

Netfilter's flowtable updates

Pablo Neira Ayuso
<pablo@netfilter.org>

Intro

- Available since Linux kernel 4.16
- Idea: “One single lookup from ingress to decide what to do with the packet”
- Documentation:
 - https://docs.kernel.org/networking/nf_flowtable.html

Updates

- New RFC round for xfrm bulking series in 2026 (17 Mar 2026) to
 - netfilter-devel@vger.kernel.org
 - netdev@vger.kernel.org
Subject: [PATCH net-next,RFC 0/8] netfilter: flowtable bulking
<https://lore.kernel.org/netfilter-devel/abk9ZQoc7GxfgODn@chamomile/>
- Performance evaluation:
 - <https://linux-ipsec.org/2025-linux-kernel-flowtable-bulk-forwarding-and-xfrm-pcpu-forwarding-testing-results.html>
- Feedback: GRO needs to be disabled, local traffic gets a penalty:
 - Try fraglist aggregation
- RFC series performance evaluation (Antony):
 - <https://linux-ipsec.org/2025-linux-kernel-flowtable-bulk-forwarding-and-xfrm-pcpu-forwarding-testing-results.html>

Updates (2)

- More preparation work: inline transmit path.
 - Push VLAN/PPPoE encap, if any + dev_queue_xmit()
 - Neigh lookup still happens
 - PPPoE GSO support:
 - Currently skb_gso_segment() from flowtable path.
 - Use recent PPPoE GSO support to from validate_xmit_skb():
 - Oneliner for netif_needs_gso() to use it: Works
 - SKB_GSO_PPPoE / NETIF_F_GSO_PPPOE to include/linux/netdev_features.h?
 - IPIP and IP6IP6 support (Lorenzo Bianconi)
 - Bugs with bridge device
- Flowtable bridge support: port to port fastpath
 - [nf-next,v5] netfilter: flowtable: initial bridge support

Roadmap?

- Flowtable xfrm only supports for TX path
 - Add support RX support path:
 - Extract xfrm RX code from RFC series and upstream it
- xfrm flowtable TX support broken with NAT?
 - Tuomo reported
- Support to recirculate packets back into GRO for local delivery?
 - Prioritize GSO forward aggregation, fallback local