

Revising BGP Maximum Prefix Limits

Job Snijders, NTT, job@ntt.net

Melchior Aelmans, Juniper Networks, maelmans@juniper.net

IETF 106, Singapore


Recap

Presented at IETF 104 in GROW


(slides: <https://datatracker.ietf.org/meeting/104/materials/slides-104-grow-bgp-maximum-prefix-limits>)

These “maximum prefix limits” are a design feature to ensure the network inherently responds in a way that will cause no or minimal harm to the network or the global Internet.

What happens when limits are applied in pre-policy during a full table leak:



What happens when limits are applied post-policy



Pre-policy vs Post-policy in EBGP-IN attachment

Pre-policy limits:

Protects against memory exhaustion (keep in mind – pre-policy limit only works if you keep all rejected routes in Adj-RIB-IN)

Protects against route leaks

Post-policy limits:

Protect against RIB & FIB exhaustion

To enforce contractual agreements

Maximum prefix limits in context of **ebgp-in**

Vendor	Pre-Policy (the most effective place)	Post-Policy
Cisco IOS XR	Not available	"maximum-prefix"
Cisco IOS XE	Not available	"maximum-prefix"
Juniper Junos	"prefix-limit"	"accepted-prefix-limit" or "prefix-limit" + "keep none"
Nokia SR-OS	"prefix-limit"	Not available
NIC.CZ's BIRD	"import keep filtered" + "receive limit"	"import limit" or "receive limit"
OpenBSD's OpenBGPD	"max-prefix"	Not available

New: Outbound maximum prefix limits

Was brought up in operational forums – it would be useful to specify how we envision outbound maximum prefix limits should work

An outbound limit would be a “self-destruct” control action, in case you end up announcing far more than is plausible on a given EBGP session

Only BIRD supports this today, hope we can get more to support the feature.

Progress & Goals

Goal: Articulate in specification the concept of “Pre Import Policy” vs “Post Import Policy” in context of Policy Information Base, and outbound limits.

→ To drive consistency across BGP implementations

Insight: previously I rejected “soft limits”, but that was my failing understanding of the core BGP spec wording. So that’s back!

Progress: Moved target working group from GROW to IDR

(Susan Hares kindly articulated how extensive the changes to the core BGP spec are)

Next step

Clean up language, ask for IDR WG adoption

Draft: <https://tools.ietf.org/html/draft-sa-idr-maxprefix-00>

Thanks!!!!