

# The FreeBSD lock order reversals (LOR) list

## The Past, the Present and the Future

Bjoern A. Zeeb  
bz@FreeBSD.org

The FreeBSD Project

BSDCAN - The Technical BSD Conference, 2007

# Outline

- 1 History
  - Why do we have a LOR list?
- 2 A 'LOR' - what is that?
  - witness(4) and LOR
- 3 From Past to Present
  - Facts and Figures
  - Formats and Tools
  - State of affairs
- 4 The Future
  - Ideas and Discussion

## In the beginning . . .

- There was the ‘Giant’-lock
- Then was 5-CURRENT and WITNESS
- Now there were mailing list postings of LORs
- There were more mailing list postings of the same LORs
- And replies: “This has been reported about  $n$  times before.”
- A FAQ entry
- An idea

# In the beginning . . .

. . . about six months later . . .

- In April 2004 - an initial LOR list  
<http://sources.zabbadoz.net/freebsd/lor.html>

# What you see

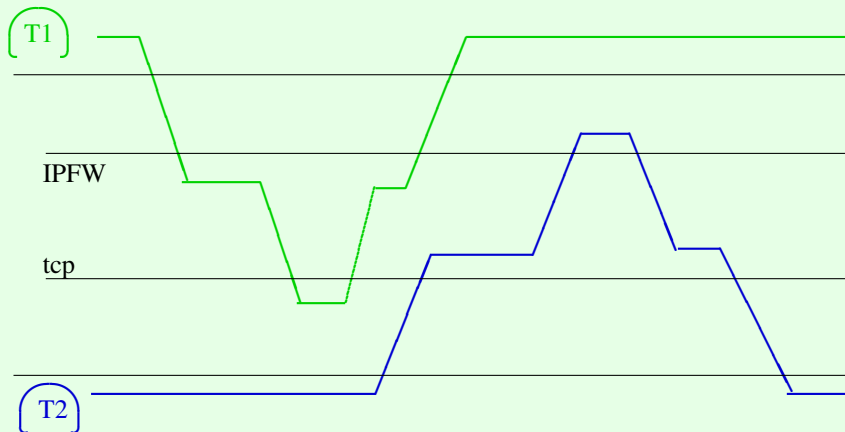
## show witness (excerpt):

```
..
1  proctree -- last acquired @ sys/kern/kern_exit.c:683
2  allproc -- last acquired @ sys/kern/sched_4bsd.c:378
11 sleep mtxpool -- last acquired @ sys/kern/kern_resource.c:1219
12  kqueue -- last acquired @ sys/kern/kern_event.c:1443
13    struct mount mtx -- (already displayed)
8   process lock -- last acquired @ sys/kern/kern_kthread.c:194
9   session -- last acquired @ sys/kern/kern_proc.c:573
10  uidinfo hash -- last acquired @ sys/kern/kern_resource.c:1181
11  uidinfo struct -- last acquired @ order list:0
12  allprison -- last acquired @ sys/kern/kern_jail.c:770
11  sleep mtxpool -- (already displayed)
10  tty -- last acquired @ sys/kern/tty.c:2793
14  vnode interlock -- (already displayed)
9   sigacts -- last acquired @ sys/kern/subr_sleepqueue.c:389
..
```

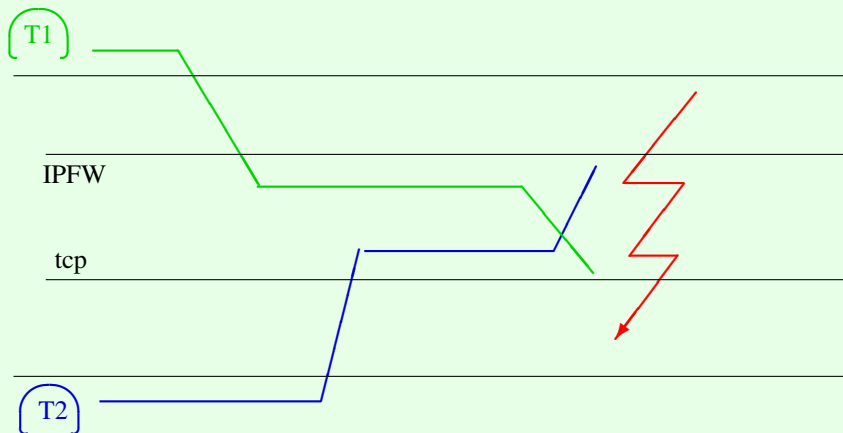
## LOR:

```
lock order reversal
 1st 0xc07287c8 IPFW static rules (IPFW static rules) @ sys/netinet/ip_fw2.c:1828
 2nd 0xc065cfcc tcp (tcp) @ sys/netinet/ip_fw2.c:1574
...
```

# The lucky one



# Lock lock deadlock



# 200 and still counting

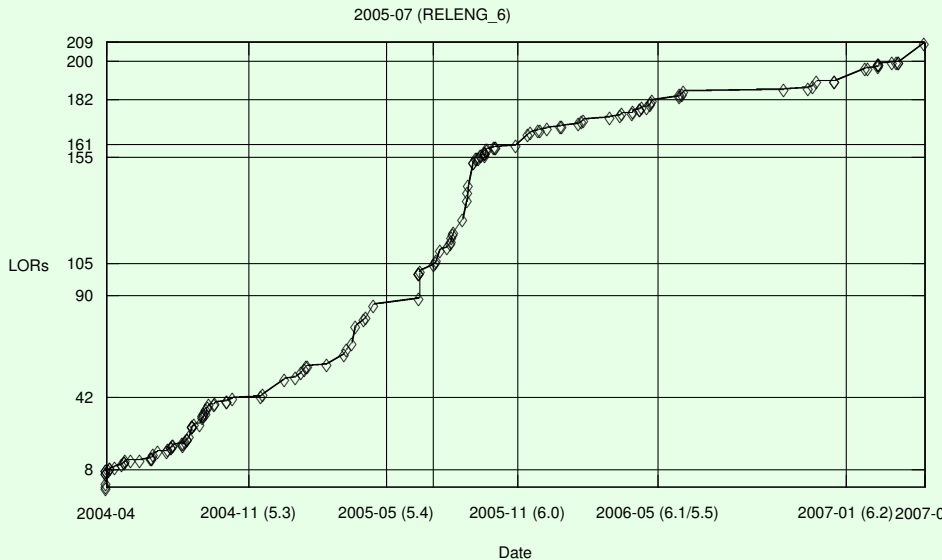
1st May 2007

## 209 LORs in total

- 132 unknown
- 63 patched
- 6 result of panic
- 2 might lead to deadlock
- 1 result of trap
- 1 fixed
- 1 result of LOR #195 (not a problem)
- 1 leads to deadlock
- 1 potentially harmful
- 1 programming error



# Updates to the LOR page



# The List version 1, 2, 3, ...

- HTML
- Single loooong page
  - LOR ID
  - traceback
  - comments, links to PRs, mailing lists and commits
- Split into overview list and details (Sept. 2005, ~160 LORs)
- Overview list and details and more information (Nov. 2005)
  - Reported first/last
  - Responsible
- Tools: cp, vi(m), rcs

# This is now

- Still HTML
- SPOF: maintainer
- ‘Batch processing’
- Still duplicates on mailing lists
- Feedback from committers and community
  - Reports of new LORs
  - Feedback on committed patches
  - General Feedback

# Ideas

- How to get better feedback?
  - Regular mails (weekly/monthly) like done for PRs?
  - More prodding needed?
  - Do we want PRs for LORs?
  - States and descriptions?
  - Kernel printf
  - ...?
- Improve handling?
  - XMLsify?
  - So it can be moved to the doc tree?
  - Have a webform for checking and submitting?
  - ...?
- Other locking 'problems' reported?
- Do we still need it?

# EOF

More feedback to:

- [bzeeb+freebsd+lor@zabbadoz.net](mailto:bzeeb+freebsd+lor@zabbadoz.net)
- [bz@FreeBSD.org](mailto:bz@FreeBSD.org)