

# Linux Crash Dump Analysis

Crash Dump Analysis 2014/2015




CHARLES UNIVERSITY IN PRAGUE

faculty of mathematics and physics

Department of  
Distributed and  
Dependable  
Systems



ORACLE®



# Agenda

- Userspace debugging
- Understanding kernel oops/panic output
- Creating kernel crash dumps
- Tools for crash dump analysis
- Debugging live systems

# User space debugging

- Crashing application
  - Through an (unhandled) signal
    - SIGABRT, SIGFPE, SIGSEGV, SIGBUS...
  - Typically produces a line in kernel log (dmesg)
    - `modprobe[833]: segfault at 7fff76200038 ip 00007f0de8422fc2 sp 00007fff761b6cb0 error 4 in ld-2.11.1.so[7f0de8420000+20000]`
  - Will produce a `core(5)` file if
    - Limits allow it (`ulimit -c; /etc/security/limits.conf`)
    - Binary is readable and not `suid`
    - ...
- Application stuck in syscall
  - `cat /proc/$PID/stack`

# User space debugging

- Executing a program under `gdb`
  - Relies on the `ptrace(2)` syscall
  - `gdb /path/to/binary`
  - `(gdb) run $param1 $param2 ...`
- Attaching to a running program
  - `gdb -p $PID`
  - We can also create core file (without crashing)
    - `(gdb) generate-core-file`
- Inspecting the core file
  - `gdb /path/to/binary /path/to/core`

# User space debugging

- **strace** – tool for tracing system calls and signals
  - Prints system call parameters and return values with symbolic translation
    - `open("/foo/bar", O_RDONLY) = -1 ENOENT (No such file or directory)`
    - Tries to keep ordering of enter/return between threads
    - Dereferences structure members
    - Can attach to a PID
  - “In some cases, strace output has proven to be more readable than the source.”
- **valgrind** – for finding memory access bugs

# Kernel oops/panic output

- Printed in console typically on fatal CPU exceptions
  - Lots of architecture-specific information
  - May contain enough information to figure out the problem without a full crash dump
- Oops leaves the system running
  - Kills just the current process (including kernel threads!)
  - System can still be left inconsistent (locks remain locked ... )
- Panic kills the system completely
  - Oops in interrupt, with `panic_on_oops` enabled, manual `panic()` calls
    - HW failure, critical memory allocation fail, init/idle task killed, int. handler killed
  - May trigger crash dump if configured, or reboot after delay

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amdkgd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLAGS: 00010286
[ 266.492285] RAX: ffff88008f6b3ba0 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492331] RDX: ffff8800914b3c98 RSI: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492376] RBP: ffff8800916fbb68 R08: 0000000000000002 R09: 0000000000000000
[ 266.492421] R10: 0000000000000008 R11: 0000000000000001 R12: ffff88008f686068
[ 266.492465] R13: ffff8800914b3c98 R14: ffff88008f6b3b90 R15: ffff88008f686000
[ 266.492513] FS: 00007fb8966f6700(0000) GS:ffff88011ed80000(0000)
knlGS:0000000000000000
[ 266.492566] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[ 266.492601] CR2: 00007f50fa190770 CR3: 0000000001b31000 CR4: 00000000000407e0
[ 266.492652] Stack:
[ 266.492665] 0000000000000000 ffff88008f686078 ffff8800916fbb68 ffff88008f686000
[ 266.492714] ffff8800916fbc08 0000000000000000 0000000000000000 ffff88008f686000
[ 266.492764] ffff8800916fbbf8 ffffffff8111ba5d 00007fb885918000 ffff88008edf3000
...
```

# Example kernel oops

```
..
[ 266.492815] Call Trace:
[ 266.492834] [] free_pgtables+0x8e/0xcc
[ 266.492873] [] exit_mmap+0x84/0x116
[ 266.492907] [] mmput+0x52/0xe9
[ 266.492940] [] do_exit+0x3cd/0x9c9
[ 266.492975] [] ? _raw_spin_unlock_irq+0x2d/0x32
[ 266.493016] [] do_group_exit+0x4c/0xc9
[ 266.493051] [] get_signal+0x58f/0x5bc
[ 266.493090] [] do_signal+0x28/0x5b1
[ 266.493123] [] ? sysret_signal+0x5/0x43
[ 266.493162] [] do_notify_resume+0x35/0x68
[ 266.493200] [] int_signal+0x12/0x17
[ 266.493235] Code: e8 03 b7 f4 ff 49 8b 47 78 4c 8b 20 48 8d 58 f0 49 83
ec 10 48 8d 43 10 48 39 45 c8 74 55 48 8b 7b 08 83 bf 8c 00 00 00 00 74 02
<0f> 0b e8 a4 fd ff ff 48 8b 43 18 48 8b 53 10 48 89 df 48 89 42
[ 266.493404] RIP [] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

(source: <https://lkml.org/lkml/2015/1/11/14>)



# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLBLADE Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 t ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0010:ffff8800916f8000 RAX: 00010286
[ 266.492285] RAX: ffff8800a3b3cf30 RCX: ffff8800a3b3cf30
[ 266.492331] RDI: ffff8800914b3c98 RDI: ffff8800914b3c98
[ 266.492376] R09: 0000000000000000 R09: 0000000000000000
[ 266.492421] R12: ffff88008f686068 R12: ffff88008f686068
[ 266.492465] R15: ffff88008f686000 R15: ffff88008f686000
[ 266.492513] FS: d80000(0000)
knlGS:0000000000000000
[ 266.492566] CS: 80050033
[ 266.492601] CR4: 00000000000407e0
[ 266.492652] Stack: 000916fba8 ffff88008f686000
[ 266.492665] 0000000000 ffff88008f686000
[ 266.492714] b885918000 ffff88008edf3000
[ 266.492764]
...
```

File + line translation enabled by  
CONFIG\_DEBUG\_BUGVERBOSE  
(implemented by \_\_bug\_table  
section on x86 - ~70-100kB)

The line in question contains:  
BUG\_ON(anon\_vma->degree);

This is essentially a hard assertion:  
if (<condition>) BUG()

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd64_edac_module amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 1155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800916f8000 b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:ffff81126630>] [<ffffffffff81126630>]
unlink_anon_vmas+0x102/0x1000
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLAGS: 00010286
[ 266.492285] RAX: ffff8800916fbb68 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
RDX: 0000000000000001 RDI: ffff8800914b3c98
RSI: 0000000000000002 R09: 0000000000000000
R10: 0000000000000001 R12: ffff88008f686068
R13: ffff88008f6b3b90 R15: ffff88008f686000
GS: ffff88011ed80000(0000)
CR0: 0000000080050033
CR2: 0000000001b31000 CR4: 00000000000407e0
RIP: 0010:ffff88008f686078 ffff8800916fbb68 ffff88008f686000
[ 266.492714] ffff8800916fbb68 0000000000000000 0000000000000000 ffff88008f686000
[ 266.492764] ffff8800916fbbf8 ffffffff8111ba5d 00007fb885918000 ffff88008edf3000
...
```

On x86, BUG() emits a standardized invalid opcode UD2 (0F 0B) triggering an exception. The exception handler checks for UD2 opcode and searches \_\_bug\_table for details.

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fb
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: ASUS BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800916f8000 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffff880091126630>] [<ffffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLAGS: 00010286
[ 266.492285] RAX: ffff8800916fbb68 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492321] RDI: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492357] RSI: 0000000000000002 R09: 0000000000000000
[ 266.492393] R10: 0000000000000001 R12: ffff88008f686068
[ 266.492429] R13: ffff88008f6b3b90 R15: ffff88008f686000
[ 266.492465] GS: ffff88011ed80000(0000)
[ 266.492501] CR0: 0000000080050033
[ 266.492537] CR2: 0000000001b31000 CR4: 00000000000407e0
[ 266.492573] CR8: 0000000000000000
[ 266.492609] f686078 ffff8800916fbb68 ffff88008f686000
[ 266.492645] 00000000 0000000000000000 ffff88008f686000
[ 266.492681] 111ba5d 00007fb885918000 ffff88008edf3000
[ 266.492717] ...
```

x86- and exception-specific error code (32-bit hex number). Typically useful for page fault exceptions where it's a mask:

Bit 0 – Present

Bit 1 – Write

Bit 2 – User

Bit 3 – Reserved write

Bit 4 – Instruction fetch



# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD Family 16h Processor/INA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3cf30 task.ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffff8800916f8000>] [<ffffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916f8000 EFLAGS: 00010286
[ 266.492285] RAX: ffff8800916f8000 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492321] R08: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492357] R09: 0000000000000000 R10: 0000000000000001
[ 266.492393] R11: ffff88008f686068 R12: ffff88008f686068
[ 266.492429] R13: ffff88008f6b3b90 R14: ffff88008f686000
[ 266.492465] R15: ffff88008f686000 CR2: ffff88011ed80000(0000)
[ 266.492501] CR0: 0000000080050033
[ 266.492537] CR2: 000000001b310000 CR4: 00000000000407e0
[ 266.492573] DR0: 0000000000000000 DR1: 0000000000000000
[ 266.492609] DR2: 0000000000000000 DR3: 0000000000000000
[ 266.492645] DR6: 0000000000000000 DR7: 0000000000000000
[ 266.492681] Call Trace:
[ 266.492717]  [ 266.492753]  [ 266.492789]  [ 266.492825]  [ 266.492861]  [ 266.492897]
[ 266.492933]  [ 266.492969]  [ 266.493005]  [ 266.493041]  [ 266.493077]
[ 266.493113]  [ 266.493149]  [ 266.493185]  [ 266.493221]  [ 266.493257]
[ 266.493293]  [ 266.493329]  [ 266.493365]  [ 266.493401]  [ 266.493437]
[ 266.493473]  [ 266.493509]  [ 266.493545]  [ 266.493581]  [ 266.493617]
[ 266.493653]  [ 266.493689]  [ 266.493725]  [ 266.493761]  [ 266.493797]
[ 266.493833]  [ 266.493869]  [ 266.493905]  [ 266.493941]  [ 266.493977]
[ 266.494013]  [ 266.494049]  [ 266.494085]  [ 266.494121]  [ 266.494157]
[ 266.494193]  [ 266.494229]  [ 266.494265]  [ 266.494301]  [ 266.494337]
[ 266.494373]  [ 266.494409]  [ 266.494445]  [ 266.494481]  [ 266.494517]
[ 266.494553]  [ 266.494589]  [ 266.494625]  [ 266.494661]  [ 266.494697]
[ 266.494733]  [ 266.494769]  [ 266.494805]  [ 266.494841]  [ 266.494877]
[ 266.494913]  [ 266.494949]  [ 266.494985]  [ 266.495021]  [ 266.495057]
[ 266.495093]  [ 266.495129]  [ 266.495165]  [ 266.495201]  [ 266.495237]
[ 266.495273]  [ 266.495309]  [ 266.495345]  [ 266.495381]  [ 266.495417]
[ 266.495453]  [ 266.495489]  [ 266.495525]  [ 266.495561]  [ 266.495597]
[ 266.495633]  [ 266.495669]  [ 266.495705]  [ 266.495741]  [ 266.495777]
[ 266.495813]  [ 266.495849]  [ 266.495885]  [ 266.495921]  [ 266.495957]
[ 266.495993]  [ 266.496029]  [ 266.496065]  [ 266.496101]  [ 266.496137]
[ 266.496173]  [ 266.496209]  [ 266.496245]  [ 266.496281]  [ 266.496317]
[ 266.496353]  [ 266.496389]  [ 266.496425]  [ 266.496461]  [ 266.496497]
[ 266.496533]  [ 266.496569]  [ 266.496605]  [ 266.496641]  [ 266.496677]
[ 266.496713]  [ 266.496749]  [ 266.496785]  [ 266.496821]  [ 266.496857]
[ 266.496893]  [ 266.496929]  [ 266.496965]  [ 266.497001]  [ 266.497037]
[ 266.497073]  [ 266.497109]  [ 266.497145]  [ 266.497181]  [ 266.497217]
[ 266.497253]  [ 266.497289]  [ 266.497325]  [ 266.497361]  [ 266.497397]
[ 266.497433]  [ 266.497469]  [ 266.497505]  [ 266.497541]  [ 266.497577]
[ 266.497613]  [ 266.497649]  [ 266.497685]  [ 266.497721]  [ 266.497757]
[ 266.497793]  [ 266.497829]  [ 266.497865]  [ 266.497901]  [ 266.497937]
[ 266.497973]  [ 266.498009]  [ 266.498045]  [ 266.498081]  [ 266.498117]
[ 266.498153]  [ 266.498189]  [ 266.498225]  [ 266.498261]  [ 266.498297]
[ 266.498333]  [ 266.498369]  [ 266.498405]  [ 266.498441]  [ 266.498477]
[ 266.498513]  [ 266.498549]  [ 266.498585]  [ 266.498621]  [ 266.498657]
[ 266.498693]  [ 266.498729]  [ 266.498765]  [ 266.498801]  [ 266.498837]
[ 266.498873]  [ 266.498909]  [ 266.498945]  [ 266.498981]  [ 266.499017]
[ 266.499053]  [ 266.499089]  [ 266.499125]  [ 266.499161]  [ 266.499197]
[ 266.499233]  [ 266.499269]  [ 266.499305]  [ 266.499341]  [ 266.499377]
[ 266.499413]  [ 266.499449]  [ 266.499485]  [ 266.499521]  [ 266.499557]
[ 266.499593]  [ 266.499629]  [ 266.499665]  [ 266.499701]  [ 266.499737]
[ 266.499773]  [ 266.499809]  [ 266.499845]  [ 266.499881]  [ 266.499917]
[ 266.499953]  [ 266.499989]  [ 266.500025]  [ 266.500061]  [ 266.500097]
[ 266.500133]  [ 266.500169]  [ 266.500205]  [ 266.500241]  [ 266.500277]
[ 266.500313]  [ 266.500349]  [ 266.500385]  [ 266.500421]  [ 266.500457]
[ 266.500493]  [ 266.500529]  [ 266.500565]  [ 266.500601]  [ 266.500637]
[ 266.500673]  [ 266.500709]  [ 266.500745]  [ 266.500781]  [ 266.500817]
[ 266.500853]  [ 266.500889]  [ 266.500925]  [ 266.500961]  [ 266.500997]
[ 266.501033]  [ 266.501069]  [ 266.501105]  [ 266.501141]  [ 266.501177]
[ 266.501213]  [ 266.501249]  [ 266.501285]  [ 266.501321]  [ 266.501357]
[ 266.501393]  [ 266.501429]  [ 266.501465]  [ 266.501501]  [ 266.501537]
[ 266.501573]  [ 266.501609]  [ 266.501645]  [ 266.501681]  [ 266.501717]
[ 266.501753]  [ 266.501789]  [ 266.501825]  [ 266.501861]  [ 266.501897]
[ 266.501933]  [ 266.501969]  [ 266.502005]  [ 266.502041]  [ 266.502077]
[ 266.502113]  [ 266.502149]  [ 266.502185]  [ 266.502221]  [ 266.502257]
[ 266.502293]  [ 266.502329]  [ 266.502365]  [ 266.502401]  [ 266.502437]
[ 266.502473]  [ 266.502509]  [ 266.502545]  [ 266.502581]  [ 266.502617]
[ 266.502653]  [ 266.502689]  [ 266.502725]  [ 266.502761]  [ 266.502797]
[ 266.502833]  [ 266.502869]  [ 266.502905]  [ 266.502941]  [ 266.502977]
[ 266.503013]  [ 266.503049]  [ 266.503085]  [ 266.503121]  [ 266.503157]
[ 266.503193]  [ 266.503229]  [ 266.503265]  [ 266.503301]  [ 266.503337]
[ 266.503373]  [ 266.503409]  [ 266.503445]  [ 266.503481]  [ 266.503517]
[ 266.503553]  [ 266.503589]  [ 266.503625]  [ 266.503661]  [ 266.503697]
[ 266.503733]  [ 266.503769]  [ 266.503805]  [ 266.503841]  [ 266.503877]
[ 266.503913]  [ 266.503949]  [ 266.503985]  [ 266.504021]  [ 266.504057]
[ 266.504093]  [ 266.504129]  [ 266.504165]  [ 266.504201]  [ 266.504237]
[ 266.504273]  [ 266.504309]  [ 266.504345]  [ 266.504381]  [ 266.504417]
[ 266.504453]  [ 266.504489]  [ 266.504525]  [ 266.504561]  [ 266.504597]
[ 266.504633]  [ 266.504669]  [ 266.504705]  [ 266.504741]  [ 266.504777]
[ 266.504813]  [ 266.504849]  [ 266.504885]  [ 266.504921]  [ 266.504957]
[ 266.504993]  [ 266.505029]  [ 266.505065]  [ 266.505101]  [ 266.505137]
[ 266.505173]  [ 266.505209]  [ 266.505245]  [ 266.505281]  [ 266.505317]
[ 266.505353]  [ 266.505389]  [ 266.505425]  [ 266.505461]  [ 266.505497]
[ 266.505533]  [ 266.505569]  [ 266.505605]  [ 266.505641]  [ 266.505677]
[ 266.505713]  [ 266.505749]  [ 266.505785]  [ 266.505821]  [ 266.505857]
[ 266.505893]  [ 266.505929]  [ 266.505965]  [ 266.506001]  [ 266.506037]
[ 266.506073]  [ 266.506109]  [ 266.506145]  [ 266.506181]  [ 266.506217]
[ 266.506253]  [ 266.506289]  [ 266.506325]  [ 266.506361]  [ 266.506397]
[ 266.506433]  [ 266.506469]  [ 266.506505]  [ 266.506541]  [ 266.506577]
[ 266.506613]  [ 266.506649]  [ 266.506685]  [ 266.506721]  [ 266.506757]
[ 266.506793]  [ 266.506829]  [ 266.506865]  [ 266.506901]  [ 266.506937]
[ 266.506973]  [ 266.507009]  [ 266.507045]  [ 266.507081]  [ 266.507117]
[ 266.507153]  [ 266.507189]  [ 266.507225]  [ 266.507261]  [ 266.507297]
[ 266.507333]  [ 266.507369]  [ 266.507405]  [ 266.507441]  [ 266.507477]
[ 266.507513]  [ 266.507549]  [ 266.507585]  [ 266.507621]  [ 266.507657]
[ 266.507693]  [ 266.507729]  [ 266.507765]  [ 266.507801]  [ 266.507837]
[ 266.507873]  [ 266.507909]  [ 266.507945]  [ 266.507981]  [ 266.508017]
[ 266.508053]  [ 266.508089]  [ 266.508125]  [ 266.508161]  [ 266.508197]
[ 266.508233]  [ 266.508269]  [ 266.508305]  [ 266.508341]  [ 266.508377]
[ 266.508413]  [ 266.508449]  [ 266.508485]  [ 266.508521]  [ 266.508557]
[ 266.508593]  [ 266.508629]  [ 266.508665]  [ 266.508701]  [ 266.508737]
[ 266.508773]  [ 266.508809]  [ 266.508845]  [ 266.508881]  [ 266.508917]
[ 266.508953]  [ 266.508989]  [ 266.509025]  [ 266.509061]  [ 266.509097]
[ 266.509133]  [ 266.509169]  [ 266.509205]  [ 266.509241]  [ 266.509277]
[ 266.509313]  [ 266.509349]  [ 266.509385]  [ 266.509421]  [ 266.509457]
[ 266.509493]  [ 266.509529]  [ 266.509565]  [ 266.509601]  [ 266.509637]
[ 266.509673]  [ 266.509709]  [ 266.509745]  [ 266.509781]  [ 266.509817]
[ 266.509853]  [ 266.509889]  [ 266.509925]  [ 266.509961]  [ 266.510000]
[ 266.510036]  [ 266.510072]  [ 266.510108]  [ 266.510144]  [ 266.510180]
[ 266.510216]  [ 266.510252]  [ 266.510288]  [ 266.510324]  [ 266.510360]
[ 266.510396]  [ 266.510432]  [ 266.510468]  [ 266.510504]  [ 266.510540]
[ 266.510576]  [ 266.510612]  [ 266.510648]  [ 266.510684]  [ 266.510720]
[ 266.510756]  [ 266.510792]  [ 266.510828]  [ 266.510864]  [ 266.510900]
[ 266.510936]  [ 266.510972]  [ 266.511008]  [ 266.511044]  [ 266.511080]
[ 266.511116]  [ 266.511152]  [ 266.511188]  [ 266.511224]  [ 266.511260]
[ 266.511296]  [ 266.511332]  [ 266.511368]  [ 266.511404]  [ 266.511440]
[ 266.511476]  [ 266.511512]  [ 266.511548]  [ 266.511584]  [ 266.511620]
[ 266.511656]  [ 266.511692]  [ 266.511728]  [ 266.511764]  [ 266.511800]
[ 266.511836]  [ 266.511872]  [ 266.511908]  [ 266.511944]  [ 266.511980]
[ 266.512016]  [ 266.512052]  [ 266.512088]  [ 266.512124]  [ 266.512160]
[ 266.512196]  [ 266.512232]  [ 266.512268]  [ 266.512304]  [ 266.512340]
[ 266.512376]  [ 266.512412]  [ 266.512448]  [ 266.512484]  [ 266.512520]
[ 266.512556]  [ 266.512592]  [ 266.512628]  [ 266.512664]  [ 266.512700]
[ 266.512736]  [ 266.512772]  [ 266.512808]  [ 266.512844]  [ 266.512880]
[ 266.512916]  [ 266.512952]  [ 266.512988]  [ 266.513024]  [ 266.513060]
[ 266.513096]  [ 266.513132]  [ 266.513168]  [ 266.513204]  [ 266.513240]
[ 266.513276]  [ 266.513312]  [ 266.513348]  [ 266.513384]  [ 266.513420]
[ 266.513456]  [ 266.513492]  [ 266.513528]  [ 266.513564]  [ 266.513600]
[ 266.513636]  [ 266.513672]  [ 266.513708]  [ 266.513744]  [ 266.513780]
[ 266.513816]  [ 266.513852]  [ 266.513888]  [ 266.513924]  [ 266.513960]
[ 266.513996]  [ 266.514032]  [ 266.514068]  [ 266.514104]  [ 266.514140]
[ 266.514176]  [ 266.514212]  [ 266.514248]  [ 266.514284]  [ 266.514320]
[ 266.514356]  [ 266.514392]  [ 266.514428]  [ 266.514464]  [ 266.514500]
[ 266.514536]  [ 266.514572]  [ 266.514608]  [ 266.514644]  [ 266.514680]
[ 266.514716]  [ 266.514752]  [ 266.514788]  [ 266.514824]  [ 266.514860]
[ 266.514896]  [ 266.514932]  [ 266.514968]  [ 266.515004]  [ 266.515040]
[ 266.515076]  [ 266.515112]  [ 266.515148]  [ 266.515184]  [ 266.515220]
[ 266.515256]  [ 266.515292]  [ 266.515328]  [ 266.515364]  [ 266.515400]
[ 266.515436]  [ 266.515472]  [ 266.515508]  [ 266.515544]  [ 266.515580]
[ 266.515616]  [ 266.515652]  [ 266.515688]  [ 266.515724]  [ 266.515760]
[ 266.515796]  [ 266.515832]  [ 266.515868]  [ 266.515904]  [ 266.515940]
[ 266.515976]  [ 266.516012]  [ 266.516048]  [ 266.516084]  [ 266.516120]
[ 266.516156]  [ 266.516192]  [ 266.516228]  [ 266.516264]  [ 266.516300]
[ 266.516336]  [ 266.516372]  [ 266.516408]  [ 266.516444]  [ 266.516480]
[ 266.516516]  [ 266.516552]  [ 266.516588]  [ 266.516624]  [ 266.516660]
[ 266.516696]  [ 266.516732]  [ 266.516768]  [ 266.516804]  [ 266.516840]
[ 266.516876]  [ 266.516912]  [ 266.516948]  [ 266.516984]  [ 266.517020]
[ 266.517056]  [ 266.517092]  [ 266.517128]  [ 266.517164]  [ 266.517200]
[ 266.517236]  [ 266.517272]  [ 266.517308]  [ 266.517344]  [ 266.517380]
[ 266.517416]  [ 266.517452]  [ 266.517488]  [ 266.517524]  [ 266.517560]
[ 266.517596]  [ 266.517632]  [ 266.517668]  [ 266.517704]  [ 266.517740]
[ 266.517776]  [ 266.517812]  [ 266.517848]  [ 266.517884]  [ 266.517920]
[ 266.517956]  [ 266.517992]  [ 266.518028]  [ 266.518064]  [ 266.518100]
[ 266.518136]  [ 266.518172]  [ 266.518208]  [ 266.518244]  [ 266.518280]
[ 266.518316]  [ 266.518352]  [ 266.518388]  [ 266.518424]  [ 266.518460]
[ 266.518496]  [ 266.518532]  [ 266.518568]  [ 266.518604]  [ 266.518640]
[ 266.518676]  [ 266.518712]  [ 266.518748]  [ 266.518784]  [ 266.518820]
[ 266.518856]  [ 266.518892]  [ 266.518928]  [ 266.518964]  [ 266.519000]
[ 266.519036]  [ 266.519072]  [ 266.519108]  [ 266.519144]  [ 266.519180]
[ 266.519216]  [ 266.519252]  [ 266.519288]  [ 266.519324]  [ 266.519360]
[ 266.519396]  [ 266.519432]  [ 266.519468]  [ 266.519504]  [ 266.519540]
[ 266.519576]  [ 266.519612]  [ 266.519648]  [ 266.519684]  [ 266.519720]
[ 266.519756]  [ 266.519792]  [ 266.519828]  [ 266.519864]  [ 266.519900]
[ 266.519936]  [ 266.519972]  [ 266.520008]  [ 266.520044]  [ 266.520080]
[ 266.520116]  [ 266.520152]  [ 266.520188]  [ 266.520224]  [ 266.520260]
[ 266.520296]  [ 266.520332]  [ 266.520368]  [ 266.520404]  [ 266.520440]
[ 266.520476]  [ 266.520512]  [ 266.520548]  [ 266.520584]  [ 266.520620]
[ 266.520656]  [ 266.520692]  [ 266.520728]  [ 266.520764]  [ 266.520800]
[ 266.520836]  [ 266.520872]  [ 266.520908]  [ 266.520944]  [ 266.520980]
[ 266.521016]  [ 266.521052]  [ 266.521088]  [ 266.521124]  [ 266.521160]
[ 266.521196]  [ 266.521232]  [ 266.521268]  [ 266.521304]  [ 266.521340]
[ 266.521376]  [ 266.521412]  [ 266.521448]  [ 266.521484]  [ 266.521520]
[ 266.521556]  [ 266.521592]  [ 266.521628]  [ 266.521664]  [ 266.521700]
[ 266.521736]  [ 266.52
```

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd64_radeon amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD B150/INA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3cf30 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffff8800916f8000+0x26630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492240] RSP: 0018:ffff8800916f8000 EFLAGS: 00010286
[ 266.492288] RAX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492336] R00: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492384] R01: 0000000000000002 R09: 0000000000000000
[ 266.492432] R02: 0000000000000001 R12: ffff88008f686068
[ 266.492480] R03: ffff88008f6b3b90 R15: ffff88008f686000
[ 266.492528] R04: ffff88011ed80000(0000)
[ 266.492576] R05: 0000000080050033
[ 266.492624] R06: 000000001b31000 CR4: 00000000000407e0
[ 266.492672] R07: 0000000036078 ffff8800916f8000 ffff88008f686000
[ 266.492720] R08: 0000000000000000 ffff88008f686000
[ 266.492768] R09: ffff8801ba5d 00007fb885918000 ffff88008edf3000
[ 266.492816] ..
```

Mostly useful when it is known which drivers are built as modules (e.g. with standard distro kernel configs).

May also contain module taint flags:

P – proprietary

O – out-of-tree

F – force-loaded

C – staging

E – unsigned

X – external

+/- – being loaded/unloaded

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800003b3c840 ti: ffff88000016f8000
[ 266.492191] RIP: 0010:[<f... ffff81126...
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800003b3c840
[ 266.492285] RAX: ffff880008f6b3c98
[ 266.492331] RDX: ffff88000914b3c98
[ 266.492376] RBP: ffff88000916fbbf8
[ 266.492421] R10: 0000000000000000
[ 266.492465] R13: ffff88000914b3c98 R
[ 266.492513] FS: 00007fb8966f6700(00
knlGS:0000000000000000
[ 266.492566] CS: 0010 DS: 0000 ES: 0
[ 266.492601] CR2: 00007f50fa190770 CR
[ 266.492652] Stack:
[ 266.492665] 0000000000000000 ffff8800003b3c840
[ 266.492714] ffff88000916fbc08 00000000
[ 266.492764] ffff88000916fbbf8 ffffffff
...
```

Information about CPU, process, kernel version, hardware.

Taint flags:

- POFCEX – same as per-module
- R – module was force-unloaded
- M – system has reported a MCE
- B – bad page was encountered
- U – userspace-defined
- D – there was an oops before
- W – there was a warning before
- A – ACPI table was overridden
- I – firmware bug workaround
- L – soft-lockup has occurred before
- K – kernel has been live patched
- S – SMP kernel on UP machine

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLA: 0000000000000000
[ 266.492285] RAX: ffff88008f6b3ba0 RBX: ffff8800916fbb68 RCX: ffff8800a3b3cf30
[ 266.492331] RDX: ffff8800914b3c98 RSI: ffff8800916fbb68
[ 266.492376] RBP: ffff8800916fbb68 R08: ffff8800916fbb68
[ 266.492421] R10: 0000000000000008 R11: ffff8800916fbb68
[ 266.492465] R13: ffff8800914b3c98 R14: ffff8800916fbb68
[ 266.492513] FS: 00007fb8966f6700(0000000000000000)
knlGS:0000000000000000
[ 266.492566] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[ 266.492601] CR2: 00007f50fa190770 CR3: 0000000001b31000 CR4: 00000000000407e0
[ 266.492652] Stack:
[ 266.492665] 0000000000000000 ffff88008f686078 ffff8800916fbb68 ffff88008f686000
[ 266.492714] ffff8800916fbc08 0000000000000000 0000000000000000 ffff88008f686000
[ 266.492764] ffff8800916fbbf8 ffffffff8111ba5d 00007fb885918000 ffff88008edf3000
...
```

Information about task that's supposed to be currently running, and whose stack we are actually running on.



# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd64_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLAGS: 00010286
[ 266.492285] RAX: ffff8800916fbb68 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492331] RDX: ffff8800916fbb68 RSI: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492376] RBP: ffff8800916fbb68 R00: 0000000000000000
[ 266.492421] R10: ffff88008f686068
[ 266.492465] R13: ffff88008f686000
[ 266.492513] FS: 0000000000000000
knlGS:0000000000000000
[ 266.492566] CS: 0000000000000000
[ 266.492601] CR2: 0000000000407e0
[ 266.492652] Stack: ffff88008f686000
[ 266.492665] 0000000000000000 ffff88008f686000
[ 266.492714] ffff88008f686000 ffff88008f686000
[ 266.492764] ffff88008f686000 ffff88008edf3000
...
```

Which instruction was executing, translated to function name + offset.

This may be different from where position where BUG\_ON() was reported, if the function containing BUG\_ON was inlined.



# Example kernel oop

Values for the rest of the registers at the trapping instruction. Some are clearly kernel addresses. Some may hold the bad value of anon\_vma->degree. Maybe RSI, R08, R10 or R11?

```
[ 266.491864] -----[ cut here ]--
[ 266.491904] kernel BUG at mm/rmap.c:395!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.10-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS BL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbb68 EFLAGS: 00010286
[ 266.492285] RAX: ffff88008f6b3ba0 RBX: ffff88008f6b3b90 RCX: ffff8800a3b3cf30
[ 266.492331] RDX: ffff8800914b3c98 RSI: 0000000000000001 RDI: ffff8800914b3c98
[ 266.492376] RBP: ffff8800916fbbba8 R08: 0000000000000002 R09: 0000000000000000
[ 266.492421] R10: 0000000000000008 R11: 0000000000000001 R12: ffff88008f686068
[ 266.492465] R13: ffff8800914b3c98 R14: ffff88008f6b3b90 R15: ffff88008f686000
[ 266.492513] FS: 00007fb8966f6700(0000) GS:ffff88011ed80000(0000)
knlGS:0000000000000000
[ 266.492566] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[ 266.492601] CR2: 00007f50fa190770 CR3: 0000000001b31000 CR4: 00000000000407e0
[ 266.492652] Stack:
[ 266.492665] 0000000000000000 ffff88008f686078 ffff8800916fbbba8 ffff88008f686000
[ 266.492714] ffff8800916fbc08 0000000000000000 0000000000000000 ffff88008f686000
[ 266.492764] ffff8800916fbbf8 ffffffff8111ba5d 00007fb885918000 ffff88008edf3000
...
```

# Example kernel oops

```
[ 266.491864] -----[ cut here ]-----
[ 266.491904] kernel BUG at mm/rmap.c:399!
[ 266.491934] invalid opcode: 0000 [#1] SMP
[ 266.491962] Modules linked in: amd_kfd amd_iommu_v2 radeon cfbfillrect cfbimgblt
cfbcopyarea drm_kms_helper ttm fuse
[ 266.492043] CPU: 3 PID: 5155 Comm: java Not tainted 3.19.0-rc3-kfd+ #24
[ 266.492087] Hardware name: AMD BALLINA/Ballina, BIOS WBL3B20N_Weekly_13_11_2
11/20/2013
[ 266.492141] task: ffff8800a3b3c840 ti: ffff8800916f8000 task.ti: ffff8800916f8000
[ 266.492191] RIP: 0010:[<ffffffff81126630>] [<ffffffff81126630>]
unlink_anon_vmas+0x102/0x159
[ 266.492249] RSP: 0018:ffff8800916fbbf8
[ 266.492285] RAX: ffff88008f6b3ba0 RBX: ffff8800916fbbf8 RCX: ffff8800916fbbf8 RDX: ffff8800916fbbf8
[ 266.492331] RDX: ffff8800914b3c98 RSI: ffff8800914b3c98 RDI: ffff8800914b3c98
[ 266.492376] RBP: ffff8800916fbbf8 R08: 00000000 R09: 00000000 R10: 00000000 R11: 00000000 R12: ffff88008f686068
[ 266.492421] R10: 00000000 R11: 00000000 R12: ffff88008f686068 R13: ffff8800914b3c98 R14: ffff8800914b3c98 R15: ffff88008f686000
[ 266.492465] R13: ffff8800914b3c98 R14: ffff8800914b3c98 R15: ffff88008f686000
[ 266.492513] FS: 00007fb8966f6700(0000) GS: ffff8800916fbbf8(0000)
knlGS:0000000000000000
[ 266.492566] CS: 0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[ 266.492601] CR2: 00007f50fa190770 CR3: 00000000b31000 CR4: 0000000000407e0
[ 266.492652] Stack:
[ 266.492665] 0000000000000000 ffff88008f686078 ffff8800916fbbf8 ffff88008f686000
[ 266.492714] ffff8800916fbc08 0000000000000000 0000000000000000 ffff88008f686000
[ 266.492764] ffff8800916fbbf8 ffffffff8111ba5d 00007fb885918000 ffff88008edf3000
```

Raw contents of top of the stack  
starting at RSP

...

# Example kernel oops

Backtrace reconstructed by unwinding the stack, showing the return addresses from individual call frames

“?” means a pointer to function is on stack but doesn't fit in the frame; could be leftover from previous execution or heuristics failure

```
..
[ 266.492815] Call Trace:
[ 266.492834] [] free_pgtables+0x8e/0xcc
[ 266.492873] [] exit_mmap+0x84/0x116
[ 266.492907] [] mmput+0x52/0xe9
[ 266.492940] [] do_exit+0x3cd/0x9c9
[ 266.492975] [] ? _raw_spin_unlock_irq+0x2d/0x32
[ 266.493016] [] do_group_exit+0x4c/0xc9
[ 266.493051] [] get_signal+0x58f/0x5bc
[ 266.493090] [] do_signal+0x28/0x5b1
[ 266.493123] [] ? sysret_signal+0x5/0x43
[ 266.493162] [] do_notify_resume+0x35/0x68
[ 266.493200] [] int_signal+0x12/0x17
[ 266.493235] Code: e8 03 b7 f4 ff 49 8b 47 78 4c 8b 20 48 8d 58 f0 49 83
ec 10 48 8d 43 10 48 39 45 c8 74 55 48 8b 7b 08 83 bf 8c 00 00 00 00 74 02
<0f> 0b e8 a4 fd ff ff 48 8b 43 18 48 8b 53 10 48 89 df 48 89 42
[ 266.493404] RIP [] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

# Example kernel oops

```
..
[ 266.492815] Call Trace:
[ 266.492834] [] ? sysret_signal+0x5/0x43
[ 266.492873] [] do_notify_resume+0x35/0x68
[ 266.492907] [] int_signal+0x12/0x17
[ 266.492940] [] ? sysret_signal+0x5/0x43
[ 266.492975] [] ? sysret_signal+0x5/0x43
[ 266.493016] [] ? sysret_signal+0x5/0x43
[ 266.493051] [] ? sysret_signal+0x5/0x43
[ 266.493090] [] ? sysret_signal+0x5/0x43
[ 266.493123] [] ? sysret_signal+0x5/0x43
[ 266.493162] [] do_notify_resume+0x35/0x68
[ 266.493200] [] ? sysret_signal+0x5/0x43
[ 266.493235] Code: e8 03 b7 f4 ff 49 8b 47 78 4c 8b 20 48 8d 58 f0 49 83
ec 10 48 8d 43 10 48 39 45 c8 74 55 48 8b 7b 08 83 bf 8c 00 00 00 00 74 02
<0f> 0b e8 a4 fd ff ff 48 8b 43 18 48 8b 53 10 48 89 df 48 89 42
[ 266.493404] RIP [] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

A bunch of instructions around the RIP.  
RIP position denoted by < >

Recall that 0F 0B is opcode for UD2

We can disassemble the code listing by  
piping the oops into  
./scripts/decodecode  
in the kernel source tree.

# Example decodecode output

```
~/linux.git> ./scripts/decodecode < oops-example.txt
```

```
[ 266.493235] Code: e8 03 b7 f4 ff 49 8b 47 78 4c 8b 20 48 8d 58 f0 49 83 ec 10 48 8d 43 10 48 39 45 c8 74 55  
48 8b 7b 08 83 bf 8c 00 00 00 00 74 02 <0f> 0b e8 a4 fd ff ff 48 8b 43 18 48 8b 53 10 48 89 df 48 89 42
```

All code

=====

```
0: e8 03 b7 f4 ff      callq 0xffffffffffff4b708  
5: 49 8b 47 78        mov 0x78(%r15),%rax  
9: 4c 8b 20          mov (%rax),%r12  
c: 48 8d 58 f0      lea -0x10(%rax),%rbx  
10: 49 83 ec 10     sub $0x10,%r12  
14: 48 8d 43 10     lea 0x10(%rbx),%rax  
18: 48 39 45 c8     cmp %rax,-0x38(%rbp)  
1c: 74 55          je 0x73  
1e: 48 8b 7b 08     mov 0x8(%rbx),%rdi  
22: 83 bf 8c 00 00 00 00  cmpl $0x0,0x8c(%rdi)  
29: 74 02          je 0x2d  
2b:* 0f 0b          ud2 <-- trapping instruction  
2d: e8 a4 fd ff ff     callq 0xffffffffffffdd6  
32: 48 8b 43 18     mov 0x18(%rbx),%rax  
36: 48 8b 53 10     mov 0x10(%rbx),%rdx  
3a: 48 89 df        mov %rbx,%rdi  
3d: 48          rex.W  
3e: 89          .byte 0x89  
3f: 42          rex.X
```

Code starting with the faulting instruction

=====

```
0: 0f 0b          ud2  
2: e8 a4 fd ff ff     callq 0xffffffffffffdab  
7: 48 8b 43 18     mov 0x18(%rbx),%rax  
b: 48 8b 53 10     mov 0x10(%rbx),%rdx  
f: 48 89 df        mov %rbx,%rdi  
12: 48          rex.W  
13: 89          .byte 0x89  
14: 42          rex.X
```



# Example decodecode output

All code

=====

```
0: e8 03 b7 f4 ff callq 0xffffffffffff4b708
5: 49 8b 47 78 mov 0x78(%r15),%rax
9: 4c 8b 20 mov (%rax),%r12
c: 48 8d 58 f0 lea -0x10(%rax),%rbx
10: 49 83 ec 10 sub $0x10,%r12
14: 48 8d 43 10 lea 0x10(%rbx),%rax
18: 48 39 45 c8 cmp %rax,-0x38(%rbp)
1c: 74 55 je 0x73
1e: 48 8b 7b 08 mov 0x8(%rbx),%rdi
22: 83 bf 8c 00 00 00 00 cmpl $0x0,0x8c(%rdi)
29: 74 02 je 0x2d
2b:* 0f 0b ud2 <-- trapping instruction
2d: e8 a4 fd ff ff callq 0xffffffffffffdd6
32: 48 8b 43 18 mov 0x18(%rbx),%rax
36: 48 8b 53 10 mov 0x10(%rbx),%rdx
3a: 48 89 df mov %rbx,%rdi
3d: 48 rex.W
3e: 89 .byte 0x89
3f: 42 rex.X
```

# Example 1: Output

```
BUG_ON(anon_vma->degree);
```

We skip the UD2 instruction if `0x8c(%rdi)` equals zero → we trap if the value is non-zero

This suggests that RDI holds the struct `anon_vma` pointer and `degree` is at offset `0x8c`

However, we can't determine the value that had been compared to zero in this case...

All code

=====

```
0:  e8 8b 7b 08 f4b708  callq 0xf4b708
5:  48 8b 7b 08      mov    0x8(%rbx),%rax
9:  48 8b 7b 08      mov    0x8(%rbx),%rax
c:  48 8b 7b 08      mov    0x8(%rbx),%rax
10: 48 8b 7b 08      mov    0x8(%rbx),%rax
14: 48 8b 7b 08      mov    0x8(%rbx),%rax
18: 48 39 45 c8      cmpl  0x38(%rbp),%rax
1c: 74 55          je     0x73
1e: 48 8b 7b 08      mov    0x8(%rbx),%rdi
22: 83 bf 8c 00 00 00 00  cmpl  $0x0,0x8c(%rdi)
29: 74 02          je     0x2d
2b:* 0f 0b          ud2                                <-- trapping instruction
2d: e8 a4 fd ff ff      callq 0xffffffffffffdd6
32: 48 8b 43 18      mov    0x18(%rbx),%rax
36: 48 8b 53 10      mov    0x10(%rbx),%rdx
3a: 48 89 df          mov    %rbx,%rdi
3d: 48              rex.w
3e: 89              .byte 0x89
3f: 42              rex.X
```



# Example decodecode output

All code

=====

```
0: e8 03 b7 f4 ff
5: 49 8b 47 78
9: 4c 8b 20
c: 48 8d 58 f0
10: 49 83 ec 10
14: 48 8d 43 10
18: 48 39 45 c8
1c: 74 55
1e: 48 8b 7b 08
22: 83 bf 8c 00 00 00 00
29: 74 02
2b:* 0f 0b
2d: e8 a4 fd ff ff
32: 48 8b 43 18
36: 48 8b 53 10
3a: 48 89 df
3d: 48
3e: 89
3f: 42
```

```

struct anon_vma_chain *avc;
list_for_each_entry_safe(avc, ...)
    struct anon_vma *anon_vma = avc->anon_vma;
    BUG_ON(anon_vma->degree);

su
lea
cmp
je
mov 0x8(%rbx),%rdi
cmpl $0x0,0x8c(%rdi)
je 0x2d
ud2 <-- trapping instruction
callq 0xffffffffffffdd6
mov 0x18(%rbx),%rax
mov 0x10(%rbx),%rdx
mov %rbx,%rdi
rex.W
.byte 0x89
rex.X
```

```
struct anon_vma_chain *avc;
list_for_each_entry_safe(avc, ...)
    struct anon_vma *anon_vma = avc->anon_vma;
    BUG_ON(anon_vma->degree);
```

Suggests that RBX holds the struct anon\_vma\_chain pointer avc and anon\_vma member is at offset 0x8



# Verifying structure offsets

- We can use pahole from dwarves package
  - May depend on GCC version, .config options
    - rwsem size depends on CONFIG\_DEBUG\_SPINLOCK, CONFIG\_DEBUG\_LOCK\_ALLOC

```
> pahole --hex -C anon_vma mm/vmscan.o
struct anon_vma {
    struct anon_vma *      root;                /*      0      0x8 */
    struct rw_semaphore    rwsem;              /*     0x8    0x80 */
    /* --- cacheline 2 boundary (128 bytes) was 8 bytes ago --- */
    atomic_t               refcount;           /*    0x88    0x4 */
    unsigned int            degree;            /*    0x8c    0x4 */
    struct anon_vma *      parent;             /*    0x90    0x8 */
    struct rb_root          rb_root;           /*    0x98    0x8 */

    /* size: 160, cachelines: 3, members: 6 */
    /* last cacheline: 32 bytes */
};
```

# Example kernel oops

```
...
[ 266.492815] Call Trace:
[ 266.492834] [] free_pgtables+0x8e/0xcc
[ 266.492873] [] exit_mmap+0x84/0x116
[ 266.492907] [] mmput+0x52/0xe9
[ 266.492940] [] do_exit+0x3cd/0x9c9
[ 266.492975] [] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

The most important registers again, with higher printk level, or in case the details had scrolled away.

# Example kernel oops

```
...
[ 266.492815] Call Trace:
[ 266.492834] [] free_pgtables+0x8e/0xcc
[ 266.492873] [] exit_mmap+0x84/0x116
[ 266.492907] [] mmput+0x52/0xe9
[ 266.492940] [] do_exit+0x3cd/0x9c9
[ 266.492975] []
[ 266.493200] [] in
[ 266.493235] Code: e8 03 b7 f4 ff 49 4c 8b 20 48 8d 58 f0 49 83
ec 10 48 8d 43 10 48 39 45 c8 74 55 48 08 83 bf 8c 00 00 00 00 74 02
<0f> 0b e8 a4 fd ff ff 48 8b 43 18 48 8 10 48 89 df 48 89 42
[ 266.493404] RIP [] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

Randomization to distinguish reports of same bug instance from separate instances.

# Example kernel oops

The task was already exiting when it oopsed. In this case it's clearly graceful exit (from the backtrace), but it could be exiting due to previous oops. It's safer to leave task as zombie than to risk infinite loops in the exit path.

```
...
[ 266.49294] [ffff81043918] do_exit+0x3cd/0x9c9
[ 266.49297] [ffff8170c1ec] ? _raw_spin_unlock_irq+0x2d/0x32
[ 266.493016] [ffff81044d7f] do_group_exit+0x4c/0xc9
[ 266.493051] [ffff8104eb87] get_signal+0x58f/0x5bc
[ 266.493090] [ffff810022c4] do_signal+0x28/0x5b1
[ 266.493123] [ffff8170ca0c] ? sysret_signal+0x5/0x43
[ 266.493162] [ffff81002882] do_notify_resume+0x35/0x68
[ 266.493200] [ffff8170cc7f] int_signal+0x12/0x17
[ 266.493235] : e8 03 b7 f4 ff 49 8b 47 78 4c 8b 20 48 8d 58 f0 49 83
ec 10 48 8d 43 8 39 45 c8 74 55 48 8b 7b 08 83 bf 8c 00 00 00 00 74 02
<0f> 0b e8 a4 f0 ff 48 8b 43 18 48 8b 53 10 48 89 df 48 89 42
[ 266.493404] RSP [ffffffffff81126630] unlink_anon_vmas+0x102/0x159
[ 266.493447] RSP <ffff8800916fbb68>
[ 266.508877] ---[ end trace 02d28fe9b3de2e1a ]---
[ 266.508880] Fixing recursive fault but reboot is needed!
```

# How is stack unwinding implemented?

- Start at value of RSP and increment in a loop
  - Check if stack contains kernel text address
  - Print with translation to function name+offset
  - When RSP matches  $RBP + \text{sizeof}(\text{long})$ , consider address *reliable* (i.e. without “?”) and update RBP from the address it points to
- Not fully reliable, even with frame pointers
  - Cannot be relied upon functionally (live patching?)
  - Assembler functions now audited for missing frame pointers
    - Planned: runtime checks + DWARF validations

# How is stack unwinding implemented?

- For perf callgraph sampling, this would be slow
  - Therefore, fully rely on frame pointer walk there
- Alternative approach: use DWARF2 exception handler (EH) frame info
  - Patch in SUSE kernels, rejected upstream
  - Also not fully reliable, and more complex

# What else can produce oops/panic?

- BUG\_ON seen in the example – hard assertion
- Memory paging related faults
  - “BUG: unable to handle kernel paging request”
    - “... handle NULL pointer dereference” (when `bad_addr < PAGE_SIZE`)
  - Corrupted page table
  - Kernel trying to execute NX-protected page
  - Kernel trying to execute userspace page (Intel SMEP)
  - Failed bounds check in kernel mode (Intel MPX feature)
  - General protection fault, unhandled double fault

# What else can produce oops/panic?

- Soft lockup
  - CPU spent 20s in kernel without reaching a schedule point
  - A warning, unless config/bootparam `softlockup_panic` enabled
    - Soft lockup can be harmless, so not good idea in production
- Hard lockup
  - CPU spent 10s with disabled interrupts
- Detection of both combines several generic mechanisms
  - High priority kernel watchdog thread updates soft lockup timestamp
  - hrtimer set to deliver periodic interrupts, increments hard lockup counter and wakes up the watchdog thread
  - NMI perf event checks if hrtimers interrupts were processed and if watchdog thread was scheduled



# What else can produce oops/panic?

- Hung task check
  - “INFO: task ... blocked for more than 120 seconds”
  - khungtaskd periodically processes tasks in uninterruptible sleep and checks if their switch count changed
- RCU stall detector
  - Detects when RCU grace period is too long (21s)
    - CPU looping in RCU critical section or disabled interrupts, preemption or bottom halves, no scheduling points in non-preempt kernels
    - RT task preempting non-RT task in RCU critical section
- Several other debugging config options (later)

# Obtaining crash dumps

- Several historical methods
  - diskdump, netdump, LKCD
  - Not very reliable (some parts of crashed kernel must still work) or universal, needs dedicated server on same network etc.
  - Out of tree patches, included in old enterprise distros
- Current solution: kexec-based **kdump**
  - Crash kernel loaded into a boot-reserved memory area
  - On panic, kexec switches to the crash kernel without reboot
  - Memory of crashed kernel available as /proc/vmcore
  - Kdump utility can save to disk, network, filter pages...
    - kexec(8), kdump(5), makedumpfile(8)

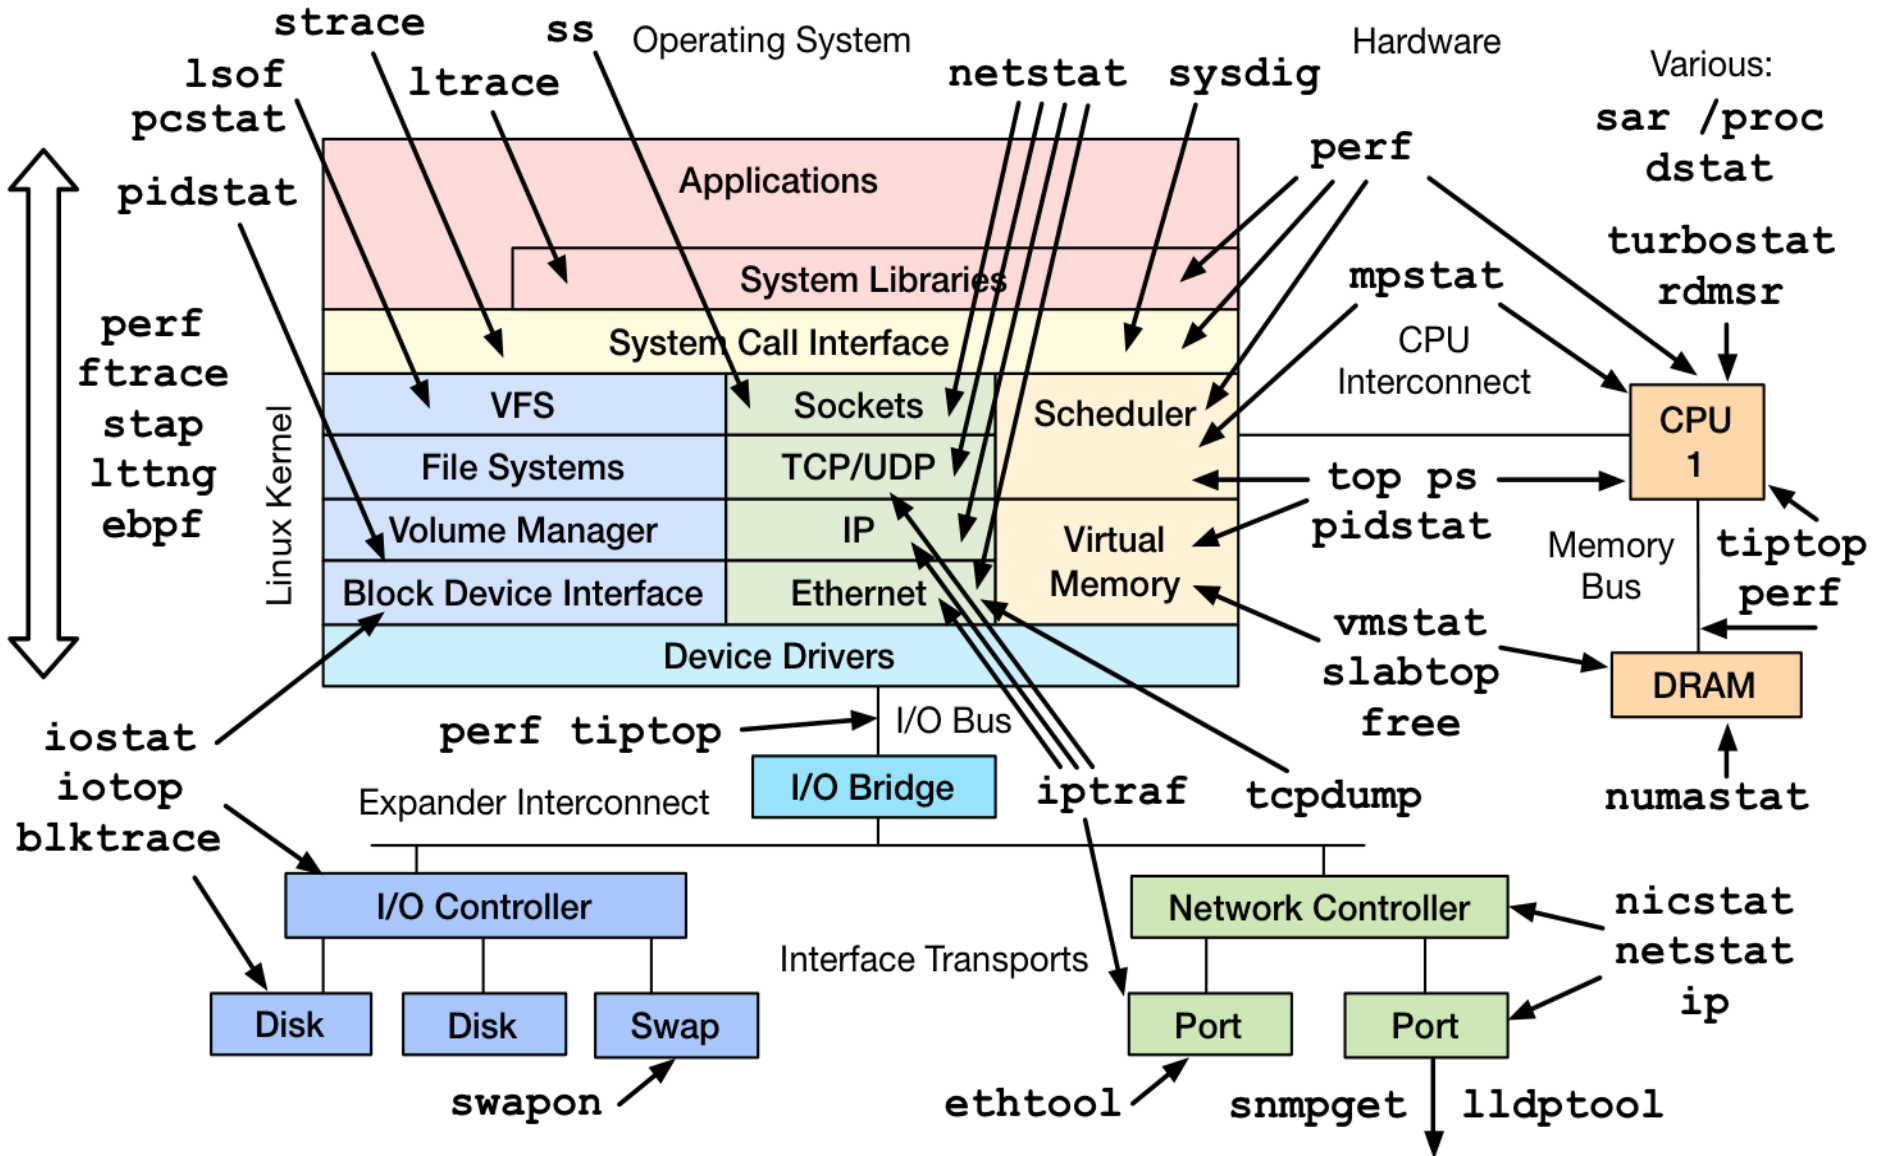
# Analyzing crash dumps with gdb

- gdb can be used to open ELF based dumps
  - But those are not easily compressed and filtered
- gdb has no understanding of kernel internals or virtual/physical mapping
  - There are some recently added Python scripts under `scripts/gdb` in the Linux source
    - Can obtain per-cpu variables, dmesg, modules, tasks
- A better tool for Linux kernel crash dumps - **crash**

# Debugging live systems

- Observability tools – separate lecture?
- Sysrq magic host keys
- Live debugger support
- Ftrace, kprobes, SystemTap, perf – separate lecture?
- Debugging config options

# Linux Performance Observability Tools



<http://www.brendangregg.com/linuxperf.html> 2015

# Standard debugging means

- Magic SysRq hot keys
- For dealing with hangs and security issues
  - Operator's intervention to the running system
  - Can be enabled/disabled by `/proc/sys/kernel/sysrq`
    - Alt + SysRq + 0 .. 9 set console logging level
    - Alt + SysRq + H show help
    - Alt + SysRq + C crash by a NULL pointer dereference
    - Alt + SysRq + B immediate reboot
    - Alt + SysRq + O immediate shutdown
    - Alt + SysRq + S sync all mounted filesystems
    - Alt + SysRq + U remount all filesystems read-only
    - Alt + SysRq + J freeze filesystems by FIFREEZE ioctl

# Standard debugging means (2)

- Alt + SysRq + P dump registers to console
- Alt + SysRq + T dump process information to console
- Alt + SysRq + L dump stack traces of running threads
- Alt + SysRq + M dump memory statistics to console
- Alt + SysRq + D dump locked locks to console
- Alt + SysRq + K kill all processes on the current console
- Alt + SysRq + E terminate all processes except `init`
- Alt + SysRq + I kill all processes except `init`
- Alt + SysRq + F execute the OOM killer
- Alt + SysRq + N reset nice level of all real-time processes
- Alt + SysRq + R switch off raw keyboard mode
- Alt + SysRq + Q dump armed hritmers, clockevent devices
- Alt + SysRq + V forcefully restore framebuffer console
- Alt + SysRq + W dump tasks in uninterruptible sleep
- Alt + SysRq + Z dump the ftrace buffer

# Standard debugging means (3)

- Activate from command line by writing into `/proc/sysrq-trigger`
- Activate over network by a special `sysrqd` server
- **Raising Elephants Is So Utterly Boring**  
**Reboot Even If System Utterly Broken**
  - Raw keyboard
  - Send SIGTERM to all processes
  - Send SIGKILL to all processes
  - Sync data to disk
  - Remount all filesystems read-only
  - Reboot



# Live kernel debugging - /proc/kcore

- `/proc/kcore` enabled by `CONFIG_PROC_KCORE`
  - Provides virtual ELF “core dump” file
  - Usable by `gdb` and `crash` for read-only inspection

# Live kernel debugging - kgdb

- **kgdb** was merged in 2.6.26 (2008)
- Provides a server for remote gdb client
  - Over serial port – CONFIG\_KGDB\_SERIAL\_CONSOLE
  - Over network using NETPOLL – not mainline (KDBoE)
- Enable on server
  - Boot with kgdboc=ttyS0,115200
  - echo g > /proc/sysrq-trigger or kgdbwait boot param
- Use from a client
  - % gdb ./vmlinux
  - (gdb) set remotebaud 115200
  - (gdb) target remote /dev/ttyS0
  - Allows limited gdb debugging similar to a userspace program

# Live kernel debugging - kdb

- **kdb** is a frontend for kgdb that runs in the debugged kernel (no need for other client) – since 2.6.35 (2010)
- Provides a shell accessed via serial terminal, with optional PS/2 keyboard support
  - Enabled same way as the kgdb server
  - Switch between kdb/kgdb by `$3#33` and kgdb
- Provides some kernel-specific commands not available in pure gdb
  - `lsmod`, `ps`, `ps A`, `summary`, `bt`, `dmesg`, `go`, `help`
  - Some can be executed from gdb – `monitor help`
- Out of tree discontinued version seemed to be more capable
- KMS console support was proposed, but dropped

# Live debugging - User-Mode Linux

- UML
  - Special pseudo-hardware architecture
    - Otherwise compatible with the target architecture
    - Running Linux kernel as a user space process
      - Originally a virtualization effort
    - Great for debugging and kernel development
      - A plain standard gdb can be used to attach to the running kernel
      - Guest threads are threads of the UML process
        - Slightly more complicated to follow processes

# Kernel debugging config options

- Kernel can be built with additional debugging options enabled
  - Extra checks that can catch errors sooner, or provide extra information, at the cost of CPU and/or memory overhead
  - Can also hide errors such as race conditions...
- Many of them under “Kernel hacking” in make menuconfig
  - Others placed in the given subsystem/driver

# Kernel debugging config options (VM)

- `DEBUG_VM` – enable `VM_BUG_ON(cond)` checks
- `PAGE_OWNER` – track who allocated which pages in order to find a memory leak
- `DEBUG_PAGEALLOC` – unmap (or poison) pages after they are freed
- `DEBUG_SLAB` – detect some cases of double free, or use-after-free (by poisoning)
  - `SLUB_DEBUG` variant can enable/disable debugging in runtime
- `DEBUG_KMEMLEAK` – detect leaks with a conservative garbage collection based algorithm
- `KASAN` – Find out of bounds accesses and use-after-free bugs at the cost of 1/8 memory and 3x slower performance

# Kernel debugging config options

- `DEBUG_STACKOVERFLOW` – check if random corruption involving struct `thread_info` is caused by too deep call chains
- `DEBUG_SPINLOCK` and others for different locks – catch missing init, freeing of live locks, some deadlocks
- `LOCK_STAT` – for lock contention, perf lock
- `PROVE_LOCKING` - “lockdep” mechanism for online proving that deadlocks cannot happen and report that deadlock can occur before it actually does