Contents

Introduction

This document describes the “max-range-utilization” component of the Performance Routing (PfRv2) and its implication on load balancing over multiple WAN links.

Prerequisites

Requirements

Cisco recommends that you have basic knowledge of Performance Routing (PfR).

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, make sure that you understand the potential impact of any command.

Configure

PfR allows network administrators to minimize bandwidth costs, enable intelligent load distribution, improve application performance, and deploy dynamic failure detection at the Wide Area Network (WAN) access edge. Whereas other routing mechanisms can provide both load sharing and failure mitigation, Cisco IOS PfR makes real-time routing adjustments based on criteria other than static routing metrics such as response time, packet loss, jitter, path availability, traffic load distribution, and cost minimization.

For Load balancing, PfR uses the following components:

1. Link Utilization: PfR keeps checking the utilization of the link and depending on the value set in the policy, a decision is taken to distribute the load from one link to the other. PFR also switch back the traffic flow from the secondary to the primary link when it sees the link utilization of the primary link has gone below a specified value.

2. Range: To specify the range of link utilization among the WAN links after which the PfR will apply the policy, PfR uses “max-range-utilization” component of the Performance Routing (PfRv2). The range functionality allows the network administrator to instruct Cisco PfR to keep the usage on a set of exit links with in a certain percentage range of each other. If the difference between the links becomes significant, Cisco PfR will attempt to bring the link back in to policy by distributing data traffic among the available exit links.

3. Traffic Class(TC) Performance: This enables customers to define multiple paths that a set of traffic (for example voice traffic) could use as long as all the paths maintain the performance SLA’s that are needed. 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.

Network Diagram

Following image would be used as a sample topology for rest of the document:

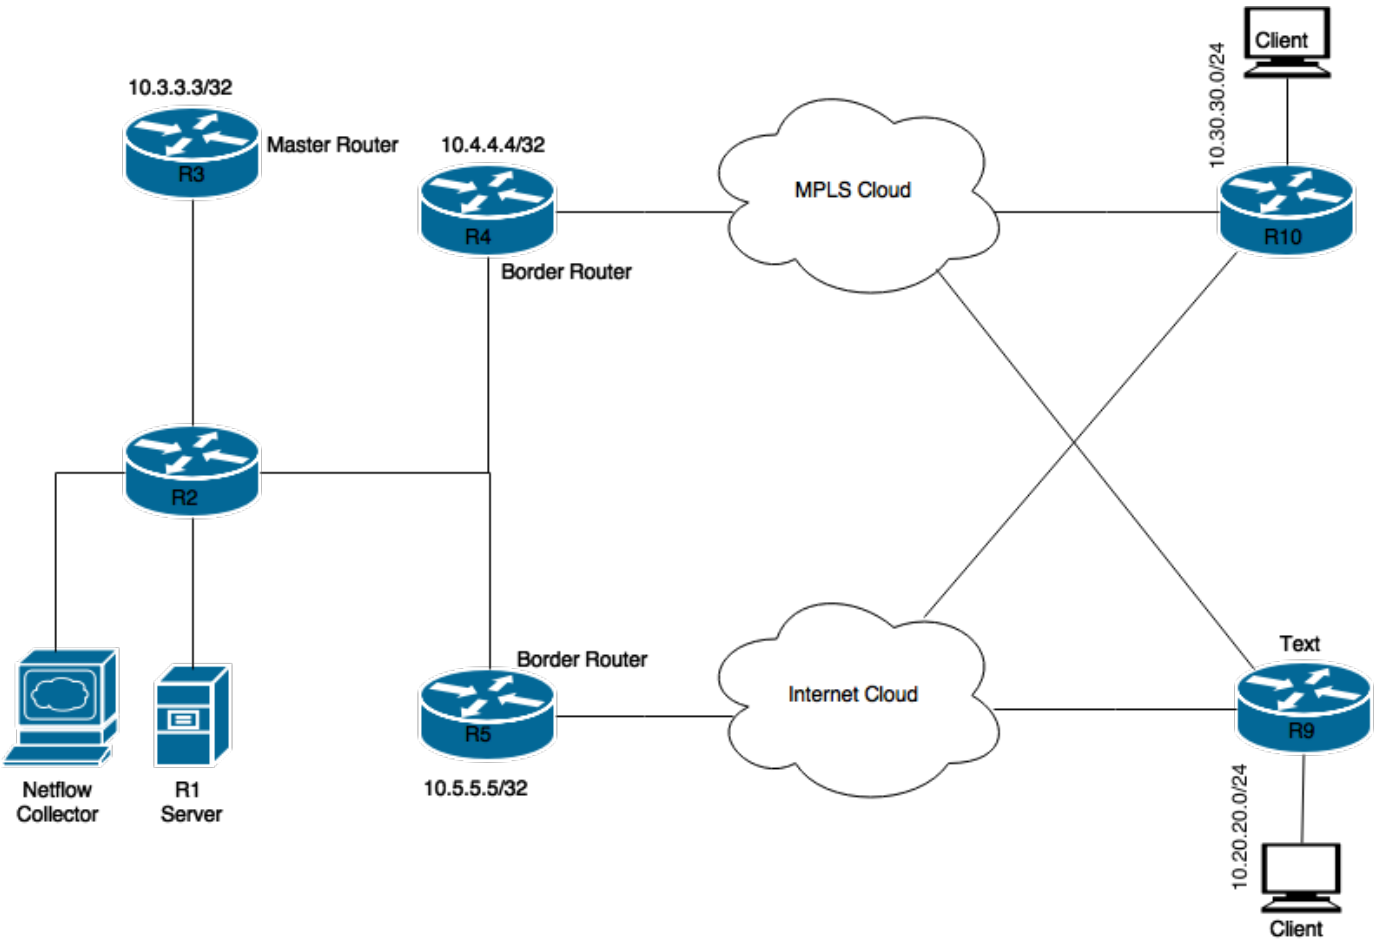

Devices shown in the diagram:

R1 Server: Initiates traffic.

R3: PfR Master Router.

R4 & R5: PfR Border Router.

Clients connected to R9 & R10 are devices receiving the traffic from the R1 server.

Relevant Configuration

R3 (Master Router)

```
hostname R3
!
!
key chain pfr
key 0
key-string cisco
!
!
pfr master
max-range-utilization percent 7
!
border 10.4.4.4 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.5.5.5 key-chain pfr
interface Ethernet0/0 internal
interface Ethernet0/1 external
!
!
interface Loopback0
ip address 10.3.3.3 255.255.255.255
!
```

R4 (Border Router)

```
hostname R3
!
!
key chain pfr
key 0
key-string cisco
!
!
pfr master
max-range-utilization percent 7
!
border 10.4.4.4 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.5.5.5 key-chain pfr
interface Ethernet0/0 internal
interface Ethernet0/1 external
!
!
interface Loopback0
ip address 10.3.3.3 255.255.255.255
!
```

R5 (Border Router)

```
hostname R3
!
!
key chain pfr
key 0
key-string cisco
!
!
pfr master
max-range-utilization percent 7
!
border 10.4.4.4 key-chain pfr
interface Ethernet0/1 external
interface Ethernet0/0 internal
!
border 10.5.5.5 key-chain pfr
interface Ethernet0/0 internal
interface Ethernet0/1 external
!
!
interface Loopback0
ip address 10.3.3.3 255.255.255.255
!
```

Verify

R3 (Master Router) has been configured to keep sending traffic for all traffic classes to selected BR till the the traffic load difference between the two BRs is or above 7%.

```
R3#show pfr master
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 3949
Version: 3.3
Number of Border routers: 2
Number of Exits: 4
Number of monitored prefixes: 2 (max 5000)
Max prefixes: total 5000 learn 2500
Prefix count: total 2, learn 2, cfg 0
PBR Requirements met
Nbar Status: Inactive
Auto Tunnel Mode: Off
Border Status UP/DOWN AuthFail Version DOWN Reason
10.4.4.4 ACTIVE UP 00:02:43 0 3.3
10.5.5.5 ACTIVE UP 00:02:43 0 3.3
Global Settings:
max-range-utilization percent 7 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
no logging
exit holddown time 60 secs, time remaining 0
```

When traffic flow is started from the server R1, on PfR master below traffic classes get created automatically:

```
R3#show 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 (percent/10000), 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
--------------------------------------------------------------------------------
10.20.20.0/24               N    N    N           N           N N
                          INPOLICY      @69          10.4.4.4 Et0/1           BGP
              U        U        0        0        0        0       49        1
              U        U        0        0        N        N        N        N
10.30.30.0/24               N    N    N           N           N N
                          INPOLICY      @69          10.4.4.4 Et0/1           BGP
              U        U        0        0        0        0        1        0
              U        U        0        0        N        N        N        N
```

As shown above, for destination prefixes, 10.20.20.0/24 and 10.30.30.0/24, the status is in INPOLICY which signifies that PfR is controlling the traffic flow for these prefixes and the exit is Border router 10.4.4.4.

Below output taken on PfR master showing link utilization on Border routers WAN link:

```
R3#show pfr master border detail
Border           Status                UP/DOWN             AuthFail  Version DOWN Reason
10.4.4.4         ACTIVE                UP       06:12:46          0  3.3
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      1000          900       106      10 UP                    4
                Rx                   1000         0       0
--------------------------------------------------------------------------------
Border           Status                UP/DOWN             AuthFail  Version DOWN Reason
10.5.5.5         ACTIVE                UP       06:12:46          0  3.3
Et0/0           INTERNAL              UP
Et0/1           EXTERNAL              UP

External            Capacity      Max BW   BW Used    Load Status          Exit Id
Interface            (kbps)       (kbps)    (kbps)    (%)
---------           --------      ------   ------- ------- ------           ------
 Et0/1           Tx      1000          900         0       0 UP                    1
                Rx                   1000         0       0
```

Above output shows all traffic going through R4 and external links ethernet0/1's load percentage is 10% and on R5 it is 0% as of now. With the above configuration in place, PfR should act and distribute some of the load on R5's currently unused WAN link.

After sometime you could stream for 10.30.30.0/24 destination has migrated to new exit:

```
R3#  show 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 (percent/10000), 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
-------------------------------------------------------------------------------
10.20.20.0/24              N    N    N           N           N N
                         INPOLICY        0          10.4.4.4 Et0/1           BGP
             U        U         0        0        0        0       32        0
            16       16         0        0        N        N        N        N
10.30.30.0/24              N    N    N           N           N N
                         INPOLICY        0          10.5.5.5 Et0/1           BGP
             U        U         0        0        0        0       32        1
             U        U         0        0        N        N        N        N
```

Real time load utilization on border routers external interfaces can also be seen below:

```
R3#show pfr master border detail
Border           Status                UP/DOWN             AuthFail  Version DOWN Reason
10.4.4.4         ACTIVE                UP       06:38:45          0  3.3
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      1000          900        52       5 UP                    4
                Rx                   1000         0       0
-------------------------------------------------------------------------------------
Border           Status                UP/DOWN             AuthFail  Version DOWN Reason
10.5.5.5         ACTIVE                UP       06:38:45          0  3.3
Et0/0           INTERNAL              UP
Et0/1           EXTERNAL              UP

External            Capacity      Max BW   BW Used    Load Status          Exit Id
Interface            (kbps)       (kbps)    (kbps)    (%)
---------           --------      ------   ------- ------- ------           ------
 Et0/1           Tx      1000          900        51       5 UP                    1
                Rx                   1000         0       0
```

Note: In above example equal load distribution on Border routers is seen but it is possible to have unequal load sharing in production setups.