

netdev 0x17

Jakub Kicinski

netdev process update,
development stats

❖ Process

❖ Work

❖ Development statistics

Process

Process

- 7 weeks of development + 2 weeks of merge window
 - net-next closed for 2 weeks
 - no plans to keep net-next open
- min 24 hour re-posting delay
- 24-72 hour review cycle
- official shutdowns for holiday periods
- vendor patch review rotation
- development stats published for each cycle
- “netdev FAQ” <https://www.kernel.org/doc/html/next/process/maintainer-netdev.html>

Challenges & barriers to participation

- High patch volume
 - lore + lei
- Lack of documentation
 - affects users users, developers, newcomers
- Lack of testing / verification
 - lower quality
 - divergent implementations
 - duplication of testing effort
 - harder for newcomers to gain confidence

Work areas

Work I've been doing

- Netlink YAML / YNL
- Page pool changes
- Queue API...
- PSP...

Configuration APIs

- Avoid implicit assumptions
 - drivers, protocols, etc make different choices for gray areas
 - makes integration / portability harder
- Avoid custom APIs (where possible)
 - stats
- Focus on real usability gaps
 - default queue counts
 - IRQ pinning
 - capability reporting
- Programming language accessibility
 - Netlink YAML - hide the netlink, keep the features
 - reasonable adoption kernel side, 12 specs upstream
 - user space side taking off more slowly, C/Python in the tree, hacky C++ at Meta

Netlink YAML / Configuration APIs

- Netlink Handbook
[Netlink Handbook — The Linux Kernel documentation](#)
- LPC talk
[Machine readable description for netlink protocols](#)

- kernel policy / op generation
- various quality of life improvements
- user space binding generation

Netlink YAML

- adoption - relative newcomers scared of netlink
- dev efficiency - “wasted” development time
- documentation / self-documenting
- user space integration

Netlink YAML - adoption

| | | |
|--------------|------------------|-------------------|
| 38.66% | Donald | Hunter |
| 22.40% | Jiri | Pirko |
| 15.82% | Stanislav | Fomichev |
| 9.61% | Jakub | Kicinski |
| 5.72% | Vadim | Fedorenko |
| 4.69% | Davide | Caratti |
| 1.53% | Chuck | Lever |
| 1.07% | Lorenzo | Bianconi |
| 0.36% | Arkadiusz | Kubalewski |
| 0.07% | Shay | Agroskin |
| 0.06% | Maciej | Fijalkowski |
| 0.01% | Tariq | Toukan |

Netlink YAML

- Dev efficiency
 - Python cli seems to be the new go-to
 - ...to a fault.
- Documentation
 - Human-friendly rendering work-in-progress
- User space integration
 - C code gen in the kernel
 - Hacky Python support (cli.py)
 - Hacky C++ code gen work-in-progress

Page pool API

- Get structured stats
- Ties in with the memory provider and queue APIs

```
$ ./page-pool
eth0[2] page pools: 10 (zombies: 0)
        refs: 41984 bytes: 171966464 (refs: 0 bytes: 0)
        recycling: 90.3% (alloc: 656:397681 recycle: 89652:270201)
```

Queue API

- Extract queue management and configuration from the drivers
- Allow more flexible and dynamic configuration
- The GET API (already posted) is very useful for configuring IRQ pinning!

Development stats

6.7 stats - top reviewers

Top reviewers (thr):

```
1 (   ) [24] Simon Horman
2 (   ) [23] Jakub Kicinski
3 (   ) [12] Andrew Lunn
4 ( +1) [10] Paolo Abeni
5 ( +1) [ 8] Eric Dumazet
6 ( +1) [ 7] David Ahern
7 (+21) [ 7] Kees Cook
8 (+12) [ 6] Jiri Pirko
9 (+14) [ 6] Florian Fainelli
10 (+19) [ 6] Jacob Keller
11 ( -2) [ 4] Russell King
12 ( -2) [ 4] Willem de Bruijn
13 ( +2) [ 4] Vladimir Oltean
14 ( -1) [ 4] Rob Herring
15 (**) [ 3] Wojciech Drewek
```

Top reviewers (msg):

```
1 ( +1) [43] Jakub Kicinski
2 ( -1) [33] Simon Horman
3 (   ) [33] Andrew Lunn
4 ( +5) [15] Eric Dumazet
5 ( +1) [14] David Ahern
6 ( +2) [13] Paolo Abeni
7 (+21) [13] Florian Fainelli
8 ( +5) [12] Jiri Pirko
9 (+23) [12] Jacob Keller
10 ( +8) [10] Vladimir Oltean
11 (+24) [ 9] Kees Cook
12 ( +5) [ 8] Rob Herring
13 ( -8) [ 8] Russell King
14 ( +1) [ 8] Jason Wang
15 ( -4) [ 7] Willem de Bruijn
```


6.7 stats - company scores

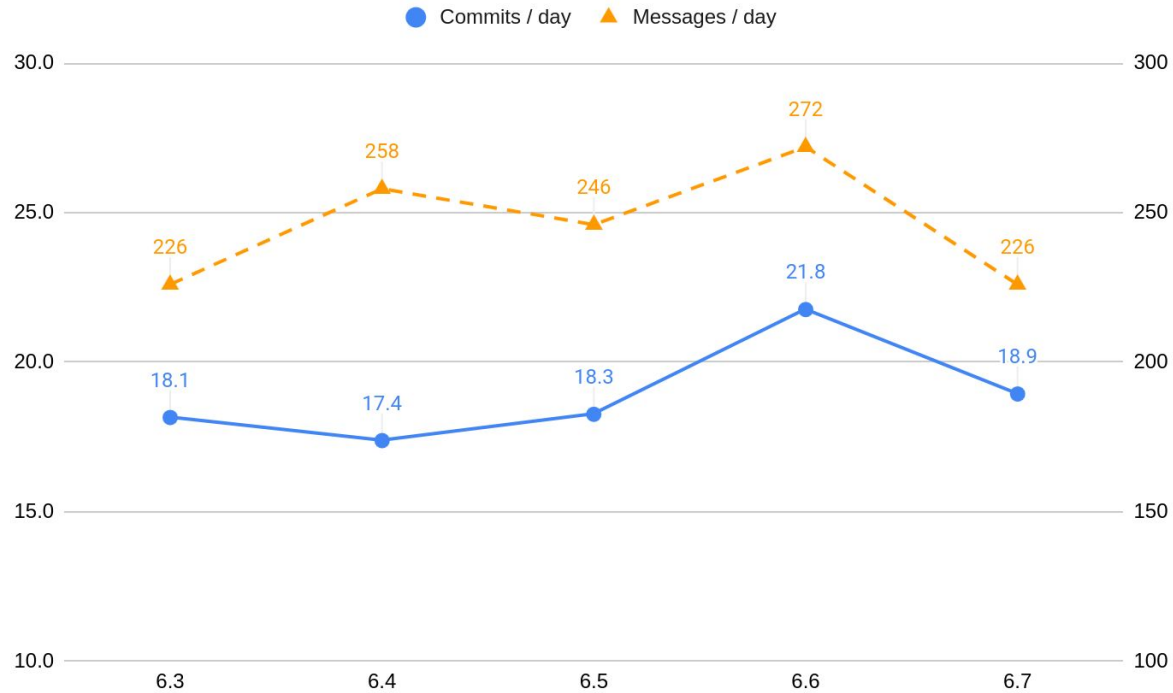
Top scores (positive):

| | | | |
|----|-------|-------|------------------|
| 1 | (+3) | [341] | RedHat |
| 2 | () | [219] | Meta |
| 3 | () | [183] | Andrew Lunn |
| 4 | (+2) | [95] | Enfabrica |
| 5 | (+4) | [59] | Broadcom |
| 6 | (+17) | [52] | Isovalent |
| 7 | (+3) | [47] | ARM |
| 8 | (+5) | [46] | Oracle |
| 9 | (-1) | [44] | Linux Foundation |
| 10 | (-5) | [32] | Linaro |

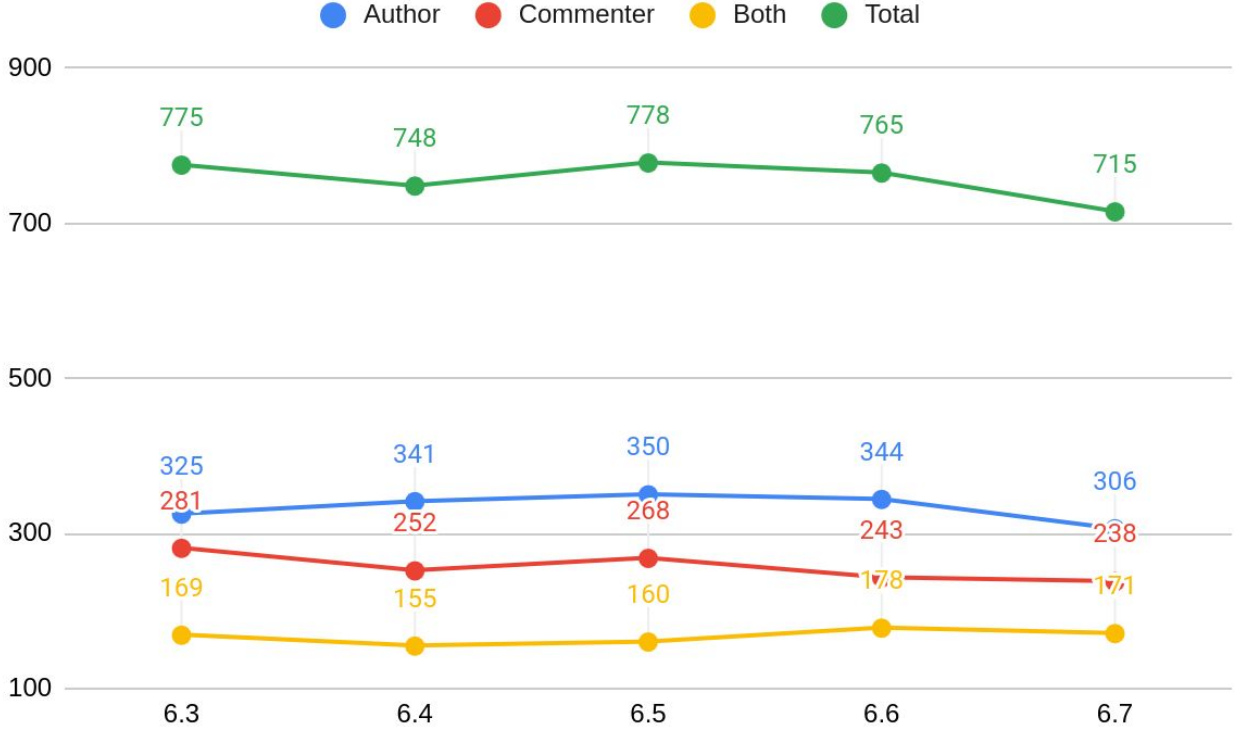
Worst scores (negative):

| | | | |
|---|-------|------|-------------|
| 1 | (+13) | [97] | Bootlin |
| 2 | (***) | [72] | nVidia |
| 3 | (-2) | [66] | Huawei |
| 4 | (+4) | [59] | Arista |
| 5 | (+10) | [48] | Alibaba |
| 6 | (***) | [41] | Pengutronix |

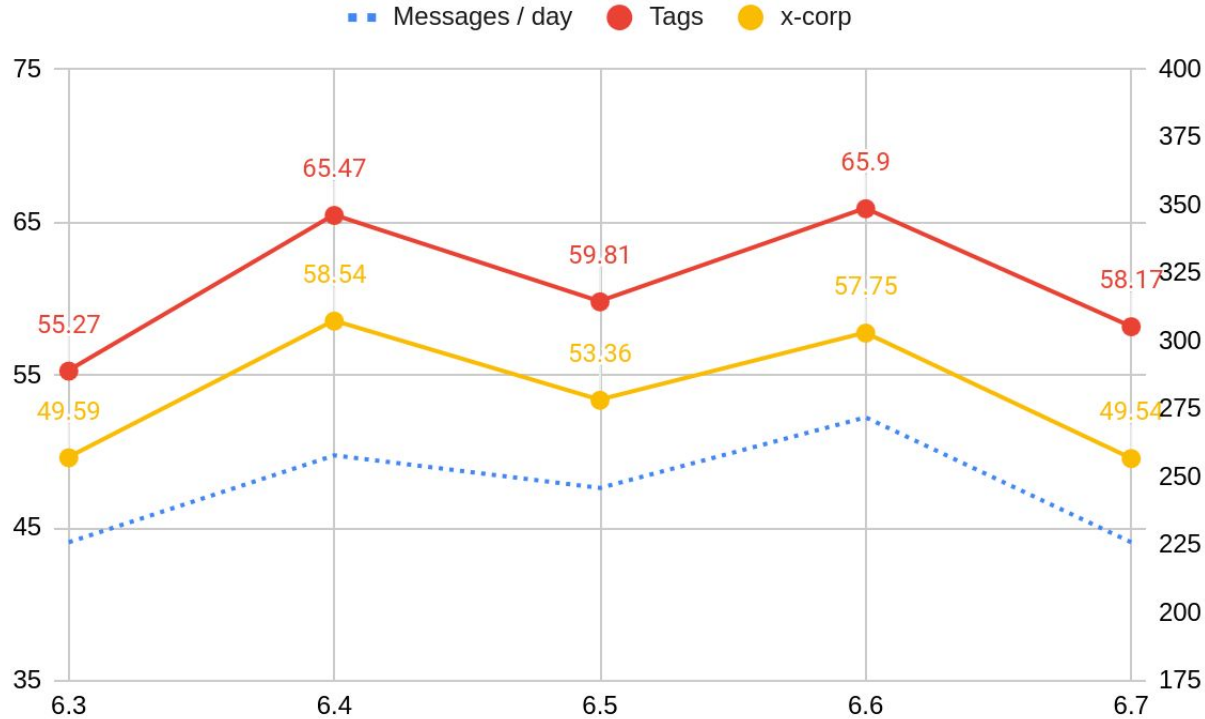
Commits & emails / day



Number of participants



Review tags



Fin

- review stuff!
- side projects [are available](#)
- feedback is welcome