

Reverse Engineering NAND Flash

Adapted from

Josh 'm0nk' Thomas's Black Hat Presentation

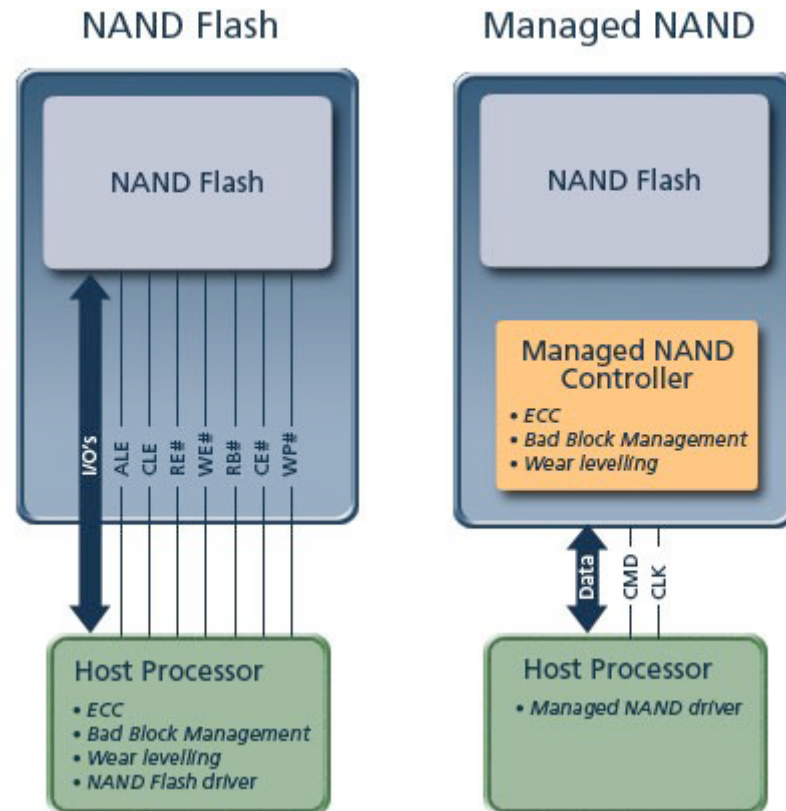
Andrew 'bunnie' Huang & Sean 'xobs' Cross 30c3 Presentation

Presented by Ben Ruktantichoke

NAND:Hard It Work

- Floating gate transistor
- Pages – Typically 512, 2048, or 4096
- Blocks – Typically 16kb – 512kb
- Shifting to 0 is easy
- Shifting to 1 is hard

NAND:Hard it Work





Faking Reliability

- Flash memory is “unreliable”
- You are not storing data, you are storing probabilistic approximations of your data
- Workaround: computational error correction (ECC)
- Flash geometry changes
 - New ECC rules, page size, block mapping, etc.

Also, Bad Blocks

- TLC/MLC Flash is super cheap
- Work around: bad block remapping
 - In some cases, over 80% of blocks are bad (e.g. 16GB chip sold as 2GB)
- Also, blocks go bad with P/E cycles

What's inside



ifixit

NAND:Soft it Works

- RAW NAND vs. MMC/eMMC
 - Complex Driver vs. Simple Driver
- Proprietary (closed) wear leveling algorithms are normally embedded

NAND:Soft it Works


- MTD Subsystem
- Kind of a meta-driver
- Used heavily for boot partitions on Android



Related Works

- MTD at the Driver Level
- Flash Transition Layers and Reverse the Embedded Controllers

Digilent Nexys™2 Spartan-3E FPGA + Schmartboard



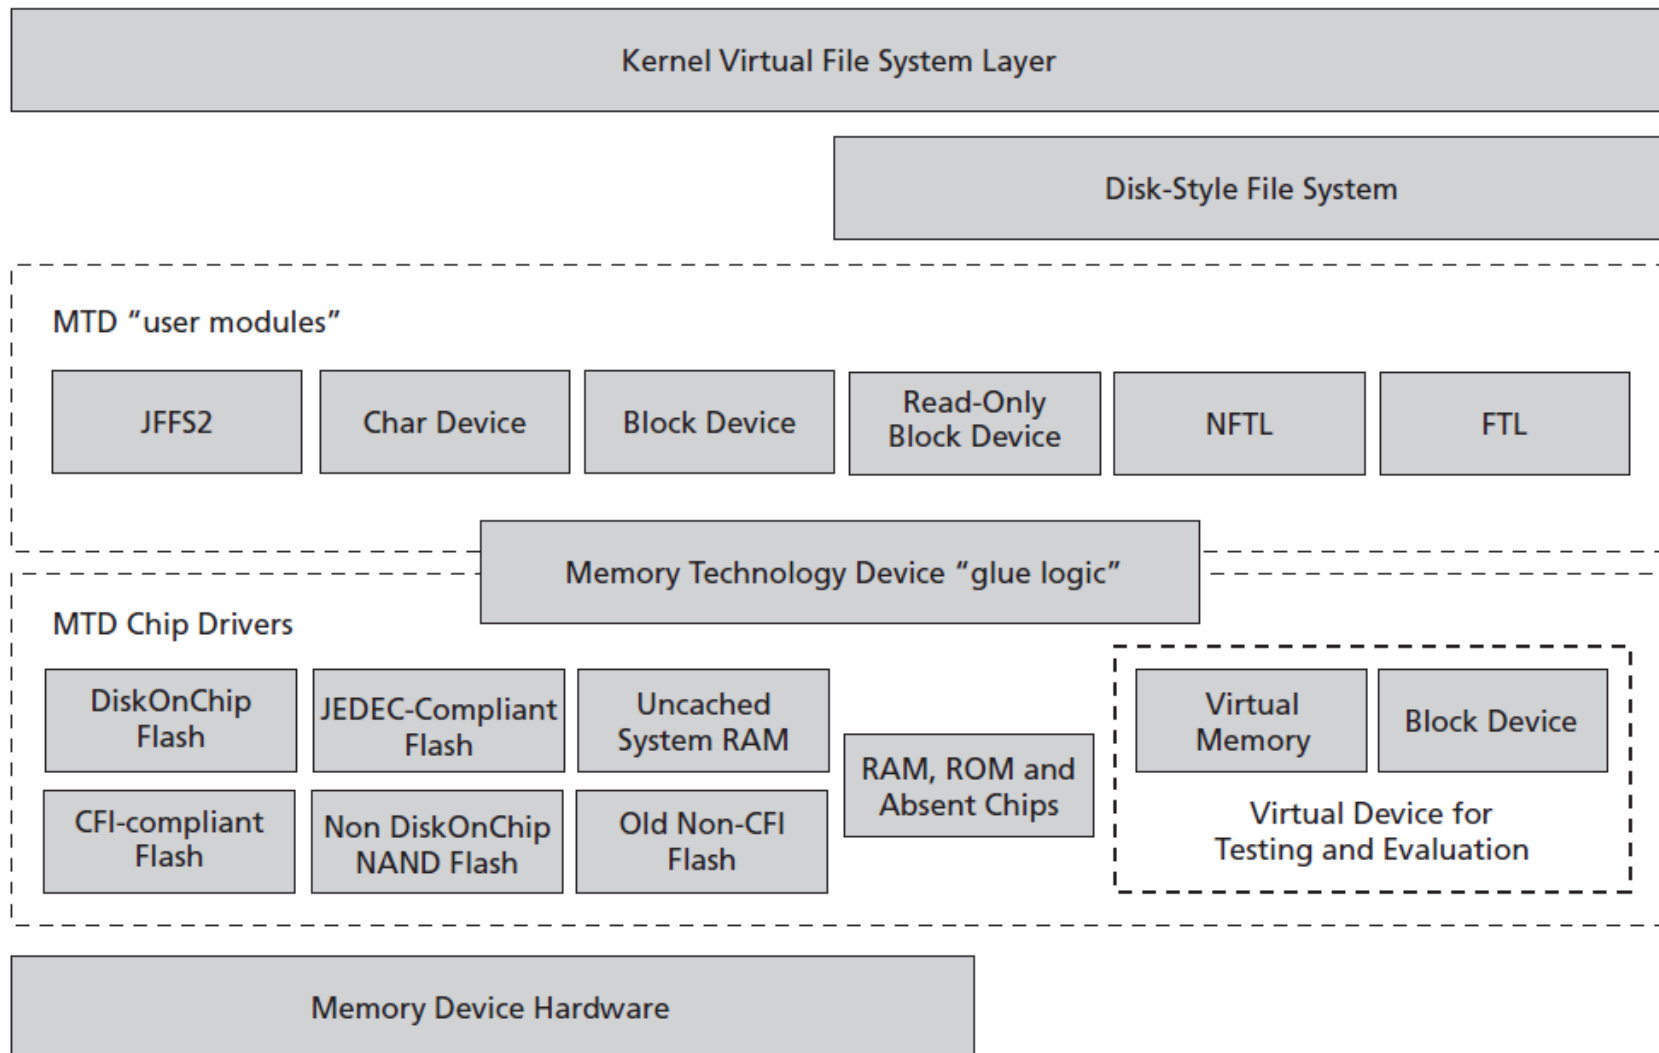
What's an initial RAM disk?

- The *initial RAM disk (initrd)* is an initial root file system that is mounted prior to when the real root file system is available. The *initrd* is bound to the kernel and loaded as part of the kernel boot procedure. The kernel then mounts this *initrd* as part of the two-stage boot process to load the modules to make the real file systems available and get at the real root file system.
- The *initrd* contains a minimal set of directories and executables to achieve this, such as the *insmod* tool to install kernel modules into the kernel.
- In the case of desktop or server Linux systems, the *initrd* is a transient file system. Its lifetime is short, only serving as a bridge to the real root file system. In embedded systems with no mutable storage, the *initrd* is the permanent root file system.

Booting with an initial RAM disk


- The boot loader, such as GRUB, identifies the kernel that is to be loaded and copies this kernel image and any associated initrd into memory.
- After the kernel and initrd images are decompressed and copied into memory, the kernel is invoked. Various initialization is performed and, eventually, you find yourself in `init/main.c:init()` (`subdir/file:function`). This function performs a large amount of subsystem initialization. A call is made here to `init/do_mounts.c:prepare_namespace()`, which is used to prepare the namespace (mount the dev file system, RAID, or md, devices, and, finally, the initrd). Loading the initrd is done through a call to `init/do_mounts_initrd.c:initrd_load()`.
- The `initrd_load()` function calls `init/do_mounts_rd.c:rd_load_image()`, which determines the RAM disk image to load through a call to `init/do_mounts_rd.c:identify_ramdisk_image()`. This function checks the magic number of the image to determine if it's a minux, etc2, romfs, cramfs, or gzip format. Upon return to `initrd_load_image`, a call is made to `init/do_mounts_rd.c:rd_load()`. This function allocates space for the RAM disk, calculates the cyclic redundancy check (CRC), and then uncompresses and loads the RAM disk image into memory. At this point, you have the initrd image in a block device suitable for mounting.
- Mounting the block device now as root begins with a call to `init/do_mounts.c:mount_root()`. The root device is created, and then a call is made to `init/do_mounts.c:mount_block_root()`. From here, `init/do_mounts.c:do_mount_root()` is called, which calls `fs/namespace.c:sys_mount()` to actually mount the root file system and then `chdir` to it. This is where you see the familiar message shown in Listing 6: VFS: Mounted root (ext2 file system).
- Finally, you return to the `init` function and call `init/main.c:run_init_process`. This results in a call to `execve` to start the `init` process (in this case `/linuxrc`). The `linuxrc` can be an executable or a script (as long as a script interpreter is available for it).

MTD Subsystem Architecture



Graphical Analysis

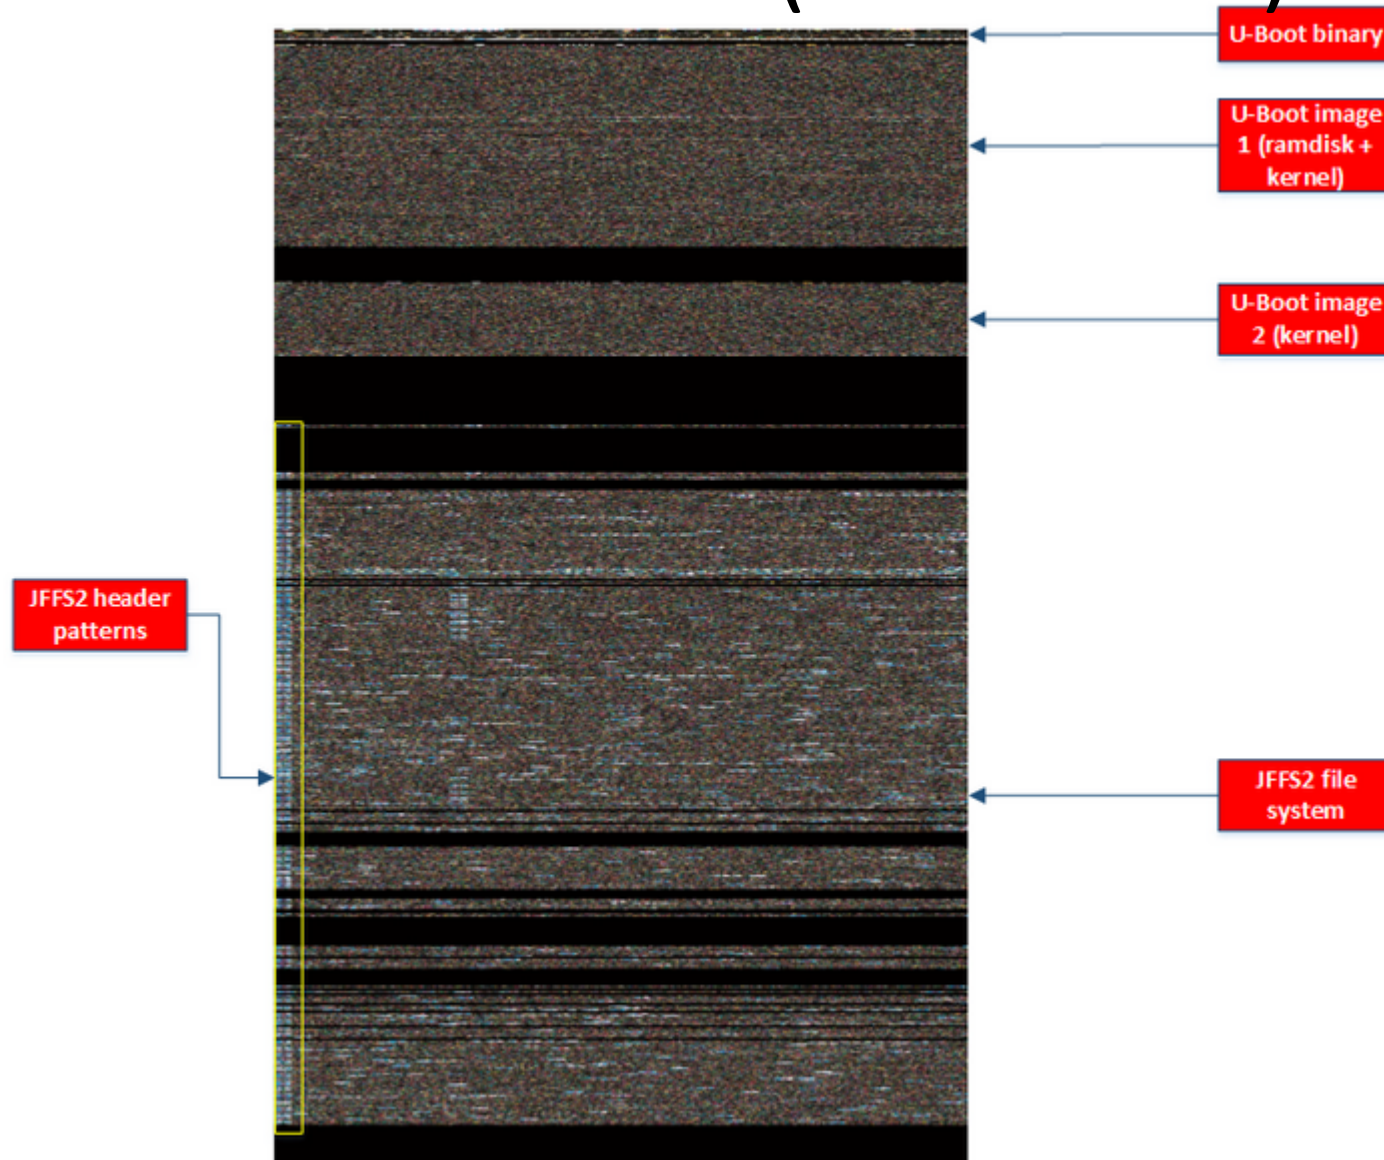
Graphical representation of the bytes in the firmware



Graphical representation of the bytes in the firmware




Graphical representation of the bytes in the firmware (zoomed in)



JFFS2 Header Fields

Offset (h)	00	01	02	03	04	05	06	07	08	09	0A	0B	0C	0D	0E	0F	
004FBF90	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004FBFA0	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004FBFB0	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004FBFC0	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004BFD0	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004BFEO	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004BFFO	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	FF	yyyyyyyyyyyyyyyy
004FC000	85	19	02	E0	67	02	00	00	A3	2F	F4	29	41	07	00	00	...äg...é/ó) Å...
004FC010	03	00	00	00	ED	81	00	00	F4	01	01	00	00	10	00	00	...i...ô.....
004FC020	D3	7D	20	4A	D3	7D	20	4A	D3	7D	20	4A	DD	0D	00	00	ó) JÓ) JÓ) JÝ...
004FC030	23	02	00	00	23	02	00	00	00	00	00	00	69	74	2A	A1	#...#.....it*;
004FC040	D9	0D	A6	F6	9E	9F	0F	B7	D3	9A	CE	25	90	35	52	57	Û;ozÿ. -óšĭ%.5RW
004FC050	B9	D9	90	ED	06	64	AB	C0	5E	B5	35	25	F8	6C	DE	B1	'Û.i.d«À^µ5%elP±
004FC060	F4	70	FE	94	23	3B	12	5D	55	1F	AA	33	91	FC	59	9F	ópp~#;.]U.*3`uYÿ
004FC070	90	55	43	8D	FE	D9	9A	56	CE	4C	5D	2F	75	C9	E0	DB	.UC.pÛsVĪL]/uĒaŮ
004FC080	F8	F3	4F	C4	42	CE	44	28	79	C1	09	E3	7F	43	77	CA	óOÅBĪD(yÁ.Å.CwĒ
004FC090	55	E9	9E	D0	B2	81	74	73	13	F6	F5	AF	8C	1A	57	4E	UéžĐ*.ts.óš`Ē.WN
004FC0A0	6E	46	C6	C3	81	0F	C2	9D	4F	10	01	CC	FF	20	E1	FD	nFĒĀ...Ā.O..ÿy áy



Bootloader Environment Variable Analysis

U-Boot code showing default_environment

```
uchar default_environment[] = {
    #ifdef CONFIG_BOOTARGS
    "bootargs=" CONFIG_BOOTARGS      "\0"
    #endif
    #ifdef CONFIG_BOOTCOMMAND
    "bootcmd=" CONFIG_BOOTCOMMAND    "\0"
    #endif
    #ifdef CONFIG_RAMBOOTCOMMAND
    "ramboot=" CONFIG_RAMBOOTCOMMAND "\0"
    #endif
    #ifdef CONFIG_NFSBOOTCOMMAND
    "nfsboot=" CONFIG_NFSBOOTCOMMAND "\0"
    #endif
    #if defined(CONFIG_BOOTDELAY) && (CONFIG_BOOTDELAY >= 0)
    "bootdelay=" MK_STR(CONFIG_BOOTDELAY) "\0"
    #endif
    #if defined(CONFIG_BAUDRATE) && (CONFIG_BAUDRATE >= 0)
    "baudrate=" MK_STR(CONFIG_BAUDRATE) "\0"
    #endif
    #ifdef CONFIG_LOADS_ECHO
    "loads_echo=" MK_STR(CONFIG_LOADS_ECHO) "\0"
    #endif
    #ifdef CONFIG_ETHADDR
    "ethaddr=" MK_STR(CONFIG_ETHADDR) "\0"
    #endif
}
```

Loading default_environment variables

```

ROM:30F8EE78 sub_30F8EE78 ; CODE XREF: sub_30F8BF24+24↑j
ROM:30F8EE78 ; ROM:30F8C02C↑j ...
ROM:30F8EE78 var_4 = -4
ROM:30F8EE78
ROM:30F8EE78 STR LR, [SP,#var_4]!
ROM:30F8EE78 LDR R3, [R8,#0x18]
ROM:30F8EE7C CMP R3, #0
ROM:30F8EE80
ROM:30F8EE84 LDREQ R3, -aBootargsRootDe ; "bootargs=root=/dev/mtdblock3 rootfstype"...
ROM:30F8EE88 LDREQB R0, [R3,R0]
ROM:30F8EE8C LDREQ PC, [SP+4+var_4],#4
ROM:30F8EE90 BL sub_30F8F0B8
ROM:30F8EE94 AND R0, R0, #0xFF
ROM:30F8EE98 LDR PC, [SP+4+var_4],#4
ROM:30F8EE98 ; End of function sub_30F8EE78

```

Loading bootloader environment variables

bootargs=root=/dev/mtdblock3 rootfstype=jffs2 noinitrd ramdisk_size=4096 mem=32M mtdparts=s3c2410-nand:16k(boot),176k(u-boot),4912k(linux-img), 27104K(rootfs),-(extra);phys_mapped_flash:-(all)

mtdparts string	Name	Blocks
<i>16k(boot)</i>	1st stage bootloader	Block 1 (1 block)
<i>176k(u-boot)</i>	U-Boot code	Block 2 – 12 (11 blocks)
<i>4912k(linux-img)</i>	U-Boot images	Block 13 – 319 (307 blocks)
<i>27104K(rootfs)</i>	JFFS2 file system	Block 320 – 2013 (1694 blocks)

U-boot sub-images

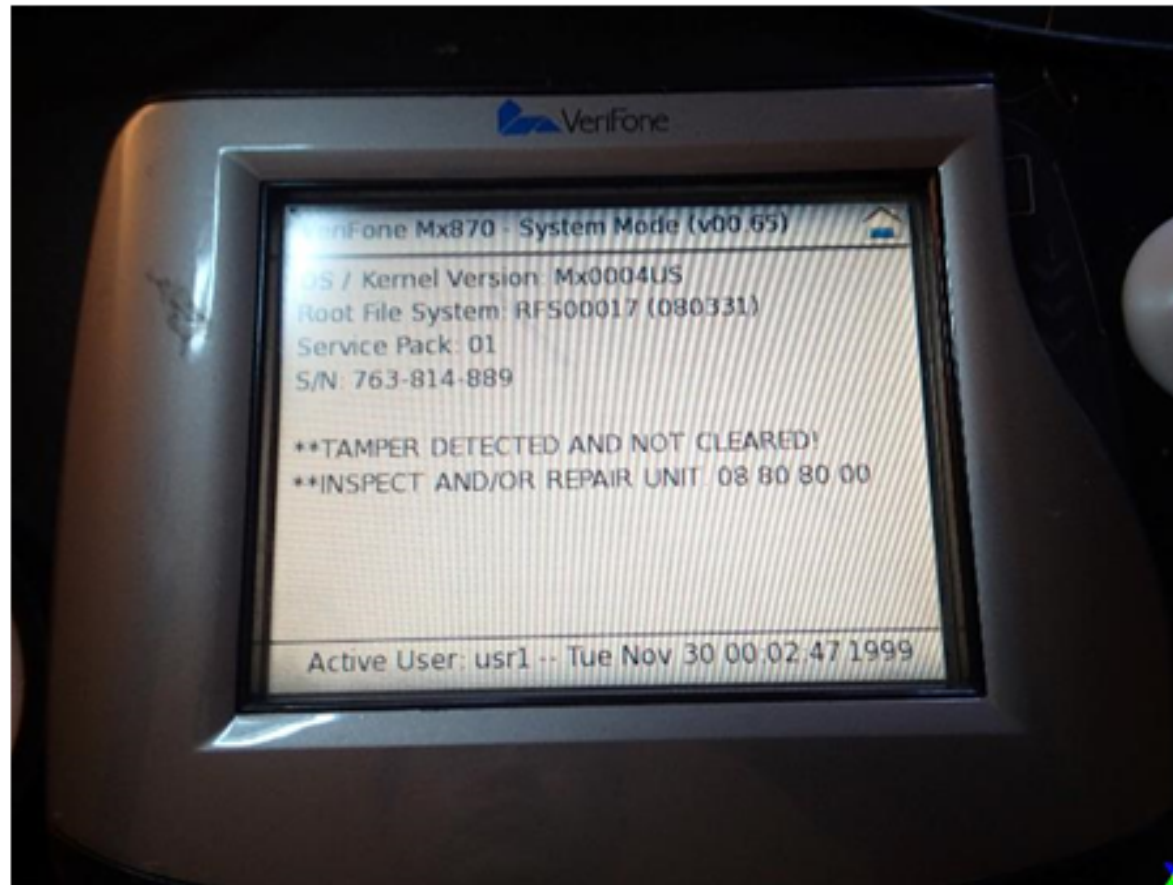
```
C:\mat\Analysis\NAND Flash\DumpFlash>c:\python27\python DumpFlash.py -U
U-Boot Image found at block 0xc
Magic: 0x27051956
HCRC: 0xa05da14d
Time: 0x496669a5
Size: 0x28a03b
Load: 0x30108000
EP: 0x30108000
DCRC: 0x2975d991
OS: 0x5 <Linux>
Arch: 0x2 <ARM>
Type: 0x4 <Multi-File Image>
Comp: 0x0 <None>
Name: Mx0004US 03.00 011f Alt

Found multi image of length 0xe9118
Found multi image of length 0x1a0f17
Extracting to U-Boot-00.dmp-00
Extracting to U-Boot-00.dmp-01
U-Boot Image found at block 0xcc
Magic: 0x27051956
HCRC: 0xbbaceac7
Time: 0x496669a6
Size: 0xe9118
Load: 0x30108000
EP: 0x30108000
DCRC: 0x2855374a
OS: 0x5 <Linux>
Arch: 0x2 <ARM>
Type: 0x2 <OS Kernel Image>
Comp: 0x0 <None>
Name: Mx0004US 03.00 011f


Extracting to U-Boot-01.dmp-00
```

Tamper Detection

Tamper protection in action



Front and Back panels



Circuits inside plastic padding



Before

```
00019354
00019354
00019354
00019354      sub_19354
00019354
00019354      var_1AC= -0x1AC
00019354      var_1A8= -0x1A8
00019354      s= -0x1A4
00019354      var_A4= -0xA4
00019354      var_A3= -0xA3
00019354      var_A2= -0xA2
00019354      var_A1= -0xA1
00019354      var_9C= -0x9C
00019354      set= -0x98
00019354      var_18= -0x18
00019354
00019354 70 40 2D E9 STHFD      SP!, {R4-R6,LR}
00019358 48 01 9F E5 LDR       R0, -aDevSpectrum ; "/dev/spectrum"
0001935C 02 10 A0 E3 MOV       R1, #2 ; oflag
00019360 67 DF 4D E2 SUB       SP, SP, #0x19C
00019364 1C C1 FF EB BL       open
00019368 00 60 50 E2 SUBS      R6, R0, #0
0001936C 3D 00 00 BA BLT      loc_19468
```

```
00019370 6D 1C A0 E3+MOV      R1, #0x6003 ; request
00019378 00 20 A0 E3 MOV      R2, #0
0001937C 29 C0 FF EB BL      ioctl
00019380 00 00 50 E3 CMP      R0, #0
00019384 18 00 00 DA BLE      loc_193EC
```

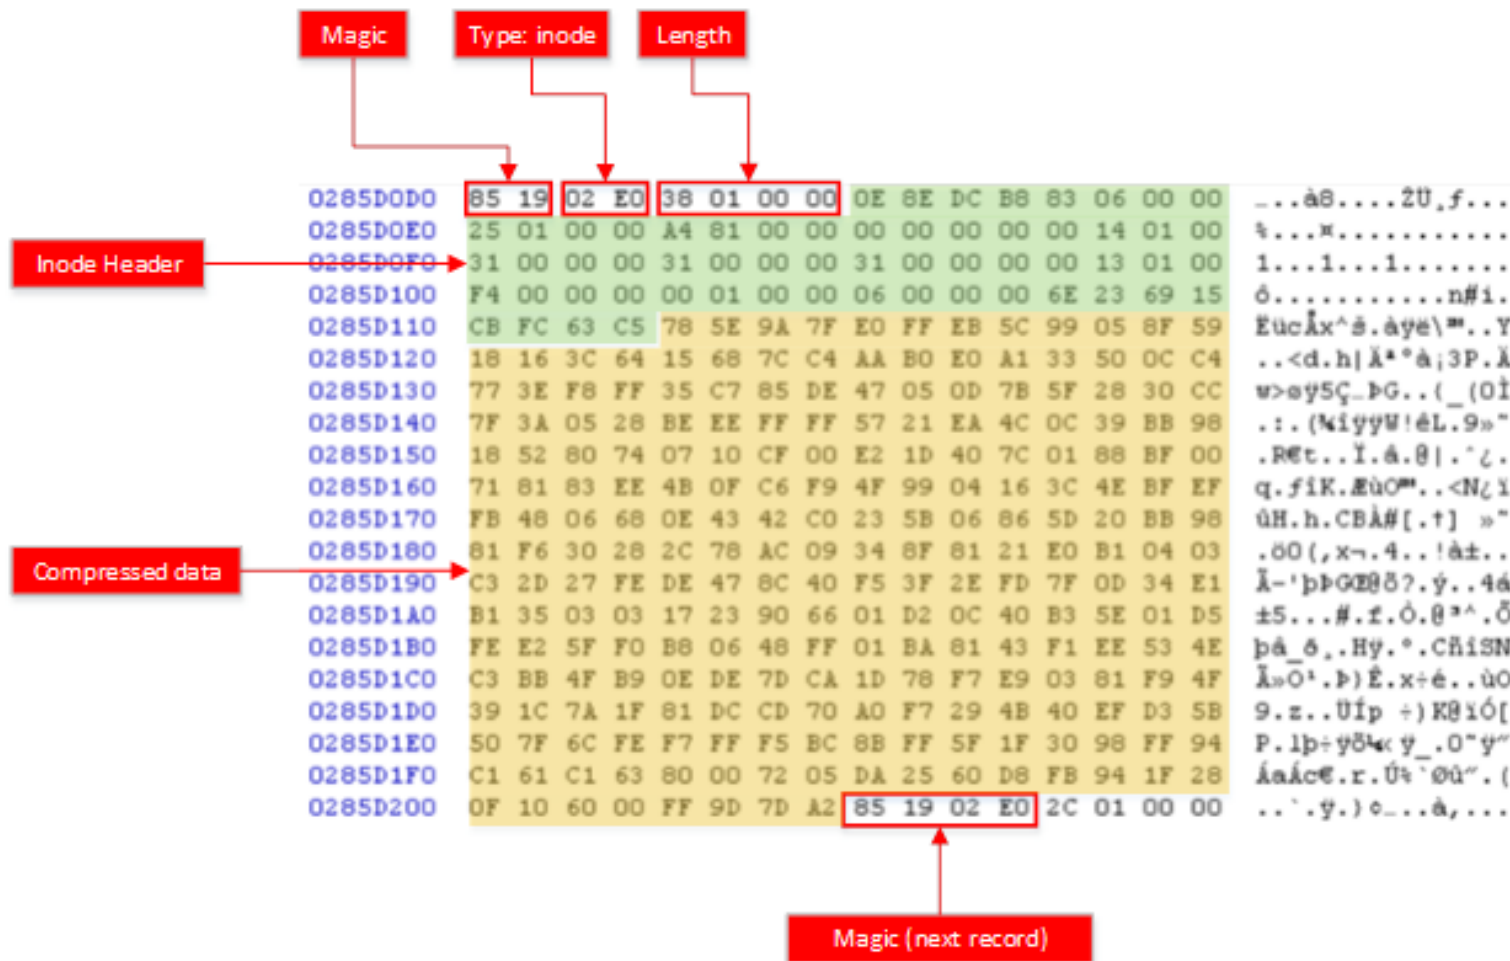
After

Patched

```
00019354
00019354
00019354
00019354      sub_19354
00019354
00019354      var_1AC= -0x1AC
00019354      var_1A8= -0x1A8
00019354      s= -0x1A4
00019354      var_A4= -0xA4
00019354      var_A3= -0xA3
00019354      var_A2= -0xA2
00019354      var_A1= -0xA1
00019354      var_9C= -0x9C
00019354      set= -0x98
00019354      var_18= -0x18
00019354
00019354      70 40 2D E9 STHFD      SP!, {R4-R6,LR}
00019358      48 01 9F E5 LDR      R0, -aDevSpectrum ; "/dev/spectrum"
0001935C      02 10 A0 E3 MOV      R1, #2 ; oflag
00019360      67 DF 4D E2 SUB      SP, SP, #0x19C
00019364      1C C1 FF EB BL      open
00019368      00 60 50 E2 SUBS     R6, R0, #0
0001936C      3D 00 00 BA BLT     loc_19468
```

```
00019370      6D 1C A0 E3+MOV     R1, #0x6D003 ; request
00019378      00 20 A0 E3 MOV      R2, #0
0001937C      29 C0 FF EB BL      ioctl
00019380      01 00 70 E3 CMN      R0, #1
00019384      18 00 00 DA BLE     loc_193EC
```

Original JFFS2 Record



Modified JFFS2 Record

