

Cisco Performance Routing (PfR) Solution Guides

PfR - Enterprise Fast Failover Using Fast Mode

Navigation

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

Contents

- [1 PfR Features that Enable Traffic Optimization](#)
 - ◆ [1.1 Link Utilization](#)
 - ◆ [1.2 Range](#)
 - ◆ [1.3 Traffic Class Performance](#)
- [2 Enterprise Failover Needs](#)
- [3 PfR Solution Used](#)
- [4 PfR Network Topology Used](#)
- [5 PfR Configuration](#)
 - ◆ [5.1 Master Controller Configuration](#)
 - ◆ [5.2 Border Routers Configuration](#)
- [6 PfR Configuration Verification](#)
- [7 Master Controller](#)
 - ◆ [7.1 Border Routers](#)
- [8 Verify Traffic Classes Statistics](#)
 - ◆ [8.1 Traffic Classes](#)
 - ◆ [8.2 A closer look at the results](#)
 - ◆ [8.3 How PfR modifies Paths](#)
- [9 Add Blackouts and observe how PfR reacts](#)
 - ◆ [9.1 Drop Test Description](#)
 - ◆ [9.2 Drop Test Execution](#)

- ◆ [9.3 Blackout Test Conclusion](#)
- [10 Add delay and observe how PfR reacts](#)
 - ◆ [10.1 Delay Test Description](#)
 - ◆ [10.2 Delay Test Execution](#)
 - ◆ [10.3 Delay Test Conclusion](#)
- [11 Possible Optimization](#)

PfR Features that Enable Traffic Optimization

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 Failover Needs

Critical, voice or video applications need high availability. A fast failover mechanism should be deployed to allow re-routing sensitive application traffic to a fallback link in case of a link or node failure but also in case of soft errors, brownouts or blackouts. Typical routing protocols cannot cope with soft errors.

PfR Solution Used

To achieve a fast failover when a Traffic Class is Out of Policy, PfR mode **fast** has to be used. Fast failover monitoring is designed for traffic classes that are very sensitive to performance issues or congested links, and voice or video traffic is very sensitive to any dropped links. But Fast Failover can also be used for very delay sensitive business applications.

Fast mode gives PfR the ability to re-route a sub-optimally performing application, very quickly. In best case, we have observed re-route within 3 seconds of event detection.

How does fast mode work?

- In fast monitoring mode, all exits are continuously probed using active probing. This ensures that the latest performance information is always available and it enables PfR to react instantly to a network event.
- The probe frequency can be set to a lower frequency in fast failover monitoring mode than for other monitoring modes, to allow a faster failover capability.
- Fast failover monitoring can be used with all types of active probes: ICMP echo, jitter, TCP connection, and UDP echo.
- Performance data from active probing is used for routing decisions while the throughput data from passive probing is used to load balance the links.
- PfR is not a fast re-route technology. Fast mode should be used only for performance sensitive traffic like real time communication applications.

Notes:

- In fast monitoring mode, probe targets are learned as well as learned prefixes. To avoid triggering large numbers of probes in the network, use fast monitoring mode only for real time applications and critical applications with performance sensitive traffic.
- User can configure 2 sec probe frequency in fast mode only.
- if < 4 second probe freq is configured, the user is not allowed to set any other mode. Error is displayed.

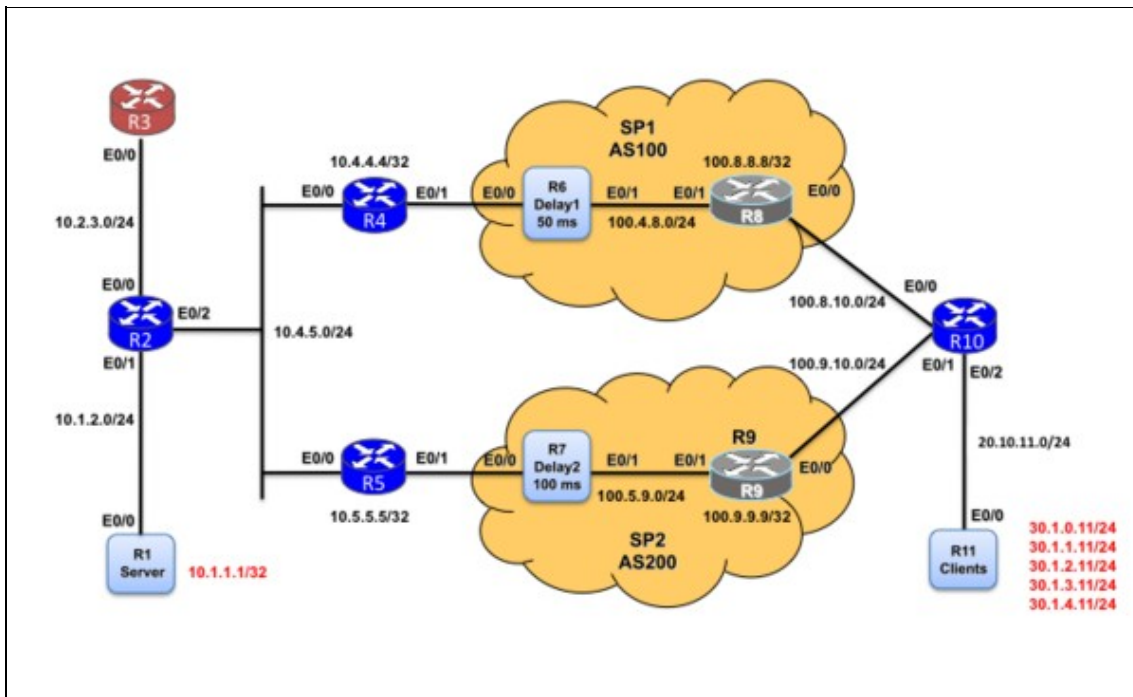
In this example, the fast failover monitoring mode is enabled and the critical traffic to be monitored is identified using DSCP AF31. To reduce some of the overhead that fast failover monitoring produces, we statically assigned a force target.

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
- 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).



PfR Configuration

Master Controller Configuration

- Basic configuration to establish session between MC and BR

```
! This is the basic configuration required to establish session
! between MC and BR. It includes
! ? Key-chain configuration for authentication.
! ? Specification of BR's IP Address and internal/external interface
! on the BR.
! ? Specification of Maximum transmit utilization of 90%
!
<pre>
!
key chain pfr
  key 0
    key-string cisco
!
pfr master
  logging
!
border 10.4.5.5 key-chain pfr
```

PfR:Solutions:EnterpriseFastFailover

```
interface Ethernet0/1 external
  max-xmit-utilization percentage 90
interface Ethernet0/0 internal
!
border 10.4.5.4 key-chain pfr
interface Ethernet0/1 external
  max-xmit-utilization percentage 90
interface Ethernet0/0 internal
!
!
```

- Enable logging

```
! Following configuration is
! - to enable logging. This will print PfR related syslog messages on
!   the console.
! ? to disable default policy of load balancing.
!
pfr master
  no max-range-utilization
  logging
```

- Define learning parameters, disable global learning

```
! - To enable continuous learn cycle, each 1 minute duration ?
!   configure periodic-interval value 0 and monitor-period value 1 (minutes)
!   By default each cycle is 5 minute and occurs ever 2 hrs.
!
! - Sort the traffic-class 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.
!   To disable global learn 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
    throughput
    periodic-interval 0
    monitor-period 1

    ! 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 sensitive applications

PfR:Solutions:EnterpriseFastFailover

```
! Following is a learn list configuration for sensitive application
! specific to one branch.
! This should be repeated for each branch. It includes
! ? Here CRITICAL (assuming it is marked as af31) is learned using
!   access-list CRITICAL and branch specific filter BRANCH1.
!
pfr master
learn
  list seq 10 refname BRANCH1_CRITICAL
  traffic-class access-list CRITICAL filter BRANCH1
  aggregation-type prefix-length 32
  throughput
!
!
! Access-list for CRITICAL
!
ip access-list extended CRITICAL
  permit ip any any dscp af31
!
! Prefix-list for BRANCH1
!
ip prefix-list BRANCH1 seq 5 permit 30.1.0.0/16
!
```

- Policy configuration for sensitive applications and specific for a branch

```
! Following is a policy configuration sensitive application 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-class learned under list BRANCH1_CRITICAL
!
! - Re-evaluate exit every 90 sec (periodic 90)
!
! - holddown timer set to 90 to shorten the time needed to move to INPOLICY state
!
! ? delay threshold is configured as 300 msec. The delay measured by PfR is
!   Round-Trip-Time. For example, video conference delay higher than 150 ms one-way
!   decreases the Quality-of-Experience.
!
! ? Loss is set to 50,000 (packets-per-million). i.e 5%
!
! ? Resolver setting is configure to set the priority in the order of
!   loss and delay. Range and utilization are DISABLED for sensitive application.
!
! ? Probe packets are set to 20 to reduce the probe traffic.
!   Instead of configuring only 1 probe (with 60 probe packet) it is better to configure
!   2 or 3 probes with 20 probe packets (resulting into lower or same number of total probe packet)
!
pfr-map MYMAP 10
  match pfr learn list BRANCH1_CRITICAL
  set periodic 90
  set holddown 90
  set delay threshold 300
  set mode route control
  set mode monitor fast
  set resolve loss priority 2 variance 5
  set resolve delay priority 5 variance 5
```

```
no set resolve range
no set resolve utilization
set loss threshold 50000
set active-probe echo 30.1.1.11
set probe frequency 2
!
```

- Assign the policy to the PfR Master Controller

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

Border Routers Configuration

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

```
! This is the minimum configuration required to BR. 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
  logging
  local Ethernet0/0
  master 10.2.3.254 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.

PfR:Solutions:EnterpriseFastFailover

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: 8 (max 5000)

Max prefixes: total 5000 learn 2500

Prefix count: total 8, learn 4, cfg 0

PBR Requirements met

Nbar Status: Inactive

Border	Status	UP/DOWN		AuthFail	Version
10.4.5.4	ACTIVE	UP	00:04:08	0	3.0
10.4.5.5	ACTIVE	UP	00:04:08	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 300 3000 300

delay relative 50

holddown 300

periodic 0

probe frequency 56

number of jitter probe packets 100

mode route observe

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

resolve range priority 12 variance 0

resolve utilization priority 13 variance 20

Learn Settings:

current state : STARTED

time remaining in current state : 92 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 BRANCH1_CRITICAL

Configuration:

Traffic-Class Access-list: CRITICAL

Filter: BRANCH1

Aggregation-type: prefix-length 32

Learn type: throughput

Session count: 50 Max count: 100

Policies assigned: 10


```
Status: ACTIVE
Stats:
Traffic-Class Count: 4
MC#
```

What to check:

- Both Border Routers are up and running
 - Number of Monitored Prefixes: 8 and 4 are automatically 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, prefix-list and 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 (here 4 TCs are learnt under BRANCH1_CRITICAL).
-
- Check the prefixes/application automatically learnt under each learn-list

```
MC#sh pfr master learn list

Learn-List seq 10 refname BRANCH1_CRITICAL
Configuration:
Traffic-Class Access-list: CRITICAL
Filter: BRANCH1
Aggregation-type: prefix-length 32
Learn type: throughput
Session count: 50 Max count: 100
Policies assigned: 10
Status: ACTIVE
Stats:
Traffic-Class Count: 4
Traffic-Class Learned:
Appl Prefix 30.1.2.11/32 af31 256
Appl Prefix 30.1.1.11/32 af31 256
Appl Prefix 30.1.3.11/32 af31 256
Appl Prefix 30.1.0.11/32 af31 256
MC#
```

- Check the policy associated

```
MC#sh 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 BRANCH1_CRITICAL
backoff 300 3000 300
*delay threshold 300
*holddown 90
*periodic 90
```

```
*probe frequency 2
  number of jitter probe packets 100
*mode route control
*mode monitor fast
  mode select-exit good
*loss threshold 50000
  jitter threshold 20
  mos threshold 3.60 percent 30
  unreachable relative 50
  next-hop not set
  forwarding interface not set
*resolve loss priority 2 variance 5
*resolve delay priority 5 variance 5

  Forced Assigned Target List:
  active-probe echo 30.1.1.11 target-port 0
MC#
```

- Check the active probes

```
MC#sh 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
10.4.5.5	ACTIVE	10	echo	30.1.1.11	N	af31
10.4.5.4	ACTIVE	10	echo	30.1.1.11	N	af31

MC#

Note: active probes use the DSCP value of the Traffic Classes in the learn-list. If there is multiple DSCP used, then multiple probes (one for each DSCP value) will be defined.

Border Routers

Active probes are defined on the Master Controller but are forged from each Border Routers. Because we are using fast mode, each Border Router generates active probes.

- On Border Router R4

```
R4#sh pfr border active-probes
      OER Border active-probes
Type     = Probe Type
Target   = Target IP Address
```

PfR:Solutions:EnterpriseFastFailover

TPort = Target Port
Source = Send From Source IP Address
Interface = Exit interface
Att = Number of Attempts
Comps = Number of completions
N - Not applicable

Type	Target	TPort	Source	Interface	Att	Comps
DSCP						
echo	30.1.1.11	N	100.4.8.4	Et0/1	346	346
104						

R4#

What to check:

- There are active probes running on the Border Router
 - Att column: gives the number of attempts
 - Comps column: gives the number of completed probes
 - if "comps" is not equal or greater than Att then do the below command to check ip sla configuration.
-
- Look at the IP SLA configuration generated by PfR on R4:

```
R4#sh ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 3267
Owner: Optimized Edge Routing (OER)
Tag:
Operation timeout (milliseconds): 1000
Type of operation to perform: icmp-echo
Target address/Source address: 30.1.1.11/100.4.8.4
Type Of Service parameter: 0x68
Request size (ARR data portion): 28
Verify data: No
Vrf Name:
Schedule:
  Operation frequency (seconds): 2 (not considered if randomly scheduled)
  Next Scheduled Start Time: Start Time already passed
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): Forever
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 1000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
Enhanced History:
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None
```

R4#

- Look at the IP SLA statistics on R4, and note the delay of ~50ms:

```
R4#
R4#sh ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 3267
    Latest RTT: 49 milliseconds
Latest operation start time: *13:47:44.831 MET Mon Mar 28 2011
Latest operation return code: OK
Number of successes: 420
Number of failures: 0
Operation time to live: Forever

R4#
```

- On Border Router R5

```
R5#sh pfr border active-probes
    OER Border active-probes
Type      = Probe Type
Target    = Target IP Address
TPort     = Target Port
Source    = Send From Source IP Address
Interface = Exit interface
Att       = Number of Attempts
Comps    = Number of completions
N - Not applicable

Type      Target          TPort Source          Interface          Att   Comps
DSCP
echo      30.1.1.11              N 100.5.9.5        Et0/1              364   364
104

R5#
```

What to check:

- There are active probes running on the Border Router
 - Att column: gives the number of attempts
 - Comps column: gives the number of completed probes
 - if "comps" is not equal or greater than Att then do the below command to check ip sla configuration.
-
- Look at the IP SLA configuration generated by PfR on R5:

PfR:Solutions:EnterpriseFastFailover

```
R5#sh ip sla configuration
IP SLAs Infrastructure Engine-III
Entry number: 3267
Owner: Optimized Edge Routing (OER)
Tag:
Operation timeout (milliseconds): 1000
Type of operation to perform: icmp-echo
Target address/Source address: 30.1.1.11/100.5.9.5
Type Of Service parameter: 0x68
Request size (ARR data portion): 28
Verify data: No
Vrf Name:
Schedule:
  Operation frequency (seconds): 2 (not considered if randomly scheduled)
  Next Scheduled Start Time: Start Time already passed
  Group Scheduled : FALSE
  Randomly Scheduled : FALSE
  Life (seconds): Forever
  Entry Ageout (seconds): never
  Recurring (Starting Everyday): FALSE
  Status of entry (SNMP RowStatus): Active
Threshold (milliseconds): 1000
Distribution Statistics:
  Number of statistic hours kept: 2
  Number of statistic distribution buckets kept: 1
  Statistic distribution interval (milliseconds): 20
Enhanced History:
History Statistics:
  Number of history Lives kept: 0
  Number of history Buckets kept: 15
  History Filter Type: None
```

R5#

- Look at the IP SLA statistics on R5 and note the delay of ~250ms:

```
R5#sh ip sla statistics
IPSLAs Latest Operation Statistics

IPSLA operation id: 3267
  Latest RTT: 251 milliseconds
Latest operation start time: *13:50:28.825 MET Mon Mar 28 2011
Latest operation return code: OK
Number of successes: 502
Number of failures: 0
Operation time to live: Forever
```

R5#

Verify Traffic Classes Statistics

As soon as you see:

```
MC#
*Mar 25 15:14:33.544: %OER_MC-5-NOTICE: Prefix Learning WRITING DATA
*Mar 25 15:14:33.618: %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.

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	Flags	Appl_ID	Dscp	Prot	SrcPort	DstPort	SrcPrefix	Protocol
	PasSDly	PasLDly	PasSUn	PasLUn	PasSJos	PasLJos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSJos	ActLJos
30.1.0.11/32			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	@0		10.4.5.4	Et0/1	PBR
	51	51	0	0	0	0	70	7
	53	52	0	0	N	N	N	N
30.1.1.11/32			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	@0		10.4.5.4	Et0/1	PBR
	52	52	0	0	0	0	68	8
	53	52	0	0	N	N	N	N
30.1.2.11/32			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	@0		10.4.5.4	Et0/1	PBR
	52	52	0	0	0	0	73	7
	53	52	0	0	N	N	N	N
30.1.3.11/32			N af31	256		N	N 0.0.0.0/0	
			INPOLICY	@0		10.4.5.4	Et0/1	PBR
	52	52	0	0	0	0	66	8
	53	52	0	0	N	N	N	N

MC#

A closer look at the results

Let's have a look at a few entries.

From the command `show oer master traffic-class`, let's focus on a specific TC: `30.1.0.11/32 dscp af31`:

```
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
-----
30.1.0.11/32          N af31  256          N          N 0.0.0.0/0
                    INPOLICY      @0          10.4.5.4 Et0/1          PBR
                    51      51      0      0      0      0      70      7
                    53      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) as this is the path with the lowest delay (remember that policies are based on loss then delay and remember the results from the IP SLA probes).
- **Line4: Active results**
- **Line4 (ActSDly, ActLDly):** short-term and long-term active delay, measured by active probes that we configured. 50 ms reported, coming from the IP SLA probe results from R4.
- **Line4 (ActSUn, ActLUn):** short-term and long-term Unreachable results.
- **Line4 (ActSLos, ActLLos):** short-term and long-term Loss results.

How PfR modifies Paths

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.

As shown before, CRITICAL traffic (dscp af31) is being forced out R4 due to delay policy.

PfR creates several dynamic route-maps to enforce the path for the sensitive applications. Depending on two dynamic ACLs, PBR enforce a next-hop to WAN1 or WAN2.

Let's check on the first border router R5 where we clearly see that packets are redirected to R4:

```
R5#sh route-map dynamic
route-map OER_INTERNAL_RMAP, permit, sequence 0, identifier 4278190085
  Match clauses:
    ip address (access-lists): oer#1
  Set clauses:
    ip next-hop 10.4.5.4
    interface Ethernet0/0
  Policy routing matches: 89574 packets, 61575074 bytes
Current active dynamic routemaps = 1
R5#
```

We can also check the dynamic ACL on R5 that matches sensitive traffic with DSCP af31.

```
R5#sh ip access-lists dynamic
Extended IP access list oer#1
  134217727 permit ip any host 30.1.3.11 dscp af31 (22980 matches)
  268435455 permit ip any host 30.1.0.11 dscp af31 (22993 matches)
  536870911 permit ip any host 30.1.2.11 dscp af31 (22972 matches)
  1073741823 permit ip any host 30.1.1.11 dscp af31 (23035 matches)
R5#
```

Add Blackouts and observe how PfR reacts

Drop Test Description

We are simulating a soft error, where the bgp control packets make it through but the application traffic gets dropped. BGP still believes that the path is the best path when in reality, the application packets are getting dropped.

Drop Test Execution

- Applying filter on the WAN1 path (primary path for application traffic):

```
*17:42:31.781 MET Mon Mar 28 2011
```

You should see PfR messages, stating that AF31 TCs are Out Of Policies (OOP), due to unreachable reason. This happens 3-5 sec after the initial drop.

```
MC#
*Mar 28 17:42:35.501: %OER_MC-5-NOTICE: Active REL Unreachable OOP Appl Prefix 30.1.1.11/32 af31 2
*Mar 28 17:42:35.501: %OER_MC-5-NOTICE: Active REL Unreachable OOP Appl Prefix 30.1.0.11/32 af31 2
*Mar 28 17:42:35.501: %OER_MC-5-NOTICE: Active REL Unreachable OOP Appl Prefix 30.1.3.11/32 af31 2
MC#
*Mar 28 17:42:35.501: %OER_MC-5-NOTICE: Active REL Unreachable OOP Appl Prefix 30.1.2.11/32 af31 2
*Mar 28 17:42:35.705: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.1.11/32 af31 256, BR 10.4.5
*Mar 28 17:42:35.705: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.0.11/32 af31 256, BR 10.4.5
*Mar 28 17:42:35.705: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.3.11/32 af31 256, BR 10.4.5
MC#
*Mar 28 17:42:35.705: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.2.11/32 af31 256, BR 10.4.5
MC#
```

- Traffic has been forced out R5 due to Out of Policy. TCs are in the holddown state to avoid flapping. Holddown timer is 300sec by default and can be changed in the PfR configuration (here changed to 90sec).

```
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
	PasSDly	PasLDly	PasSUn	PasLUn	PasSJos	PasLJos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSJos	ActLJos
30.1.0.11/32			N af31 256		N	N 0.0.0.0/0		
			HOLDDOWN	@85		10.4.5.5 Et0/1		PBR
	U	U	0	0	0	0	82	5
	246	246	0	0	N	N	N	N
30.1.1.11/32			N af31 256		N	N 0.0.0.0/0		
			HOLDDOWN	@85		10.4.5.5 Et0/1		PBR
	U	U	0	0	0	0	70	8
	246	246	0	0	N	N	N	N
30.1.2.11/32			N af31 256		N	N 0.0.0.0/0		

PfR:Solutions:EnterpriseFastFailover

		HOLDDOWN		@78	10.4.5.5 Et0/1		PBR
U	U	0	0	0	0	80	5
246	246	0	0	N	N	N	N
30.1.3.11/32		N af31 256		N	N 0.0.0.0/0		
		HOLDDOWN		@84	10.4.5.5 Et0/1		PBR
U	U	0	0	0	0	80	7
246	246	0	0	N	N	N	N

MC#

- After the holddown timer expires, Traffic Classes move to INPOLICY state:

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
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos	EBw	IBw
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos

30.1.0.11/32			N af31	256	N	N 0.0.0.0/0		
			INPOLICY	@79		10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	24	3
	252	252	0	0	N	N	N	N
30.1.1.11/32			N af31	256	N	N 0.0.0.0/0		
			INPOLICY	@1		10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	24	3
	252	252	0	0	N	N	N	N
30.1.2.11/32			N af31	256	N	N 0.0.0.0/0		
			INPOLICY	@85		10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	26	2
	252	252	0	0	N	N	N	N
30.1.3.11/32			N af31	256	N	N 0.0.0.0/0		
			INPOLICY	@78		10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	24	3
	252	252	0	0	N	N	N	N

MC#

Blackout Test Conclusion

PfR is able to detect blackouts in the WAN. In this test, only Application traffic is dropped and therefore BGP still has a route through the path with blackouts.

Without PfR, the application traffic would have been dropped without further notice.

Add delay and observe how PfR reacts

Delay Test Description

This scenario is one of the worst case. This is not a failure scenario, not even a blackout condition but a scenario where soft errors or brownouts are observed in the primary WAN, leading to an increased delay. This delay increase would have a huge impact on the critical applications running over the WAN.

The initial state is that sensitive traffic is forced out R4 (WAN1), delay is 50 ms initially as measured by active probes.

Delay Test Execution

- Add delay 500 ms to R4

```
R6#sh clock
*17:29:47.731 MET Mon Mar 28 2011
```

- You should see PfR messages, stating that AF31 TCs are Out Of Policies (OOP), due to delay reason. This happens 7-10 sec after the initial delay increase.

```
MC#
*Mar 28 17:29:53.267: %OER_MC-5-NOTICE: Active ABS Delay OOP Appl Prefix 30.1.1.11/32 af31 256, de
*Mar 28 17:29:53.267: %OER_MC-5-NOTICE: Active ABS Delay OOP Appl Prefix 30.1.0.11/32 af31 256, de
*Mar 28 17:29:53.267: %OER_MC-5-NOTICE: Active ABS Delay OOP Appl Prefix 30.1.3.11/32 af31 256, de
*Mar 28 17:29:53.267: %OER_MC-5-NOTICE: Active ABS Delay OOP Appl Prefix 30.1.2.11/32 af31 256, de
MC#
*Mar 28 17:29:53.268: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.1.11/32 af31 256, BR 10.4.5
*Mar 28 17:29:53.268: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.0.11/32 af31 256, BR 10.4.5
*Mar 28 17:29:53.269: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.3.11/32 af31 256, BR 10.4.5
*Mar 28 17:29:53.269: %OER_MC-5-NOTICE: Route changed Appl Prefix 30.1.2.11/32 af31 256, BR 10.4.5
MC#
```

PfR:Solutions:EnterpriseFastFailover

- Traffic has been forced out R5 due to Out of Policy. TCs are in the holddown state to avoid flapping. Holddown timer is 300sec by default and can be changed in the PfR configuration (here changed to 90sec).

```
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		Protocol
	Flags						State	Time	
	PasSDly	PasLDly	PasSUn	PasLUn	PasSLos	PasLLos			
	ActSDly	ActLDly	ActSUn	ActLUn	ActSJit	ActPMOS	ActSLos	ActLLos	
30.1.0.11/32		N	af31	256	N	N	0.0.0.0/0		
		HOLDDOWN		@35		10.4.5.5	Et0/1		PBR
	127	127	0	0	0	0	75		6
	252	251	0	0	N	N	N		N
30.1.1.11/32		N	af31	256	N	N	0.0.0.0/0		
		HOLDDOWN		@34		10.4.5.5	Et0/1		PBR
	144	144	0	0	0	0	69		6
	252	251	0	0	N	N	N		N
30.1.2.11/32		N	af31	256	N	N	0.0.0.0/0		
		HOLDDOWN		@34		10.4.5.5	Et0/1		PBR
	145	145	0	0	0	0	80		6
	252	251	0	0	N	N	N		N
30.1.3.11/32		N	af31	256	N	N	0.0.0.0/0		
		HOLDDOWN		@38		10.4.5.5	Et0/1		PBR
	144	144	0	0	0	0	75		6
	252	251	0	0	N	N	N		N

MC#

- After the holddown timer expires, Traffic Classes move to INPOLICY state:

```
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		Protocol
	Flags						State	Time	

PfR:Solutions:EnterpriseFastFailover

	PasSDly ActSDly	PasLDly ActLDly	PasSUn ActSUn	PasLUn ActLUn	PasSJos ActSJit	PasLJos ActPMOS	EBw ActSJos	IBw ActLJos
30.1.0.11/32			N af31 256 INPOLICY	@61	N	N 0.0.0.0/0 10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	24	2
	253	252	0	0	N	N	N	N
30.1.1.11/32			N af31 256 INPOLICY	@55	N	N 0.0.0.0/0 10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	26	2
	253	252	0	0	N	N	N	N
30.1.2.11/32			N af31 256 INPOLICY	@56	N	N 0.0.0.0/0 10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	25	2
	253	252	0	0	N	N	N	N
30.1.3.11/32			N af31 256 INPOLICY	@63	N	N 0.0.0.0/0 10.4.5.5 Et0/1		PBR
	252	252	0	0	0	0	24	3
	253	252	0	0	N	N	N	N

MC#

Delay Test Conclusion

Using PfR allows soft errors and brownout detection. In this test, there is no way BGP can re-reroute based on a delay increase which can deeply impact the critical applications. PfR can detect this kind of soft errors and can re-route based on user policies.

Possible Optimization

Use jitter probes to get more accurate statistics on loss, delay and jitter. Configure 3 probes instead of only one for better reachability.

- Policy configuration for sensitive applications and specific for a branch

```
! Following is a policy configuration sensitive application 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-class learned under list BRANCH1_CRITICAL
!
! ? delay threshold is configured as 300 msec. The delay measured by PfR is
!   Round-Trip-Time. For example, video conference delay higher than 150 ms one-way
!   decreases the Quality-of-Experience.
!
```

PfR:Solutions:EnterpriseFastFailover

```
! ? Loss is set to 50,000 (packets-per-million). i.e 5%
!  
! ? Resolver setting is configure to set the priority in the order of  
!   loss and delay. Range and utilization are DISABLED for sensitive application.  
!  
! ? Probe packets are set to 20 to reduce the probe traffic.  
!   Instead of configuring only 1 probe (with 60 probe packet) it is better to configure  
!   2 or 3 probes with 20 probe packets (resulting into lower or same number of total probe packet)  
!  
pfr-map MYMAP 10  
match pfr learn list BRANCH1_CRITICAL  
set delay threshold 300  
set mode route control  
set mode monitor fast  
set resolve loss priority 2 variance 5  
set resolve delay priority 5 variance 5  
no set resolve range  
no set resolve utilization  
set loss threshold 50000  
set active-probe udp-jitter 30.1.1.11 target-port 2001 codec g729a  
set active-probe udp-jitter 30.1.1.11 target-port 2002 codec g729a  
set active-probe udp-jitter 30.1.1.11 target-port 2003 codec g729a  
set probe frequency 2  
!
```