Maximum size of the IPv6 routing table

Santiago Lorente - 2020-04-29 - 0 Comments - in BGP

The maximum number of IPv6 routes allowed in the kernel is set at */proc/sys/net/ipv6/route/max_size*. That *max_size* value defines the number of slots that the routing table can contain.

In the case you experience IPv6 BGP peers dropping or IPv6 routes not being installed in the kernel, it might be because a too low value is set as *max_size*.

If you are experiencing that problem, maybe because you are using a large number of interfaces and/or routes, you might want to increase that value.

However, having *max_size* with a very high value might lead to other kind of problems. When increasing the value, the Linux kernel inserts a temporary "cache" entry into the IPv6 route table radix tree, and that uses RAM, so a too high value may take you to RAM exhaustion. When there is no RAM left, the kernel may misbehave.

A good practice could be to set the value just big enough so that you don't have to worry for the next 2 or 3 years. You can set that value based on the current rate of growth of the default-free zone.

You could set the *max_size* value to, for instance, 262144, with the command below:

```
vyos@vyos# set system sysctl custom net.ipv6.route.max_size value
$262144
```

The "cache" entries will be cleaned up when the IPv6 table is full (and also periodically), but there will still be the limit of not being able to install more than *max_size* routes in the kernel — whether they have come from FRR, kernel routes for interfaces or static routes.

Text based on Marek Isalski's 29/04/2020 reports.