

Netfilter kernel updates since last workshop
(only highlights not to get this too boring...;)
Pablo Neira Ayuso <pablo@netfilter.org>

Git stats since june 2015

- A bit of stats:
 - 276 commits
 - 49 unique contributors
 - Florian ~25% (84 commits)
 - Pablo ~20% (65 commits)

June 2015

- x_tables: avoid percpu ruleset duplication from Florian.
 - Florian Westphal: “On my test system (144 possible cpus, 400k dummy rules) this change saves close to 9 Gigabyte of RAM.”
- Warn if better conntrack protocol helper to use from Marcelo R. Leitner.
 - Helpers such as DCCP, SCTP and UDPLITE loadable via modprobe.
 - Break the Internet a bit for new protocols.
- Many ipset fixes from Sergey Popovich, a bit problematic since not willing to upstream things the way Jozsef needed.
- RCU for ipset from Jozsef.

June 2015 (2)

- Add secctx to nfnetlink_queue from Samsung.
- Add -m socket --transparent --restore-skmark from Qualcomm.
- From E.Bierdeman:
 - netns nf_tables fixes for nf_tables_core
 - nf_queue: Drop entries on nf_unregister_hook
- Reported by E. Leblond: keep processing nf_tables batch till end, then abort if missing modules, kills $O(n^m)$ where n is rules and m modules.

July 2015

- netfilter: Per network namespace netfilter hooks by Eric Bierderman.
- Static key for rare reentrancy case (tee) in x_tables from Florian: Avoid saves and restore of jumpstack.
- Fix netns dependencies with conntrack templates leading to CPU 100% spinning on exit path of conntrack ns.
- Support expectation in different ct zones, from Joe Stringer.

July 2016 (2)

- SCTP minimal multihoming support From M. Kubecek.
 - Currently secondary flows enter as NEW. This should be converted to RELATED.
 - Discussion on using expectations, without triggering large core rewrite for this specific case.
 - Can we do better?
- nf_tables optional rule counter to per-cpu.
- nft_limit: token-based limiting at nanosecond granularity.
- IPVS fixes from Julian.

August 2015

- nft_limit per byte support and burst parameter.
- Nfacct per ns support from Andreas Schultz.
- Add direction support for zones: NAT from isolated tenants with same private IP address scheme. From Daniel Borkmann.
 - Still missing work to avoid the existing adhoc change in the core to restore the mapping quickly in iptables
 - Proposed to avoid this by per-cpu conntrack template as scratchpad area to annotate mark to map it to zone.

August 2015 (2)

- VLAN header stripping for nft_payload from Florian Westphal.
- IPVS weighted overflow scheduling and fixes for sync daemon.
- Fix wrong endianness in nf_tables res_id field
 - Not many subsystem, so handle this value as alias.
- ICMP scheduling for IPVS from Alex (FB)
 - Using /proc interface to enable this...

September 2015

- Don't zap all loggers on module removal, from Florian Westphal.
- Another sysctl to ignore tunnel packets (and avoid lookups), from Alex FB.
- Large batch to pass net pointer to netfilter hooks, from Biederman.
- use y2038 safe timestamp from Arnd Bergmann.

October 2015

- Conntrack integration for nflog, from Ken-ichirou Matsuzawa.
- More pass net pointer to functions, from Biederman.
- Don't use ->prev pointer in hooks in nfqueue, from Florian.

November 2015

- Fix wrong permissions for iptables /proc entries to get users working, from Philip Whineray.
- Add nftables payload mangling support from Patrick.
- New extended tracing infrastructure for nftables, from Florian.

December 2015

- Add netns support to ctimeout.
- Add cgroup2 support to xt_cgroup from Tejun Heo.
- nftables skb->pkttype support, from Florian.
- Always include direction in nftables ct expression, also from Florian.
- More nftables netns: Missing release of objects at exit.
- nft_limit inversion support.

January 2016

- Packet duplication and forwarding for nftables netdev family.
- Don't break atomicity on errors in nflog and more interface fixes for nfqueue, from Kenichirou Matzuzawa.
- nft_byteorder 64 bits and ct counter match from Florian.
- Add NFTA_SET_USERDATA from Carlos Falgueras.
- Revisit lock all buckets in conntrack resize, from Sasha Levin
- Conntrack flush from workqueue in Ipv6, from Florian.

February 2016

- Remove nfnetlink mmap support from Florian.
- Add random support for nft_meta, also Florian.
- On-dmain hook registration in x_tables.

March 2016

- nft_masq port range support.
- IPVS fixes for SIP helper, from Marco Angaroni.
- Honor nfqueue fail-open flag netlink unicast fails.
- nftables bridge support for nfqueue, from Stephan Bryant.

April 2016

- Large batch from Florian to validate x_tables blob.
- One-Packet-Scheduler scheduler for IPVS, from Marco Angaroni.
- Allow adjacent intervals with dynamic updates in nft_rbtree.
- Disable automatic helper assignment, four years later. See E. Leblond secure usage of ct helpers.
- Connlabel support for nft, from Florian.

May 2016

- Rework conntrack netns support from Florian: One single table for all netns.
- Clash resolution on ct confirmation for connectionless protocols.
- More validation for missing netlink attributes, from Phil Turnbull.
- Bail out on duplicated ports in ct helpers via modprobe.

June 2016

- Several nftables from Liping Zhang:
 - Wrong genmask check in packet classifier.
 - Loop detection from set element.
 - Memory leak on error path.
 - Disable tracing via “meta trace set 0”.
- nf_tables generation mask to table, chain and sets.