Max Prefixes

Section eight of RFC 7454 discusses limiting the prefixes from a peering through the use of a Maximum Prefixes configuration setting. The goal is to protect your network against large increases from peers & upstreams during events in which they suddenly advertise a large number of additional routes.

Examples of large-scale route table leaking are not infrequent and taking action to ensure your routing continues to function optimally during them are well advised. Section eight recommends that, for your peers, you use a maximum prefix threshold lower than "the number of routes in the Internet." This will shut down the peerings during major leakage incidents. For upstream providers its suggested that the limit be set to be more than he number of routes in the Internet, and somewhere near the technical ability of the router to handle the routes. In both cases per-peer setting also also suggested as alternatives, with the setting being based on the current number of routes with some headroom for growth added in.

Warning levels can be set on most platforms, to log messages when the prefixes received go over a certain threshold. For example, "Shut down the BGP session at 10,000 routes and begin logging warning when the session reaches 80% of that number." (Note that the peering DB currently recommends a 95% warning threshold.)

It is absolutely critical that the session limits be monitored, either through the use of the warning log messages or through other means. The limits are meant to protect from unexpected sudden increases. The slow gradual growth of prefixes received, should they eventually grow enough to reach the limits, can cause an unwatched session to shut down. This sort of "false positive" would be an unfortunate outcome of a technology meant to protect.

You may also want to consider what would happen if your session did shut down. If you primarily rely on one provider, or a default provider, perhaps you would NOT want to shut down the session even if they did flood your BGP with routes. You may also want to consider an auto-restart timer, reenabling the session after 15 minutes or so.

And don't forget your IPv6 sessions!

Juniper Example

In the Juniper example below we're setting the IPv6 maximum number of prefixes received, for each member of the group CONNECTORS6, to be 50. Should we receive 51 prefixes the session will be torn down. The "teardown 90" command indicates that the session should be torn down when the limit is reached, and that warnings should be logged when the 90% threshold is reached.

```
Protocols {
  bgp {
   group CONNECTORS6 {
     type external;
      metric-out igp;
      family inet6 {
          unicast {
              prefix-limit {
                  maximum 50;
                  teardown 90;
```

Juniper documentation on the premix-limit can be found at:

https://www.juniper.net/documentation/en_US/junos12.3/topics/reference/configuration-statement/prefix-limit-edit-protocols-bgp.html

Cisco Example

In the Cisco example below the IPv4 session has been configured with a teardown value of 1000 and a warning/logging percentage of 90%.

```
router bgp 65500
 neighbor 192.0.2.1
```

```
    remote-as 65555
    address-family ipv4 unicast
      maximum prefix 1000 90
```

Cisco documentation on Maximum prefix can be found at:

http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/25160-bgp-maximum-prefix.html

Brocade Example

The brocade example sets the teardown threshold at 100,000 and begins issues warnings at 80%.

```
  router bgp
    neighbor 10.0.0.1 maximum-prefix 100000 threshold 80
```

You can find documentation on the Brocade configurations at:

http://www.brocade.com/content/html/en/command-reference-guide/fastiron-08040-commandref/GUID-78FEBC9A-6F1B-4328-B975-E64A3FE8442E.html