

Deterministic Route Redistribution into BGP

draft-chen-bgp-redist-05.txt
IETF-120

Enke Chen (enchen@paloaltonetworks.com)
Jenny Yuan (jyuan@paloaltonetworks.com)

Problem Statement (Recap)

- **Uncorrelated route selection** by RIB and BGP
 - RIB: Admin-distance
 - BGP: LOCAL_PREF, AS-PATH, et al
- When redistributing a backup route into BGP
 - Non-deterministic results - depending on the order of path arrivals
- Single router and multiple-routers in a network
 - Detailed in the draft
- Day-1 issue
 - **Lack of specification in BGP**
 - Issue exists in a number of implementations
 - Requires configuration **if one is aware of** the issue and solution
 - Can be frustrating in network operation

Proposed Solution (Recap)

- Make it deterministic among paths with known admin-distances
 - Locally redistributed, EBGP, and aggregate
 - **Compare the admin-distance first in BGP route selection**
- What about IBGP routes (network-wide)?
 - Use the LOCAL_PREF for network-wide influence
 - Lower the LOCAL_PREF for redistributed backup routes
- Configuration option
 - Can be used to override the automatic setting

Updates in -05

- Addressed comments from the previous WG Adoption
- No change to the admin-distance algorithm
- Simplified the solution for the network-wide issues
 - Removed automatic mapping from "admin-distance" to LOCAL_PREF
 - Primary/backup working automatically, and config for more complex scenarios
- Enhanced backward compatibility
 - Evaluate config, and then adjust the LOCAL_PREF **only if needed and safe**
-
- Posted to IDR WG a reference implementation in FRR

Next Steps

- Request IDR WG Adoption