

Netfilter userspace updates
since June 2016
NFWS 2017 – Faro, Portugal
Pablo Neira Ayuso
<pablo@netfilter.org>

iptables

- 1.6.1 release: 27 Jan 2017
- 121 commits
- 35 contributors
 - Liping: 28 commits
 - Pablo Bermudo: 17 commits
 - Phil Sutter: 7 commits
 - Pablo Neira: 6 commits
- 1.6.1 release

iptables

- Many updates for iptables-translate (Phil Sutter, Shivani, Gargi Sharma)
- Updates for the userspace iptables lock (Liping, Lorenzo Colitti)
 - --with-xt-lock-name
 - iptables-restore holds lock
 - Revisit to fix corner cases...
- ebpf pinned objects (Willem de Bruijn)
- On revision mismatch, do not call print/save (Willem)
- Fix build with musl (Baruch Siach)

nftables

- 317 commits
- 26 contributors
 - Pablo: 112 commits
 - Florian: 71 commits
 - Elise (Outreachy): 31 commits
 - Phil Sutter: 19 commits
 - Arturo: 18 commits
 - Liping: 10 commits
- One release:
 - 0.7 in Dec, 20th 2016

nftables

- xt compat support
 - ./configure --with-xtables
- Working incremental updates via nft -f
- meta random (Florian)
- Set non-byte header fields (Florian)
 - ip ecn set 1
- Checksum fixup with odd-sizes (Florian), ie. 16 bit aligned fetch.
 - Also, related to 'ip ecn set 1'

nftables (2)

- Quote user-defined strings when used from rule selectors.
- RFC2732 IPv6 address (brackets)
- POSIX.1-2008: Allowed characters in strings
- Parse meta priority support using tc classid
- Allow numeric conntrack labels (Florian)
- Create set command
 - `nft create set x y { type ipv4_addr\; }`
- Create element command
 - `nft create element x y { 1.1.1.1 }`

nftables (3)

- add quota statement
 - nft add rule filter input \
flow table http { \
ip saddr timeout 60s \
quota over 50 mbytes } drop
- add numgen expression
 - nft add rule nat prerouting \
dnat to numgen inc mod 2 map { \
0 : 192.168.10.100, \
1 : 192.168.20.200 }

nftables (4)

- src: add hash expression
 - nft add rule x y dnat to \
jhash ip saddr mod 2 seed 0xdeadbeef map { \
0 : 192.168.20.100, \
1 : 192.168.30.100 \
}
- Allow variable references in set elements definition
 - define s-ext-2-int = {
10.10.10.10 . 25, \
10.10.10.10 . 143 }

```
table inet forward {  
    set s-ext-2-int {  
        type ipv4_addr . inet_service  
        elements = $s-ext-2-int  
    }  
}
```

nftables (5)

- Many documentation updates (many)
- Fix endianness when printing/set II address (Florian)
 - ether daddr set 00:03:2d:2b:74:ec
- Variable to add/create/delete elements:
 - define whitelist_v4 = { 1.1.1.1 }
table inet filter {
 set whitelist_v4 { type ipv4_addr; }
}
add element inet filter whitelist_v4 \$whitelist_v4
- add offset attribute for numgen/hash expression
 - ct mark set numgen inc mod 2 offset 100

nftables (6)

- introduce routing expression
 - nft add rule filter postrouting \
flow table acct { \
rt nexthop timeout 600s counter }
- add fib expression (reverse path filtering)
 - nft add rule x prerouting fib saddr . iif oif eq 0 drop
- add notrack support
- add support to flush sets
 - nft flush set filter xyz

nftables (7)

- add used quota support
 - quota over 200 mbytes used 1143 kbytes drop
- add/create/delete/list/reset stateful objects
 - nft add counter filter badguy1
nft add map filter badguys { \
 type ipv4_addr : counter \; }
nft add rule filter input counter name \
 ip saddr map @badguys
nft add element filter badguys { \
 192.168.2.3 : "badguy1" }

nftables (8)

- sort set elements by default (Elise)
- add average bytes per packet counter support
 - nft add rule x y ct avgpkt \> 100
- Allow to list ruleset without stateful information
 - nft -s list ruleset
- Fix listing of icmp type in inet family (Arturo)
- TCP option matching (Manuel Messner)
- Store byteorder for set keys/data
 - nft add rule x y meta mark set meta cpu map { 0 : 1, 1 : 2 }
- Support of symmetric hash (Laura)

nftables (9)

- Boolean datatype (Phil)
 - fib daddr oif exists
- ct zone set support (Florian)
- ct helper support (Florian)

```
table ip filter {  
    ct helper ftp-standard {  
        type "ftp" protocol tcp  
    }  
    chain y {  
        tcp dport 21 ct helper set "ftp-standard"  
    }  
}
```

nftables (10)

- Missing ICMPv6 types (Phil)
- Flush flow tables (Elise)
 - nft flush flow table filter ft-https
- Ct event mask (Florian)
- No wraparounds when printing sets and maps (Arturo)
- Reject ambiguous set (Arturo)
 - tcp dport set {1, 2, 3, 4, 5}
- Include directory and glob support (Ismo Puustinen)

libnftnl

- 109 commits
- 16 authors
 - Pablo: 48 commits
 - Carlos Falgueras: 19 commits
 - Phil Sutter: 11 commits
 - Florian: 7 commits
- One release:
 - Dec 19th, 2016

conntrack-tools

- 37 commits
- 7 authors
 - Arturo: 21 commits
 - Kevin Cernekee: 7 commits
 - Pablo: 6 commits
- Releases:
 - 1.4.4, Aug 22th, 2016

conntrack-tools

- IPv6 NAT for conntrack command line (Neil Wilson)
- UPnP eventing userspace conntrack helper (Kevin Cernekee)
- Daemon runs with rt scheduler (Arturo)
- Deprecate lots of options for conntrackd (Arturo)
- Request resync at startup clause (Arturo)
- Better logging in conntrackd (Arturo)
- mdns conntrack helper (Kevin)