Rate Limiting (Draft)

Rate-limiting, in the context RFC 7454, is discussed in the context of protecting the RE from BGP events received at a high rate. The goal being the protection of the control plane of the router and not letting "bad" control traffic impact your desired BGP updates. The RFC also points to RFC 6954 for more information on general control plane protections, but in the words of the RFC:

> *In addition to strict filtering, rate-limiting MAY be configured for accepted BGP traffic. Rate-limiting BGP traffic consists in permitting only a certain quantity of bits per second (or packets per second) of BGP traffic to the control plane. This protects the BGP router control plane in case the amount of BGP traffic surpasses platform capabilities.*

The topic contains significant complexity. Setting the limit too low may impact convergence and stability. Setting the limits too high can effectively eliminates the benefits. Complicating this is a lack of good material on what effective settings, or guidelines in general. In other words, while there is academic agreement that this should be an effective tool, pragmatically there is little advice available on what settings to use in which situations.

If a device is already blocking unauthorized BGP speakers, say through the use of a dynamic filter, in combination with other techniques such as GTSM and internal border spoofing protections, then the issue may be moot as the other techniques offer significantly more protections. Newer router code from vendors also includes some protection schemes, on by default, that help protect the control plane from becoming overloaded.

Juniper Example

A Juniper example follows. A term is defined in the loopback ingress filter that allows specifies that only certain traffic types from BGP peers (and MSDP peers in this case) through at a 500k drop limit, which is defined in a separately configured policer policy. Note that the 500k limit is only a placeholder; there's no assertion that this is the correct setting.

```
term limit-tcp-syn {

  from {

    source-prefix-list {

    BGP-PEERS;

    MSDP-PEERS;                    }

     protocol tcp;

     tcp-flags "(syn & !ack) | fin | rst";              }                  then {

    policer 500K-drop;

    next term;                }
```

```
 policer 500K-drop {

  if-exceeding {

    bandwidth-limit 500k;

     burst-size-limit 62k;        }

  then discard;    }
```

Juniper has an entire series of articles available that details the (default) internal DDOS protections for their RE's. It's worth becoming familiar with the default RE protections Juniper provides. It is at:

https://www.juniper.net/documentation/en_US/junos12.3/information-products/pathway-pages/config-guide-ddos/ddos-protection.html

There's also a nice series of articles by Saku Ytti on the ddos protection systems in the Juniper. Here's a window in to one of them:

http://blog.ip.fi/2014/03/quick-look-at-trio-ddos-protection-with.html

Cisco Example

The Cisco example uses IOS-XR for Local Packet Transport Services. Hardware policers on the line cards limit traffic sent up the stack. Packets per second are defined for three types of BGP traffic. bgp-cfg-peer are new sessions and sessions that are not yet 'Established.' bgp-know are established sessions and bgp-default is the 'default' entry for bgp traffic; things not falling in to other buckets, etc.

```
lpts pifib hardware police
 flow bgp-cfg-peer 2000
 flow bgp-default 2500
 flow bgp-known 1500
```

Cisco has a feature set called Control Plane Policing, or COPP, as well as Control Plane Protection, CPPr.

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst6500/ios/12-2SX/configuration/guide/book/copp.html

There are also articles about LPTS, such as:

http://www.cisco.com/c/en/us/td/docs/routers/asr9000/software/asr9k_r4-2/addr_serv/configuration/guide/b_ipaddr_cg42a9k/b_ipaddr_cg42a9k_chapter_0111.html

Cisco has several good summary article of DDOS that describes various DDOS and protection techniques.

http://www.cisco.com/c/en/us/about/security-center/guide-ddos-defense.html

http://www.cisco.com/c/en/us/about/security-center/copp-best-practices.html

Brocade Example

Brocade has a variety of solutions for filtering traffic via ACL's to the local CPU, as well as protecting the CPU from various Layer-2 events, but it has no specific abilities with regard to BGP rates. There are some built in protections against SMURF and TCP SYN attacks. In addition to these, Brocades "Receive ACL's" can offer an uplift in protection to the CPU.

http://www.brocade.com/content/html/en/configuration-guide/NI_05800a_SECURITY/GUID-4DDF53B1-14EB-45AB-9B23-EA7E5A52C2CE.html