

BGP Maximum Prefix Limits

Job Snijders, NTT Communications, job@ntt.net

Melchior Aelmans, Juniper Networks, maelmans@juniper.net


draft-sa-grow-maxprefix

Maximum Prefix Limits


These 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.

An operator defines a *high water mark*, where the assumption is that when this threshold is reached, something somewhere is wrong – usually a full table route leak.

What happens when limits are applied in Pre-Policy during a full table leak:



What happens when limits are applied Post-Policy



Pre- vs Post-Policy prefix limits in **ebgp-in**

Pre-policy maximum prefix limits:

- Protect against memory exhaustion
 - Keep in mind: a pre-policy limit only works if the router remembers the list of rejected routes
- Protect against (full table) route leaks

Post-policy maximum prefix limits:

- Protect against RIB+FIB exhaustion
- To enforce contractual agreements (Layer-3 VPN context)

Maximum prefix limits @ various vendors

Vendor	Type A Pre-Policy	Type B Post-Policy	Outbound
Cisco IOS XR	Not available	"maximum-prefix"	Not available
Cisco IOS XE	Not available	"maximum-prefix"	Not available
Juniper Junos	"prefix-limit"	"accepted-prefix-limit" or "prefix-limit" + "keep none"	Not available
Nokia SR-OS	"prefix-limit"	Not available	Not available
NIC.CZ's BIRD	"import keep filtered" + "receive limit"	"import limit" or "receive limit"	"export limit"
OpenBSD's OpenBGPD	"max-prefix"	Not available	Not available
Arista EOS	"maximum routes"	"maximum-accepted-routes"	Not available
Huawei VRP v5	"peer route-limit"	Not available	Not available
Huawei VRP v8	"peer route-limit"	"peer route-limit accept-prefix"	Not available

Outbound maximum limits?

This was raised before on nanog@nanog.org – we should work to get *outbound* maximum prefix limits. I'd rather self-destruct than leak a full table to a peering partner.

At least a “self-destruct the session” control action, in case you end up announcing far more than plausible. Implementations could choose additional “soft” control actions (but that is out of scope for this draft).

Only BIRD supports this today. We'd need to standardize this in IETF.

WG Adopt?