nft connlimit

Pablo Neira Ayuso
<pablo@netfilter.org>
NFWS 2018, Berlin, Germany

xt_connlimit

- ratelimit number of connections per src/dst host
- uses conntrack
- Datastructure: hashtable + rbtree + linked list
- Datapath operation:
  - hash(key) => rbtree
  - lookup(key) in rbtree => list
  - count number of conn nodes in that list
    - conn nodes look like ct tuples (no pointer to conntrack objects!)
  - garbage collection of stale objects from packet path
- Key = saddr or daddr
- Key & mask also possible

nf_conncount

- Datastructure + API (Florian)
  - data = init(family, keylen)
  - destroy(data)
  - count(data, key, ct_tuple, ct_zone)
- List API (Pablo)
  - lookup(list, ct_tuple, ct_zone, &addit)
  - add(list, ct_tuple, ct_zone)
  - free(list)

nft connlimit support

- New extension: nft_connlimit
  - Stores just a list of conn objects
    - No hashtable + rbtree
  - nft add set x y { type ipv4_addr\; }
  - nft add rule x y ct count over 2 drop
- Combine them with ‘meters’

  nft add rule filter input ct state new tcp dport 22 \
  meter xyz { ip saddr ct count over 2 } counter reject
- Fixed policies (TODO):

  nft add ct counter filter bad-guy1 { over 2 }
  nft add ct counter filter bad-guy2 { over 6 }
  nft add rule filter input ct state new tcp dport 22 \
  ct count name ip saddr map { \
  1.2.3.4 : “bad-guy1”, 2.3.4.5 : “bad-guy2” } reject
- Open issues: New default set size.