RPKI Notes

Refer to RPKI-Based Policy Without Route Refresh for context.

Basically BGP implementations should/must not send a route refresh when receiving updated RPKI data, and are recommended instead to retain the prefix marked as invalid should the future RPKI state change.

Also presented at RIPE 83 for additional background and context.

It has been noted by several operators that their Cisco routers implementing ROV were bombarding peers with Route Refresh requests. This is challenging for those routers which are “control plane challenged” and can be construed as a denial of service on those peering routers.

ROV

The following table documents ROV behaviours on receipt of updated RPKI information from validators.

“Adj-RIB-In” is the BGP table as received from BGP peers, prior to processing by inbound policy. Retaining this BGP table requires extra memory (not a hardship in this day and age), and makes process BGP policy changes simple. Without Adj-RIB-In, the router has to send a Route Refresh to the peer to request all BGP updates again. Which can be exciting when today's IPv4 table is heading to 900k prefixes, and IPv6 table is heading to 150k prefixes.

Implementation	Adj-RIB-In	ROV behaviour	Notes
Cisco IOS-XE	No	VRP update triggers a route-refresh	Workaround is to turn on “soft-reconfiguration in”
Cisco IOS-XR	No	VRP update triggers a route-refresh	Workaround is to turn on “soft-reconfiguration in”
Juniper JunOS	Default	VRP update handled locally	Adj-RIB-In can be turned off by “set protocol bgp group keep none” as described here
Bird 2.0.8	?	handles VRP updates locally	“rpki reload on” is default in 2.0.8 as described here
Arista EOS	Default	VRP updated handled locally	Adj-RIB-In can be turned off
FRR 8.1	?	?	?

Back to Home page

From:
https://www.bgp4all.com/pfs/ - **Philip Smith's Internet Development Site**

Permanent link:
https://www.bgp4all.com/pfs/rpki?rev=1637798130

Last update: **2021/11/24 23:55**