

Deterministic Route Redistribution into BGP

draft-chen-bgp-redist-02.txt
IETF 111, on-line

Enke Chen
Jenny Yuan

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 (based on the admin-distance) 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
 - 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

- 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: $ibgp_admin_distance + offset$
 - LOCAL_PREF: $default_local_pref - offset$
- Configuration option
 - Can be used to override the automatic setting

Next Steps

- Welcome feedback
 - In particular the implementation perspectives
- Revision (-03)
 - Further clarifications
- Request IDR WG Adoption post -03 version