



nexthop™

Persistent Route Oscillation: Issues and Solutions

NANOG 26

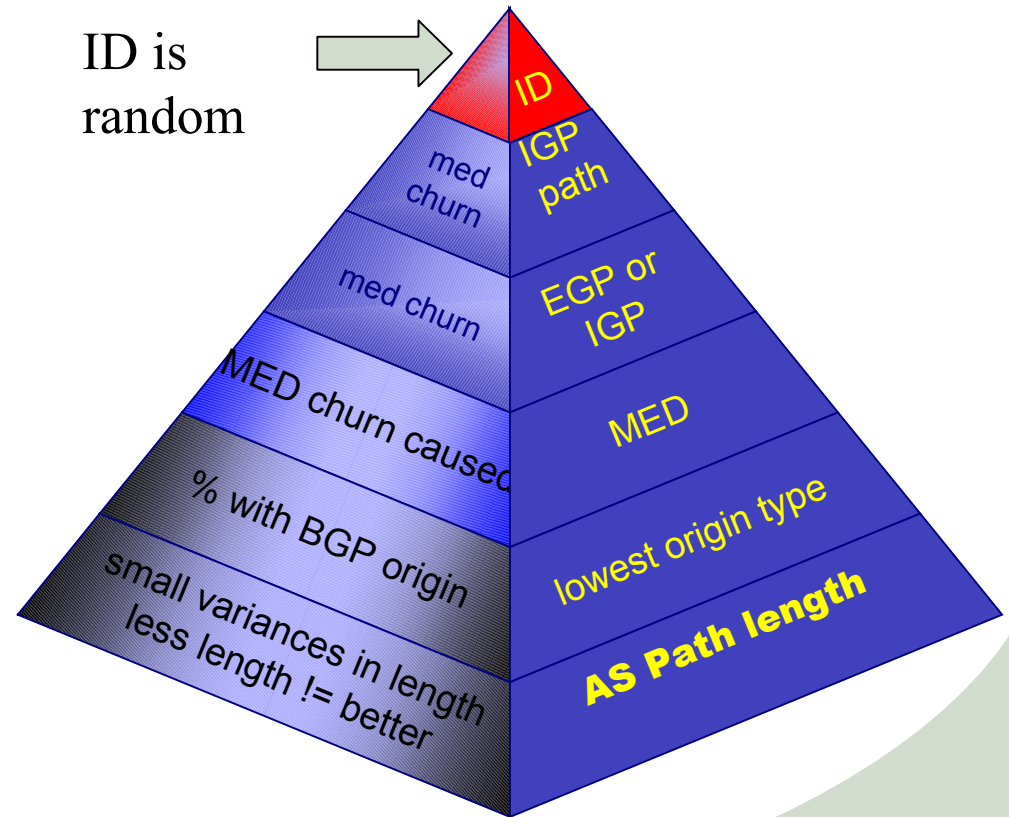
- Route Oscillation Problem - Sue Hares
(skh@nexthop.com)
- BGP Route Oscillation Reduction – Single Path Approach Enke Chen (enke@redback.com)
- Multipath for Oscillation - John Scudder
(jgs@cisco.com)
- BGP Route Oscillation Detection and Deployment Considerations – Enke Chen

Route Oscillations

- What is not being solved
- Goals for Solutions
- Analysis of persistent route oscillation
 - Without MEDs
 - With MEDs
- Route oscillation reduction
 - Advertising the best external route
 - Avoid unnecessary route withdraws

What is not being solved

- Consistent route selection
 - Good news it is done!
 - Draft-18 of BGP will have the rules consistent with deployment
- MED knobs
 - bgp-always-compare-med
 - bgp- bestpath med-confed
 - Bgp-deterministic med



tie breaking after policy

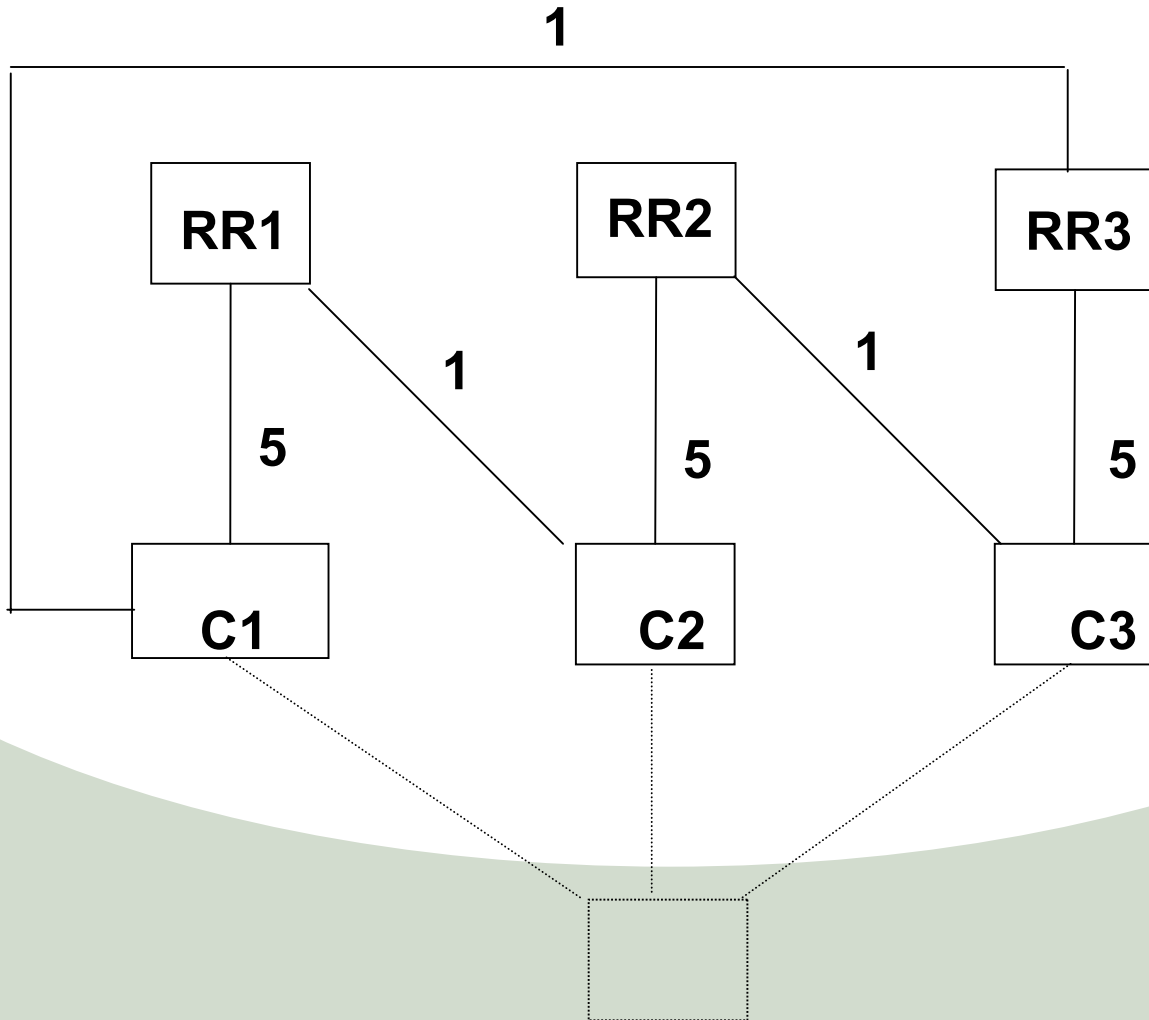
Goals for Solutions

- BGP should be stable and converge
- Less Constraints on the topology for
 - Route Reflection
 - AS Confederations
 - IGP metrics
- BGP needs to be Robust



Route Oscillation – No MEDs

Physical Topology






Route Oscillation – No MEDs

- Consider the following interactions:
 - C1, C2 and C3 advertises the EBGP routes to their respective RR
 - RR1, RR2 and RR3 advertise client's routes to each other
 - RRs re-compute best path and then withdraw their advertisement of client's routes
- The route withdraw results in the complete removal of the routing information from other clusters
- This over reduction (compared with full-mesh) of routing information is the culprit for route oscillation in this case

Route Oscillation – MEDs

Physical Topology

3rd: Tie break
Rc and Re
(same adjacent
AS) on MED




MED	10	1	0
ASPATH	10 100	6 100	6 100

Look at RFC 3345 for full example

Route Oscillation – MEDs

Physical Topology

Tie break on Rb and Re on IGP distance, Select RB



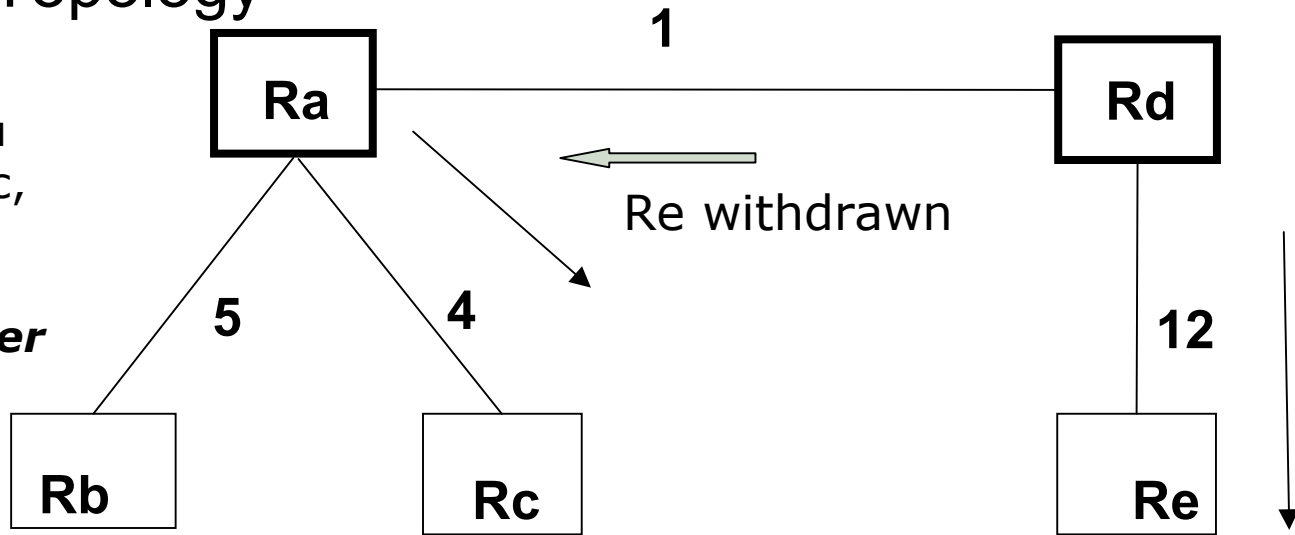
MED	10	1	0
ASPATH	10 100	6 100	6 100

Look at RFC 3345 for full example

Route Oscillation – MEDs

Physical Topology

With Re
withdrawn, you
have Rb and Rc,
Select Rc
..and start over



MED	10	1	0
ASPATH	10 100	6 100	6 100

Look at RFC 3345 for full example

Route Oscillation – MEDs

- Consider the following interactions:
 - Rb, Rc and Re advertises the EBGP routes to their respective RR (Ra and Rd)
 - Ra and Rd advertise client's routes to each other.
 - Ra then re-computes best path and advertises a different path (Rb) to Rd, which would cause Rd to withdraw its advertisement
 - The route withdraw by Rd would force Ra to advertise a different path (Rc)
- Again, route withdraw is the culprit for route oscillation



Major Factors for Route Oscillation

- Dependency of IBGP updates
 - Sometimes circular dependency
- Partial information by RR or confederation
 - Withdraws (over reduction) amplifies the issue
- Partial order (due to MED) in route selection
 - Route selection overly sensitive to route information
 - Seems to play less role than perceived
 - Not much can be done with the current protocol and deployed base

Route Reflection vs Confederation

- In principle they can be equivalent
 - Treat AS Border Routers as RRs
 - Treat IBGP as clients
 - Treat C-EBGP as non-client IBGP
- But in practice, Confederations are almost always deployed as a two-deep hierarchy
 - Because the ASBRs are not fully meshed
 - This means the probability of oscillation is higher (more hierarchy = more probability)

Look at RFC 3345 for full example



- Route Oscillation Problem - Sue Hares
(skh@nexthop.com)
- BGP Route Oscillation Reduction – Single Path Approach Enke Chen (enke@redback.com)
- Multipath for Oscillation - John Scudder
(jgs@cisco.com)
- BGP Route Oscillation Detection and Deployment Considerations – Enke Chen

Questions