

Deterministic Route Redistribution into BGP

draft-chen-bgp-redist-03.txt
IDR WG Interim, 08/2021

Enke Chen
Jenny Yuan

Goals of the Draft

- Explain and document the non-deterministic behavior
 - Not a well-known issue
- Propose a solution
- Promote implementation in software
 - Configuration remains an option

Problem Statement

- Uncorrelated route selection algorithms 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 / network-wide
 - Detailed in the draft
- Day-1 issue
 - “Under-specification” in BGP
 - Exists in a number of implementations (but not all?)
 - Requires configuration if one is aware of the issue and solution
 - Can be frustrating in network operation

Proposed Solution

- Make it deterministic among paths with known admin-distances
 - Locally redistributed, EBGP, and aggregate
 - Compare the admin-distance 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
- Compute LOCAL_PREF based on the admin-distance
 - Rank the backup routes in the network, and assign “offsets”
 - Admin_distance: $\text{ibgp_admin_distance} + \text{offset}$
 - LOCAL_PREF: $\text{default_local_pref} - \text{offset}$
- Configuration option
 - Can be used to override the automatic setting

Software Option vs Configuration Option

- Recommend implementation in software
 - Make the default behavior deterministic and work
 - Less configuration
- Use config until the algorithm is implemented in the software
 - Manually set the weight (if applicable) and LOCAL_PREF
 - Can also be used to override the automatic setting