PfR Performance Routing

Massimiliano Sbaraglia

PfR advantage

- PfR abilita i processi di routing su base best path (non più su prefix-destination) basandosi su definite policy;
- PfR potenzia i tradizionali protocolli di routing (bgp, eigrp, ospf, pbr) attraverso metriche di performance realtime attraverso IOS feature quali interfaces, netflow, IP SLA, etc….);
- Le performance di routing sono ottimizzate attraverso parametri quali:
  - Reachability
  - Delay
  - Packet Loss
  - Jitter
  - MOS (Mean Opinion Scope)
  - Troughput
  - Traffic Load
  - Cost Policies

PfR components

- Master Controller (MC): policy decision maker, dove sono definite ed applicate a varie classi di traffico che traversano il BR
  - Il MC controlla il BR e mantiene un database centrale di informazioni collezionate dal BR
  - Il MC può essere collocato insieme al BR in uffici di piccole/media dimensioni e può essere uno chassis standalone in uffici di grandi dimensioni
  - L'MC comunica con il BR attraverso un TCP socket autenticato
    - APPLY POLICY, VERIFICATION, REPORTING
    - NO PACKET FORWARDING
    - NO INSPECTION REQUIRED
- Border Router (BR): colleziona dati attraverso la loro Netflow cache and IP SLA probes, prendendo decisioni e gestendo il packet switching path tra users
    - GAIN NETWORK VISIBILITY (Learn, Measure)
    - FORWARDING PATH
    - IMPORRE MC DECISION (Path Enforcement)
    - ALMENO DUE EXTERNAL INTERFACE ED UNA INTERNAL INTERFACE

PfR design example

PfR Domain Peering:
advertise Site Prefix IP
advertise IP SLA responder IP addressess

Prefix	Responder	SITE
A	1	1
B	2	2
C D E	3 4	3

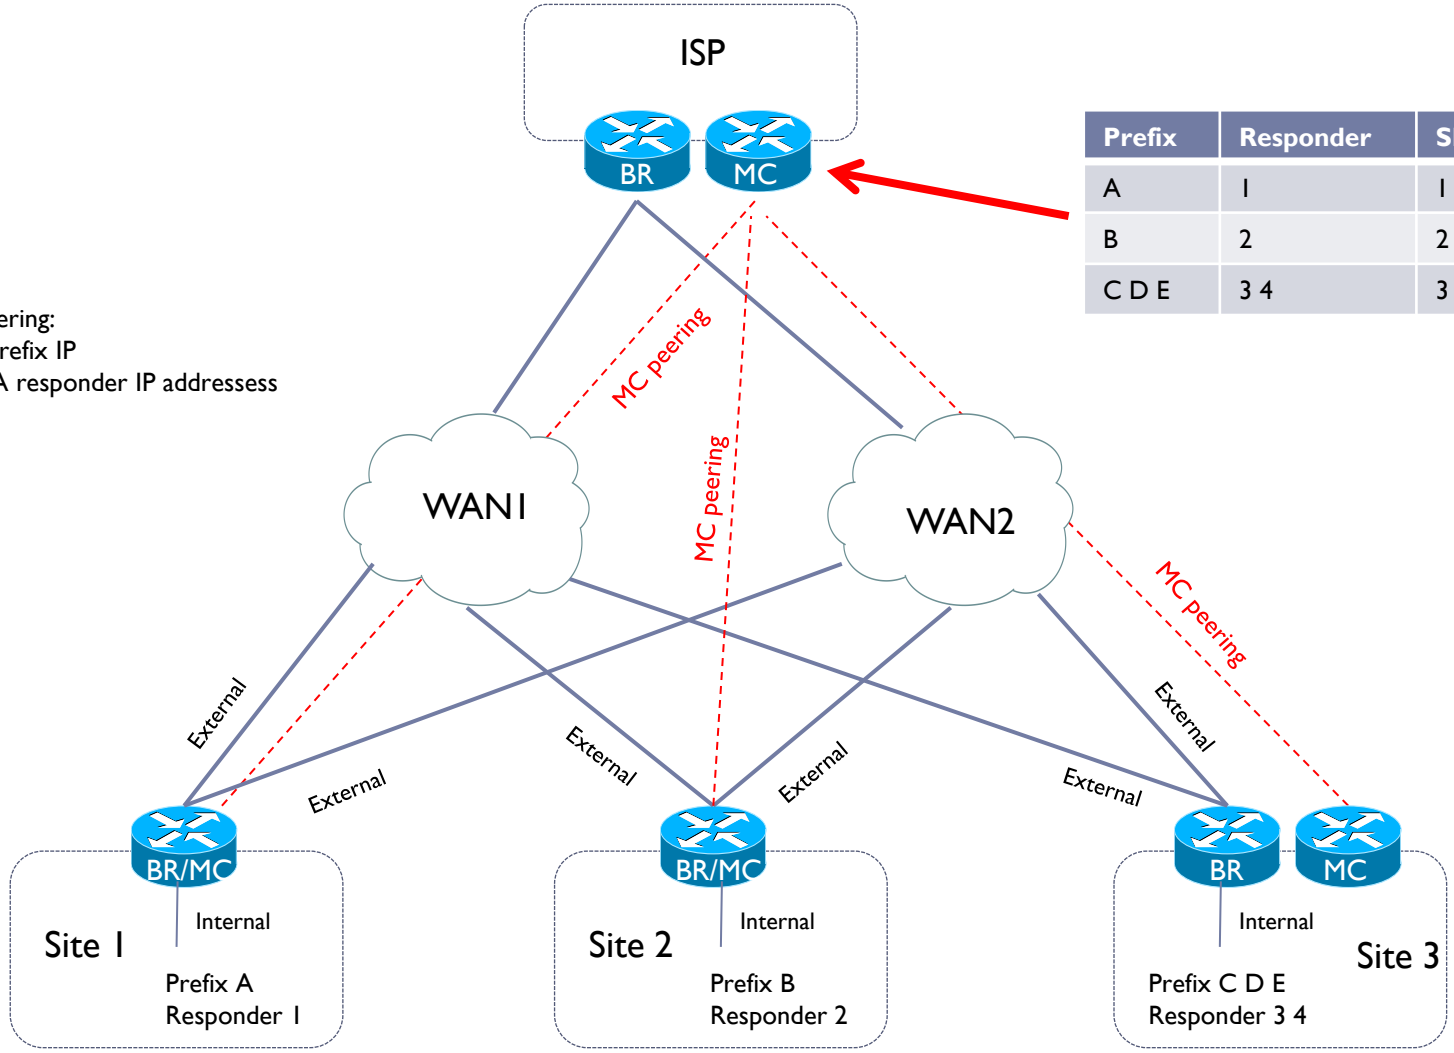

PfR policy engine

- Learn Application: compito del MC di richiedere al BR quali sono le applicazioni rispondenti alle classi di traffico interessate
  - destination prefix with or without port, dscp, source prefix or even application using NBAR
  - profiling process can be entirely automatic based on the top talkers (using Netflow) or configured manually.
- Measure Application Performance: collezione statistiche di traffico
  - Monitor Modes: Passive, Active, Both, Fast
  - Netflow for UDP (bandwidth) and TCP flows (availability, delay, bandwidth, loss)
  - IP SLA for TCP and UDP flows (Availability, delay, loss, jitter, MOS)
- Apply Policies:
  - Use measured application data to determine whether managed traffic-class is out of policy (OOP) and if an alternate path can meet the policy requirements.
- Enforce: re-route traffic
  - Prefix Control: Inject BGP or Static routes
  - Application Control: Dynamic Route-map/PBR for traffic classes defined by ACLs, NBAR, unsupported routing protocols (OSPF, ISIS) or, BRs running a mix of routing protocols.
- Verify that the new route matches the policy

PfR provisioning BR

- PfR configurazione è centralizzata nel MC; Il Border Router configurazione include solamente:
  - the source address definition (to setup the peering with its MC)
  - The MC IP address.
  - The key-chain used to secure the connection with the MC.

Example BR config:

```
key chain pfr
   key 0
      key-string cisco
!
pfr border
  logging local
  loopback0
  master 10.10.10.3 key-chain pfr
!
```

PfR provisioning MC

- PfR configurazione iniziale MC
  - the IP address of each BR and their respective internal and external interfaces. It's extremely important to list all internal and external interfaces as PfR is only going to monitor flows from internal to external interfaces.
  - The key-chain used to secure the connection with the BRs.

Example

```
key chain pfr
  key 0
   key-string cisco
!
pfr master
  logging
!
border 10.10.10.4 key-chain pfr
   interface Ethernet0/0 internal
   interface Ethernet0/1 external
!
border 10.10.10.5 key-chain pfr
   interface Ethernet0/0 internal
    interface Ethernet0/1 external
```

This concludes the initial provisioning. From now one, all PfR configuration will be defined on the Master Controller only.

PfR configurazione manuale

- Le prefixes o applicazioni sono manualmente configurate all'interno del PfR database.; non esiste nessun processo di learning

Example di configurazione di specifiche prefix:

```
pfr master
 policy-rules MYMAP
 logging
 border 10.10.10.4 key-chain key1
    interface Ethernet0/0 internal
    interface Ethernet0/1 external
!
 border 10.10.10.5 key-chain key1
    interface Ethernet0/0 internal
    interface Ethernet0/1 external
!
ip prefix-list SITE1 seq 5 permit 30.1.0.0/16
!
pfr-map MYMAP 10 match ip address prefix-list SITE1
!
```

PfR configurazione manuale

Example di configurazione di specifiche applicazioni:

```
pfr master
  policy-rules MYMAP
    logging border 10.10.10.4 key-chain key1
      interface Ethernet0/0 internal
      interface Ethernet0/1 external
!
 border 10.10.10.5 key-chain key1
    interface Ethernet0/0 internal
    interface Ethernet0/1 external
!
ip prefix-list FILTER_SITE1 seq 10 permit 10.1.1.0/24
ip prefix-list FILTER_SITE1 seq 20 permit 10.1.2.0/24
```

PfR configurazione manuale

Example di configurazione di specifiche applicazioni:

```
!
ip prefix-list FILTER_SITE2 seq 10 permit 10.1.3.0/24
ip prefix-list FILTER_SITE2 seq 20 permit 10.1.4.0/24
!
# Define FTP application
ip access-list extended MY_APP 10
   permit tcp any any eq 21
!
pfr-map MYMAP 10 match traffic-class application telnet prefix-list FILTER_SITE1
set mode select-exit good
set delay threshold 2000
set mode route control
set mode monitor both
no set resolve delay
set active-probe echo 10.1.1.10
```

PfR configurazione manuale

Example di configurazione di specifiche applicazioni:

```
pfr-map MYMAP 20
  match traffic-class application http prefix-list FILTER_SITE2
   set mode select-exit good
   set delay threshold 2000
   set mode route control
   set mode monitor both
   no set resolve delay
!
pfr-map MYMAP 30
  match traffic-class access-list MY_APP filter FILTER_SITE2
  set mode select-exit good
  set delay threshold 2000
  set mode route control
  set mode monitor both
  no set resolve delay
  set active-probe echo 10.1.3.10
```

PfR configurazione automatica

- PfR determina le classi di traffico attraverso flussi che attraverso i BR; il processo di learning è attivo.
  - PfR automatically tracks the top talkers based on netflow information received from the BRs:

```
!
key chain pfr
  key 0
    key-string cisco
!
pfr master
  logging
    border 10.10.10.4 key-chain pfr
    interface Ethernet0/0 internal
   interface Ethernet0/1 external
!
border 10.10.10.5 key-chain pfr
interface Ethernet0/0 internal
interface Ethernet0/1 external
!
```

PfR learn list

- PfR supporta in processo di learning list in configuration mode per semplificare l'apprendimento delle classi di traffico interessate:

```
ip prefix-list SITE1 seq 5 permit 30.1.0.0/16
!
pfr master max-range-utilization percent 10
policy-rules MYMAP
logging
learn
  throughput
   list seq 10 refname SITE_BUSINESS
   traffic-class application ssh filter SITE1
   throughput
   list seq 20 refname SITE_BE
   traffic-class prefix-list SITE1
   throughput
holddown 180
mode select-exit best
periodic 180
```

PfR learn list: uso di ACL

- L'uso di ACL serve per definire applicazioni:

```
ip access-list extended MY_APP 10
   permit tcp any any eq 21
!
ip prefix-list FILTER_SITE1 seq 10 permit 10.1.1.0/24
ip prefix-list FILTER_SITE1 seq 20 permit 10.1.2.0/24
!
ip prefix-list FILTER_SITE2 seq 10 permit 10.1.3.0/24
ip prefix-list FILTER_SITE2 seq 20 permit 10.1.4.0/24
!
pfr master
  learn
   list seq 10 refname SITE1 traffic-class access-list MY_APP filter FILTER_SITE1
   list seq 10 refname SITE2 traffic-class access-list MY_APP filter FILTER_SITE2
```

PfR definizione di gruppi, aka services class

- Con il processo di learn list, si possono creare differenti gruppi di traffico per classificare gruppi audio e video, gruppi business application e gruppi di traffico in best effort
- Tipicamente questo modello di classificazione utilizza DCSP value

```
pfr master
mc-peer domain 65000 head-end Loopback0
target-discovery responder-list RESPONDER_PREFIX inside-prefixes HQ_PREFIX
logging
!
border 10.10.10.4 key-chain pfr
  interface Ethernet0/1 external
  interface Ethernet0/0 internal
!
border 10.10.10.5 key-chain pfr
  interface Tunnel200 external
  interface Ethernet0/0 internal
```

PfR definizione di gruppi, aka services class

- Automatic Learning ; Define 3 Service Classes ; Deny global learning

```
learn
 throughput
 traffic-class filter access-list DENY_GLOBAL_LEARN_LIST
!
list seq 10 refname LEARN_VOICE_VIDEO
 traffic-class access-list VOICE_VIDEO filter SITE_PREFIX
 aggregation-type prefix-length 24
 throughput
!
list seq 20 refname LEARN_CRITICAL
  traffic-class access-list CRITICAL filter SITE_PREFIX
  aggregation-type prefix-length 24
  throughput
!
list seq 30 refname LEARN_BEST_EFFORT
    traffic-class access-list BEST_EFFORT filter SITE_PREFIX
    aggregation-type prefix-length 24
    throughput
```

PfR definizione di gruppi, aka services class

- HQ and BRANCH Prefixes - HQ Responder

```
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
!
ip prefix-list SITE_PREFIX seq 5 permit 20.9.0.0/16
ip prefix-list SITE_PREFIX seq 10 permit 20.10.0.0/16
!
ip prefix-list RESPONDER_PREFIX seq 5 permit 10.2.2.2/32
```

PfR definizione di gruppi, aka services class

- ACL per la classificazione

```
ip access-list extended VOICE_VIDEO
permit ip any any dscp ef
!
ip access-list extended CRITICAL
permit ip any any dscp af31
!
ip access-list extended BEST_EFFORT
permit ip any any dscp default
!
ip access-list extended DENY_GLOBAL_LEARN_LIST deny ip any any
```

PfR misurazione in modalità attiva e passiva

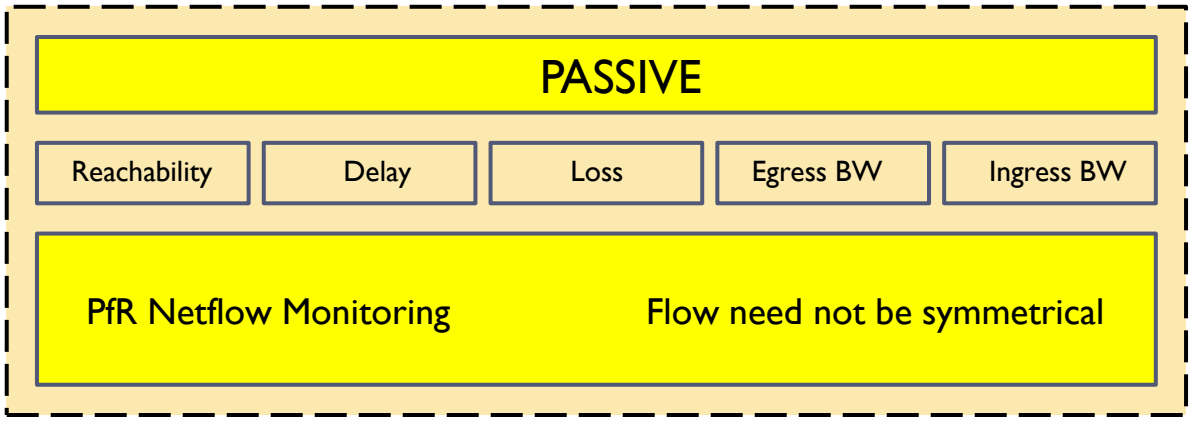

Measurements of the TCP traffic flows are characterized by:

✓ Delay: Time between TCP SYN and TCP ACK in a TCP three-way handshake.

✓ Loss: TCP sequence numbers are tracked, loss can estimated when lower sequence numbers than the highest sequence number observed are seen.

✓ Reachability: Repeated TCP SYNs without an accompanying TCP ACK identify reachability failures.

✓ Throughput: Throughput is calculated from NetFlow and measured in bits per second (bps). Measurements of non-TCP traffic flows is characterized by throughput only.

PfR misurazione in modalità attiva

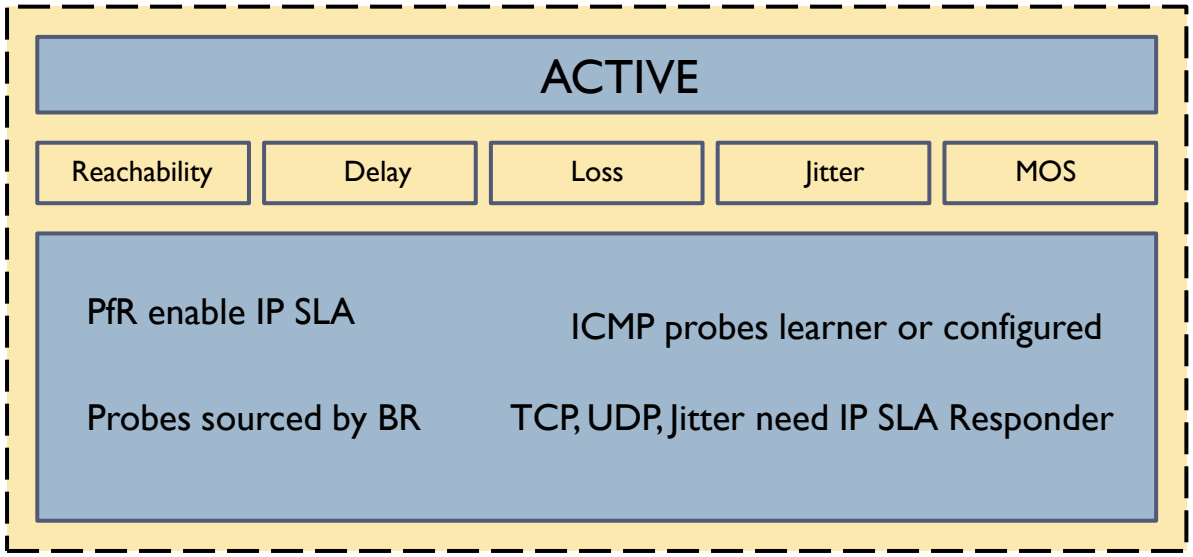


```
set active-probe jitter 30.1.0.11 target-port 33033 codec g729a
set probe frequency 2
```

In this mode, IP SLAs probes are generated by the border routers through the current exits and transmitted at the configured probe frequency value. Probes are only generated through alternate paths (exits) in the event the current path is out-of-policy.

The ICMP ECHO requests that are generated by default constitute additional background traffic on the network. When used on the Internet, activating probing may not be desirable in that ICMP packets may be blocked or administratively prohibited and may be considered a threatening or abusive posture to the target hosts.

By default, an active probe is of the type of ICMP echo. If VoIP is to be characterized, the network manager may choose to explicitly configure an active probe. Following is an example from an *pfr-map* using a traffic-class that matches VoIP streams with probes sent every 2 seconds.

PfR misurazione in modalità ibrida

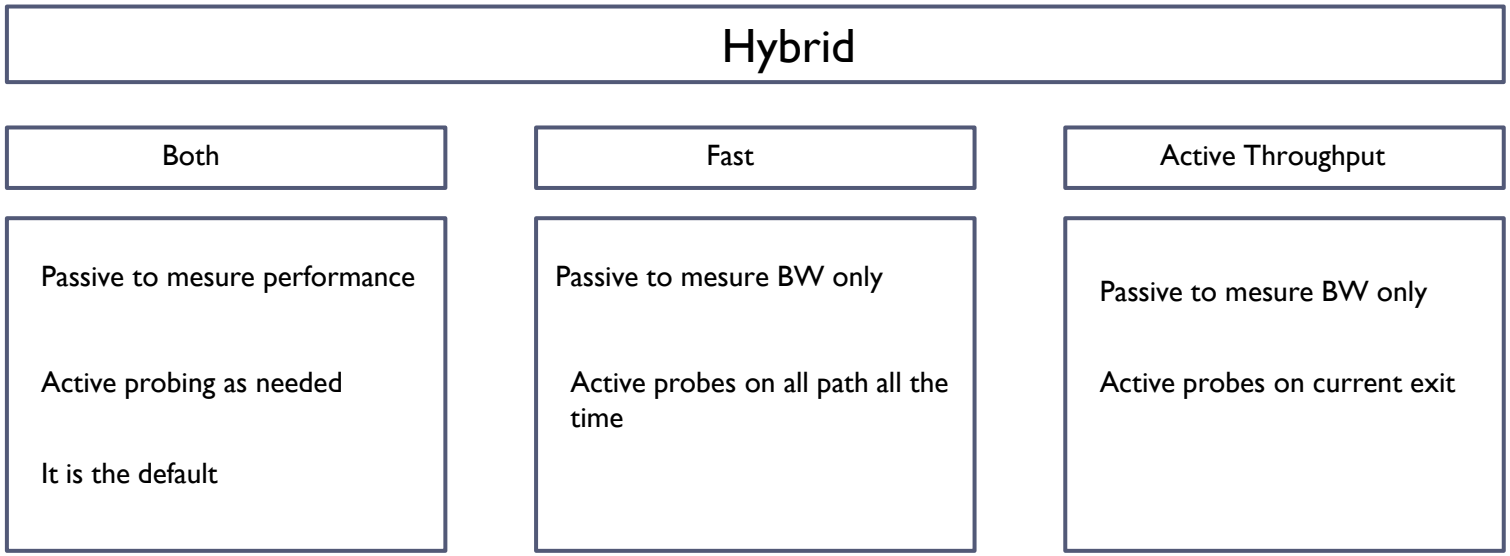

Mode monitor both is the default value and combines the capabilities of passive and active monitoring. Up to five IP addresses are actively probed for each destination prefix network learned through passive monitoring. By default, an IP SLA ICMP ECHO probe is automatically generated for the learned IP addresses.

In this mode, IP SLAs probes are generated by the border routers through the current exits and transmitted at the configured (or default) probe frequency value. Unlike mode monitor fast, active probing does not probe all exit points continuously. It probes only the current exit point provided the status is INPOLICY and probes are generated after the prefix timer value is exhausted. Probes are only generated through alternate paths (exits) in the event the current path is out-of-policy.

By monitoring both actively and passively, additional data points regarding a network prefix can be obtained through two separate and distinct tools; Netflow for passive measurements and IP SLA for the active measurements. However, the inclusion of active probing also has disadvantages.

Mode monitor both is best suited for use within the private internal network of the enterprise.

Important note: With mode monitor both, passive measurements are used to trigger the Out of Policy (OOP) state but the active probing results from the exit interfaces are used with the passive throughput measurement for the traffic class, to select the exit interface.

PfR monitor fast

This feature was introduced in Cisco IOS Release 12.4(15)T as a key component to the Fast Reroute feature.

This mode generates active probes through all exists continuously at the configured probe frequency. This differs from either active or both modes in that these modes only generate probes through alternate paths (exits) in the event the current path is out-of-policy.

With Fast Reroute, the characteristics of the alternative paths are always known, allowing immediate use as required. If unreachable is determined to be out-of-policy for the current exit, the alternate exit that is in policy is selected as the new current exit.

The unreachable threshold is calibrated in number of failed probes per million probe attempts. If the unreachable value is set to 1, a single probe fails on the current exit, an attempt is made to locate a alternate exit. However, if the alternate exits also have a single failed probe, they are not selected because they too are out-of-policy.

The Fast Reroute feature, therefore, allows rerouting actions to be taken, at an interval approaching the configured probe frequency value.

Probe frequency can now be set as low as 2 seconds if fast mode is configured.

This allows re-routing at slightly more than the configured probe frequency value. The Fast Reroute feature can reroute OOP traffic in as little as 3 seconds.

The obvious drawback to this feature is the potential for adding additional network traffic overhead associated with the probes themselves and additional CPU resources to the PfR border routers, the source of the active probes.

Probes are always generated unless the prefix is deleted or in the default state.

The active probe results are used for out-of-policy and to control routing. Passive data collected is for information only, the throughput transmit and receive Kbps values (show oer mast border detail) are used for load balancing.