5 (10.5.5.5). If the NBAR application is not supported on one or more border routers, then all the traffic classes related to that NBAR application are marked inactive and cannot be optimized using PfR.

```
R3#show pfr master nbar application
OER NBAR Applications Status:
```

NBAR Appl	10.4.4.4	10.5.5.5
bgp	Valid	Valid
bittorrent	Valid	Valid
bridge	Invalid	Invalid
bstun	Invalid	Invalid
cdp	Invalid	Invalid
citrix	Valid	Valid
clns	Invalid	Invalid
clns_es	Invalid	Invalid
clns_is	Invalid	Invalid
cmns	Invalid	Invalid
compressedtcp	Invalid	Invalid
cuseeme	Valid	Valid

[SNIP]

```
R3#
```

```
R3#show pfr master nbar application | inc http
http          Valid          Valid
secure-http   Valid          Valid
R3#
```

Check the policy associated

```
R3#show pfr master policy 10
* Overrides Default Policy Setting
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 HTTP
  backoff 90 900 90
*delay threshold 60
  holddown 90
*periodic 90
*probe frequency 8
  number of jitter probe packets 0
  mode route control
*mode monitor fast
  loss relative 10
```

```
jitter threshold 20
mos threshold 3.60 percent 30
unreachable relative 50
next-hop not set
forwarding interface not set
trigger-log percentage 30
*resolve delay priority 1 variance 5

Forced Assigned Target List:
  active-probe echo 20.9.9.9 target-port 0
R3#
```

Check the active probes

```
R3#show pfr master active-probes forced
      OER Master Controller active-probes
Border   = Border Router running this Probe
Policy   = Forced target is configure under this policy
Type     = Probe Type
Target   = Target Address
TPort    = Target Port
N - Not applicable
```

The following Forced Probes are running:

Border	State	Policy	Type	Target	TPort	Dscp
10.4.4.4	ACTIVE	10	echo	20.9.9.9	N	defa
10.5.5.5	ACTIVE	10	echo	20.9.9.9	N	defa

```
R3#
```

Verify Traffic Class Statistics

As soon as you see:

```
R3#
*Sep 20 20:24:43.515: %OER_MC-5-NOTICE: Prefix Learning WRITING DATA
*Sep 20 20:24:43.589: %OER_MC-5-NOTICE: Prefix Learning STARTED
R3#
```

You should be able to see the traffic classes on the Master Controller. You will need a few cycles before having all applications in INPOLICY state.

Verify learnt Traffic Classes

PfR:Solutions:NBAR

On the Master Controller, you have all the Traffic Classes (learnt as well as statistics): Use show pfr master traffic-class to see the state of the application:

Traffic Classes are in HOLDDOWN state first:

```
R3#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	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	CurrI/F	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos	
20.20.1.0/24		http	N	N	N	N	0.0.0.0/0		
			HOLDDOWN		@47	10.4.4.4	Et0/1		CCE
	U	U	0	0	0	0	59	4	
	51	51	0	0	N	N	N	N	
20.20.2.0/24		http	N	N	N	N	0.0.0.0/0		
			HOLDDOWN		@45	10.4.4.4	Et0/1		CCE
	U	U	0	0	0	0	59	4	
	51	51	0	0	N	N	N	N	

R3#

Then moved into INPOLICY state:

```
R3#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	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	CurrI/F	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos	
20.20.1.0/24		http	N	N	N	N	0.0.0.0/0		
			INPOLICY		@72	10.4.4.4	Et0/1		CCE
	U	U	0	0	0	0	55	4	
	49	50	0	0	N	N	N	N	
20.20.2.0/24		http	N	N	N	N	0.0.0.0/0		
			INPOLICY		@70	10.4.4.4	Et0/1		CCE
	U	U	0	0	0	0	54	4	

49 50 0 0 N N N N

R3#

What to check:

- Traffic Classes Appl_ID is http (NBAR)
- state is INPOLICY (could be HOLDOWN and then INPOLICY after 90 sec)
- Time column: remaining time in the current state
- CurrBR: Current Border Router and its external interface used
- 2 lines of results: 1st line is passive results and 2nd line is active results
- Protocol: enforcement method used. Here CCE, means classmaps with a set next-hop (NBAR within PBR is not supported yet).

A closer look at the results

Let's have a look at the show pfr master traffic-class output in more detail. We will use the show pfr master traffic-class performance command to do this. Note: This command is only available from 15.2(1)T

```
R3#show pfr master traffic-class performance ip any 20.20.1.0/24
```

```
=====
Traffic-class:
Destination Prefix : 20.20.1.0/24                    Source Prefix     : 0.0.0.0/0
Destination Port   : N/A                            Source Port       : N/A
DSCP               : N                              Protocol          : N/A
Application Name:   : http

General:
Control State                                        : Controlled using CCE
Traffic-class status                                : INPOLICY
Current Exit                                         : BR 10.4.4.4 interface Et0/1, Tie breaker was Delay
Time on current exit                                : 0d 0:2:2
Time remaining in current state                   : @52 seconds
Traffic-class type                                  : Learned
Improper config                                     : None

Last Out of Policy event:
Exit                                                 : BR 10.4.4.4 interface Et0/1
Reason                                              : Delay
Time since Out of Policy event                    : 0d 0:2:2
Active Delay Performance                          : 49 msec
Active Delay Threshold                            : 60 msec

Average Passive Performance Current Exit: (Average for last 5 minutes)
Unreachable                                         : 0% -- Threshold: 50%
Delay                                                : 0 msec -- Threshold: 60 msec
Loss                                                : 0% -- Threshold: 10%
Egress BW                                          : 55 kbps
Ingress BW                                         : 4 kbps
Time since last update : 0d 0:1:35

Average Active Performance Current Exit: (Average for last 5 minutes)
Unreachable                                         : 0% -- Threshold: 50%
Delay                                                : 49 msec -- Threshold: 60 msec
```

Last Resolver Decision:

BR	Interface	Status	Reason	Performance	Threshold
10.5.5.5	Et0/1	Eliminated	Delay	105 msecs	60 msecs
10.4.4.4	Et0/1	Best Exit	Delay	51 msecs	60 msecs

R3#

The sections under:

- Traffic-class: -- displays the application-prefix that is being learnt or monitored.
- General: -- will display whether the application is being controlled, the controlling protocol, and the state. In the above output the application is INPOLICY and is being controlled using CCE, which is the class-map based control that is used along with a set next hop, to enforce the path.
- The Active and Passive Performance measurements are also displayed in the sections that follow.
- Lastly, you will see the Resolver decision that was used to select the best exit. In the example above, we have a delay threshold of 60 ms configured in the policy MYMAP. The delay being calculated on BR(R5)- 10.5.5.5, exit Et0/1 is 105 msec and therefore it is being eliminated. The delay calculated on BR(R4)- 10.4.4.4, exit Et0/1 is 51 msec, which is within the configured delay threshold and therefore it is chosen as the best exit.

Border Routers

Check the netflow cache on the Border Routers

On Border Router that is supposed to learn the flow, issue ?show ip cache flow? and verify that the flow is being learnt.

Note: There is no need to explicitly configured NetFlow, the Master Controller tells the BRs to enable NetFlow.

How does PfR modify Paths on the Border Routers

PfR will create classmaps (CCE) on the non-forwarding Borders to forward any flows matching the application over the PFR configured internal interfaces to the forwarding Border. Finally on the forwarding Border, there will be another classmap directing traffic for that application out of the border?s external exit, that is chosen as the best exit.

In-order for the CCE based route control to be successful, the BRs have to be adjacent either through a direct connection or one hop away.

The primary next-hop from BGP is R5. Without PfR, traffic is flowing through R5 and SP2. Let's check on the Border Router - R5, where we clearly see that packets are being sent out to R4.

```
R5#show pfr border routes cce
Class-map oer-class-acl-oer_cce#2-stile-http, permit, sequence 0, mask 24
  Match clauses:
```

```
ip address (access-list): oer_cce#2
style: http
Set clauses:
ip next-hop 10.4.5.4
interface Ethernet0/0
Statistic:
Packet-matched: 18491
```

R5#

```
R5#show ip access-lists dynamic oer_cce#2
Extended IP access list oer_cce#2
 10 permit ip any 20.20.2.0 0.0.0.255 (9579 matches)
 20 permit ip any 20.20.1.0 0.0.0.255 (9801 matches)
R5#
```

The packets are being finally sent out of R4's external exit since that was chosen as the best exit.

```
R4#show pfr border routes cce
Class-map oer-class-acl-oer_cce#2-style-http, permit, sequence 0, mask 24
Match clauses:
ip address (access-list): oer_cce#2
style: http
Set clauses:
ip next-hop 100.4.81.1
interface Ethernet0/1
Statistic:
Packet-matched: 76
```

R4#

```
R4#show ip access-lists dynamic oer_cce#2
Extended IP access list oer_cce#2
 10 permit ip any 20.20.2.0 0.0.0.255 (54 matches)
 20 permit ip any 20.20.1.0 0.0.0.255 (22 matches)
R4#
```

Further Debugging Tips

The following set of debugs might help narrow down any issues seen while troubleshooting NBAR CCE control within PFR:

- On the Border

```
debug pfr border cce [detail]
debug pfr border nbar [detail]
```

- On the Master Controller

```
debug pfr master prefix [appl] [detail]
```

- Further, if you have PFR Master logging configured, you should be able to see Notice messages for indications of other actions being taken by PFR.
- The following debugs can be turned on the Master Controller to see if netflow or active probing updates are being received from the borders for ingress/egress flows.

```
MC#debug pfr master collector ?  
  active-probes  Display PfR active probe debugs  
  netflow        Display PfR NetFlow debugs
```

Conclusion

The configuration examples provided above explain how applications can be learnt and optimized using learn lists. The same configuration can be extended to even consider configured NBAR based application control ? with prefix-lists associated with them. The number of applications being optimized can be fine tuned based on the specific requirement.