Netfilter Updates since 2016

NFWS 2017 - Faro, Portugal

Pablo Neira Ayuso <pablo@netfilter.org>

Git stats since June 2016

- In net/netfilter/
  - 511 commits
  - 83 contributors
    - Pablo ~20% (103 commits)
    - Liping ~20% (100 commits)
    - Florian ~16% (83 commits)
    - Gao Feng ~6% (30 commits)

Git stats since June 2016

- In net/ipv4/netfilter/
  - 90 commits
  - 25 contributors
    - Florian ~26% (24 commits)
    - Liping ~15% (14 commits)
    - Gao Feng ~10% (9 commits)
    - Pablo ~8% (7 commits)

Git stats since June 2016

- In net/bridge/netfilter/
  - 34 commits
  - 19 contributors
    - Pablo ~23% (8 commits)
    - Willem de Bruijn ~9% (3 commits)
    - Joe Perches ~9% (3 commits)

Git stats since June 2016

- In net/ipv6/netfilter/
  - 66 commits
  - 22 contributors
    - Florian ~30% (20 commits)
    - Liping ~12% (8 commits)
    - Pablo ~9% (6 commits)

Jul 2016

- Address a race condition when resizing the conntrack table by caching the bucket size when iterating over the hashtable (Liping Zhang).
- Revisit early_drop() path to perform lockless traversal on conntrack eviction under stress (Florian)
- Use rhashtable for the by-source NAT hashtable (Florian)
- Use binary search to validate jump offset in x_tables, this addresses the O(n!) validation (Florian)
- Remove zone extension as place it in the conntrack object (Florian)

Jul 2016 (2)

- Use generation mask in table, chain and set objects in nf_tables.
- Support for matching inverted set lookups (Arturo)

Nov 2016

- Add fib lookup expression for nf_tables (Florian Westphal).
- Introduce rt expression for nf_tables (Anders K. Pedersen).
- Add notrack support for nf_tables.
- nftables netdev family logging support.
- hash expression to allow arbitrary hashing of selector Concatenations (Laura García)
- Remove ip_conntrack sysctl backward compatibility code.
- Add quota expression for nf_tables.

Nov 2016 (2)

- Add number generator expression for nf_tables (Laura García)
- Remove the per-conntrack timer, use the workqueue approach (Florian)

Dec 2016

- Add support for stateful objects. http://marc.info/?l=netfilter-devel&m=148029128323837&w=2
- On-demand registration of nf_conntrack and defrag hooks per netns (Florian)
- Improve memory locality in x_tables: Allocate 4k chunks and then use these for x_tables counter allocation requests, this improves ruleset load time and also datapath ruleset evaluation (Florian)
- Add support to flush sets in nf_tables.
- Update layer 4 checksum if any of the pseudoheader fields is updated, for stateless NAT.

Dec 2016 (2)

- Built-in DCCP, SCTP and UDPlite conntrack and NAT support.

Feb 2017

- Revisit warning message when not applying default helper assignment (Jiri Kosina)
- Stash ctinfo 3-bit field into pointer to nf_conntrack object from sk_buff so we only access one single cacheline in the conntrack hotpath (Florian)
- Don't leak pointer to internal structures when exporting x_tables ruleset back to userspace (Willem DeBruijn)
- nf_log_all_netns sysctl (Michal Kubecek)

Feb 2017 (2)

- Merge UDPlite conntrack and NAT helpers into UDP (Florian)

May 2017

- Speed up netns by selective calls of synchronize_net() (Florian)
- Simplify ct extension infrastructure: No expensive runtime time calculation of extension area (Florian)
- pernet hook whenever possible (Florian)
- Allow to get rid of unassured flows under stress in conntrack fro DCCP, SCTP and TCP protocols (Florian)

May 2017 (2)

- Kill the fake untracked conntrack objects, use ctinfo instead (Florian)
- Ct event mask filtering (Florian)