

Persistent Route Oscillation in BGP Constrained Route Distribution  
draft-idr-bgp-rt-oscillation-01


IETF 98, March 2017  
Chicago

Dongling Duan (Cisco)  
Jakob Heitz (Cisco)  
Keyur Patel (Arrcus)  
Jeff Haas (Juniper)

Best Path  
Other Path

Bestpath due  
to metric

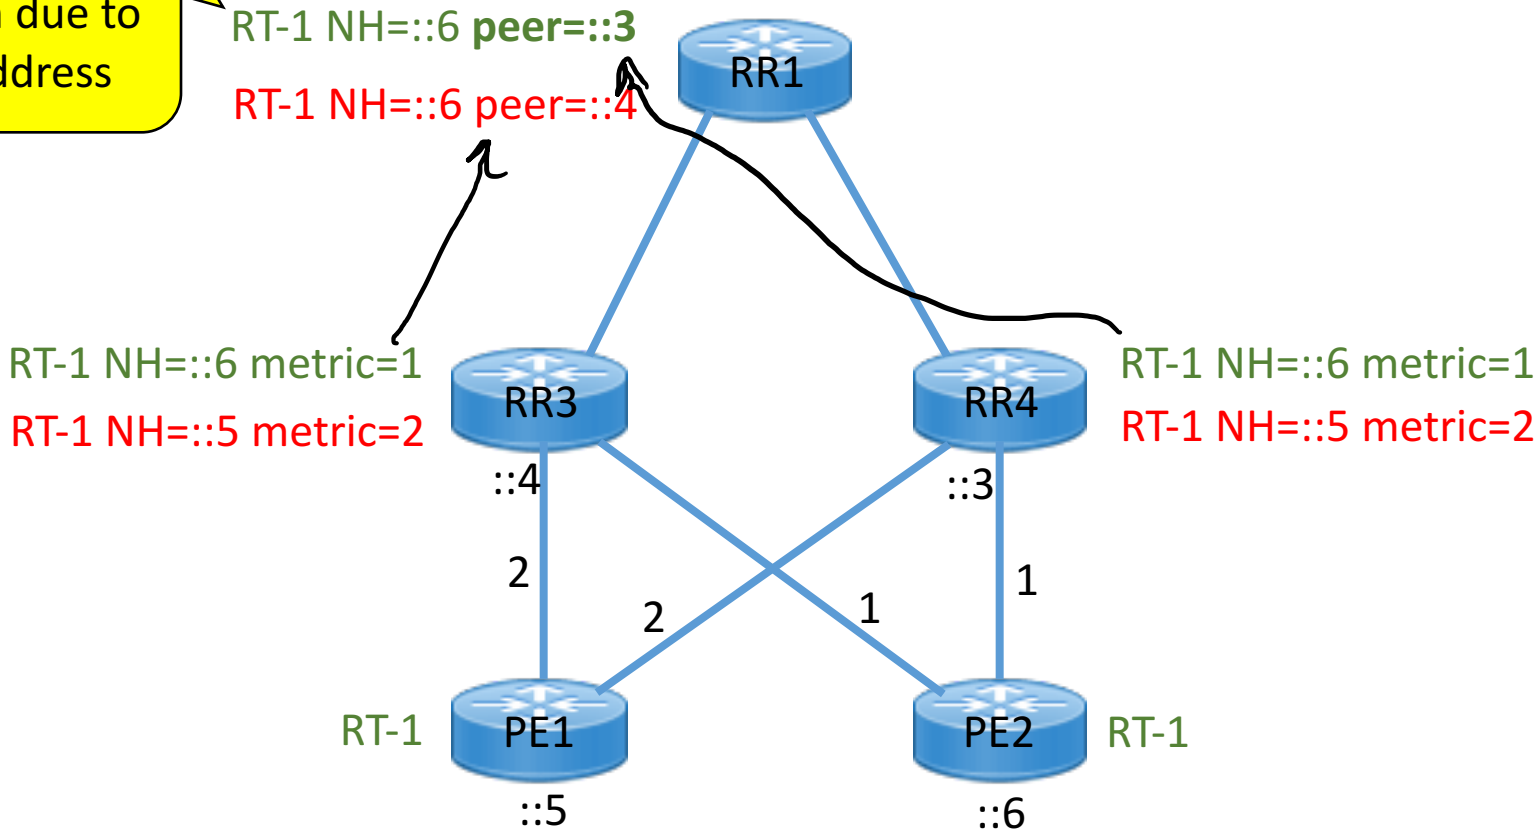
RT-1 NH>:::6 **metric=1**  
RT-1 NH>:::5 **metric=2**



RT-1 NH>:::6 **metric=1**  
RT-1 NH>:::5 **metric=2**

Metric is equal,  
bestpath due to  
peer address

Best Path  
Other Path




Best Path  
Other Path

RT-1 NH>:::6 peer>:::3  
RT-1 NH>:::6 peer>:::4

RT-1 NH>:::6 m=1 orig>:::1 peer>:::1  
RT-1 NH>:::6 m=1 orig>:::6 peer>:::6  
RT-1 NH>:::5 m=2 orig>:::5 peer>:::5

RT-1 NH>:::6 m=1 orig>:::1 peer>:::1  
RT-1 NH>:::6 m=1 orig>:::6 peer>:::6  
RT-1 NH>:::5 m=2 orig>:::5 peer>:::5



Choose this client path at random

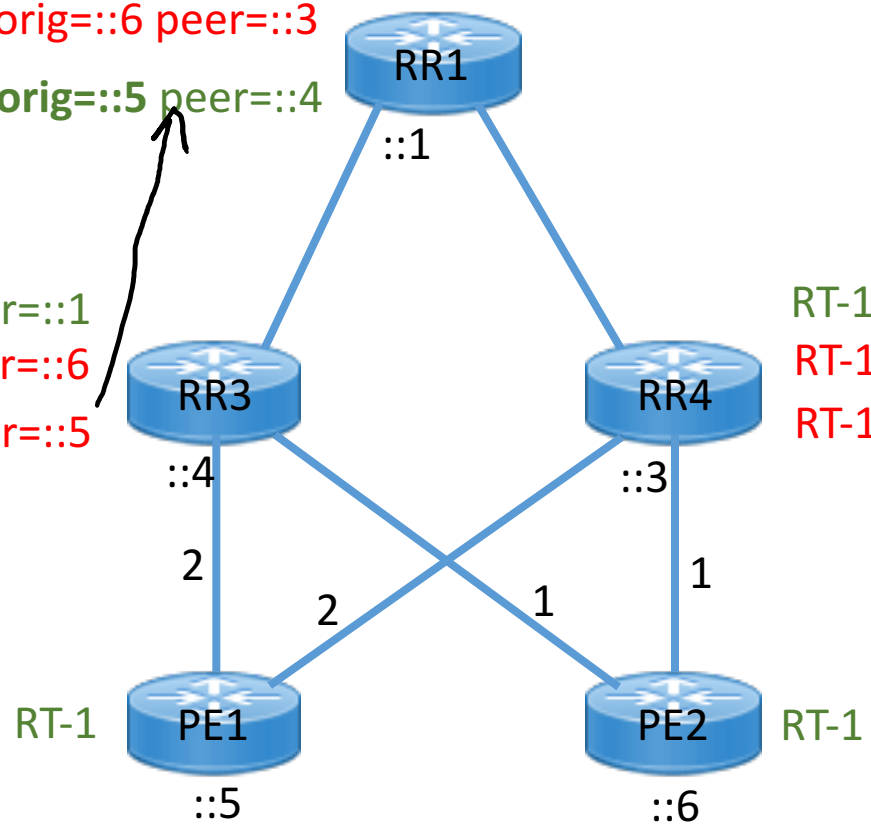
This replaces the previous path from peer ::4

RT-1 NH>:::6 orig>:::6 peer>:::3  
RT-1 NH>:::5 orig>:::5 peer>:::4

RT-1 NH>:::6 m=1 orig>:::1 peer>:::1  
RT-1 NH>:::6 m=1 orig>:::6 peer>:::6  
RT-1 NH>:::5 m=2 orig>:::5 peer>:::5

RT-1 NH>:::6 m=1 orig>:::1 peer>:::1  
RT-1 NH>:::6 m=1 orig>:::6 peer>:::6  
RT-1 NH>:::5 m=2 orig>:::5 peer>:::5

Best Path  
Other Path



Drop this path  
due to cluster  
id in list


RT-1 NH>:::6 orig>:::6 peer>:::3

RT-1 NH>:::5 orig>:::5 peer>:::4

RT-1 NH>:::5 m=1 orig>:::1 peer>:::1

RT-1 NH>:::6 m=1 orig>:::6 peer>:::6

RT-1 NH>:::5 m=2 orig>:::5 peer>:::5



RT-1 NH>:::5 m=1 orig>:::1 peer>:::1

RT-1 NH>:::6 m=1 orig>:::6 peer>:::6


RT-1 NH>:::5 m=2 orig>:::5 peer>:::5

Best Path  
Other Path

# The Problem

RFC 4684: RT Constrain, Section 3.2

ii. When advertising an RT membership NLRI to a non-client peer, if the best path as selected by the path selection procedure described in [Section 9.1](#) of the base BGP specification [[4](#)] is a route received from a non-client peer, and if there is **an** alternative path to the same destination from a client, the attributes of **the** client path are advertised to the peer.



When there's more than one client path?

# The Solution

The solution is for the Route Reflector always to prefer the client paths when selecting a best path. This preference **MUST** be expressed before step f) of the BGP Decision Tie Breaking rules in Section 9.1.2.2 of [RFC4271]. It **MAY** be expressed at a higher step.