netfilter logging infrastructure & nft

Pablo Neira Ayuso <pablo@netfilter.org>

Current state

```
# cat /proc/net/netfilter/nf_log
 0 NONE ()
 1 NONE ()
 2 NONE ()
 3 NONE ()
 4 NONE ()
 5 NONE ()
 6 NONE ()
 7 NONE ()
 8 NONE ()
 9 NONE ()
10 NONE ()
11 NONE ()
12 NONE ()
```

Current state

```
# modprobe nf_log_ipv4
# cat /proc/net/netfilter/nf_log
 0 NONE ()
 1 NONE ()
 2 nf_log_ipv4 (nf_log_ipv4)
 3 NONE ()
 4 NONE ()
 5 NONE ()
 6 NONE ()
 7 NONE ()
 8 NONE ()
 9 NONE ()
10 NONE ()
11 NONE ()
12 NONE ()
```

Current state

```
# modprobe nfnetlink_log
# cat /proc/net/netfilter/nf_log
 0 NONE (nfnetlink_log)
 1 NONE (nfnetlink_log)
 2 nf_log_ipv4 (nf_log_ipv4,nfnetlink_log)
 3 NONE (nfnetlink_log)
 4 NONE (nfnetlink_log)
 5 NONE (nfnetlink_log)
 6 NONE (nfnetlink_log)
 7 NONE (nfnetlink_log)
 8 NONE (nfnetlink_log)
 9 NONE (nfnetlink_log)
10 NONE (nfnetlink_log)
11 NONE (nfnetlink_log)
12 NONE (nfnetlink_log)
```

Current state

```
# echo nfnetlink_log > /proc/sys/net/netfilter/nf_log/2
# cat /proc/net/netfilter/nf_log
 0 NONE (nfnetlink_log)
 1 NONE (nfnetlink_log)
 2 nfnetlink_log (nf_log_ipv4,nfnetlink_log)
 3 NONE (nfnetlink_log)
 4 NONE (nfnetlink_log)
 5 NONE (nfnetlink_log)
 6 NONE (nfnetlink_log)
 7 NONE (nfnetlink_log)
 8 NONE (nfnetlink_log)
 9 NONE (nfnetlink_log)
10 NONE (nfnetlink_log)
11 NONE (nfnetlink_log)
12 NONE (nfnetlink_log)
```

Current state

- Need to be enabled via modprobe
- Control plane via /proc entry
- No real integration with tooling
- nfnetlink_log always uses group 0

Proposed solutions: Possibilities

- #1 Add standalone netlink interface
  - So iptables can use this.
- #2 Extend nft netlink API (transactions)
- Netlink format:
  - NFT_MSG_CONFIG, family via nfnlhdr->family.
    - NFTA_CFG_LOG
      - NFTA_CFG_LOG_NAME { “kernel”, “user” }
      - NFTA_CFG_LOG_DATA
        - NFTA_CFG_LOG_GROUP

Proposed solution

- nft syntax
  - From command line:

    nft default log ip kernel

    nft default log ip6 user group 0
  - From file:

    ```
    default log {
        ip kernel
        ip6 user group 0
    }
    ```
- Rule configuration stand over the default configuration

Other problems

- (“netfilter: allow logging from non-init namespaces”) from M. Kubecek:
  - /proc/sys/net/netfilter/nf_log_all_netns
- Syslog ns stuck back in 2013: https://lwn.net/Articles/527342/