



The Linux Kernel Internal

Ramin Farajpour Cami

Github : @raminfp

Twitter : @MF4rr3ll



Agenda

- Linux Kernel Booting Process
- How To Make Your Own Kernel OS
- Anatomy of Linux Kernel Development
- System Call
- Kernel Security (Bugs/Vulnerability)
And Kernel Fuzzing

Who Am I?

- Cyber Security Researcher
- Bug Bounty Program
(Google, Twitter, Yahoo, Apple, Ebay, Blackberry, etc)
- Linux Exploit Developer
- Malware Analysis
- Linux System Programmer
- Linux Kernel/Device Developer
- Windows System Programmer
- Django Contributor

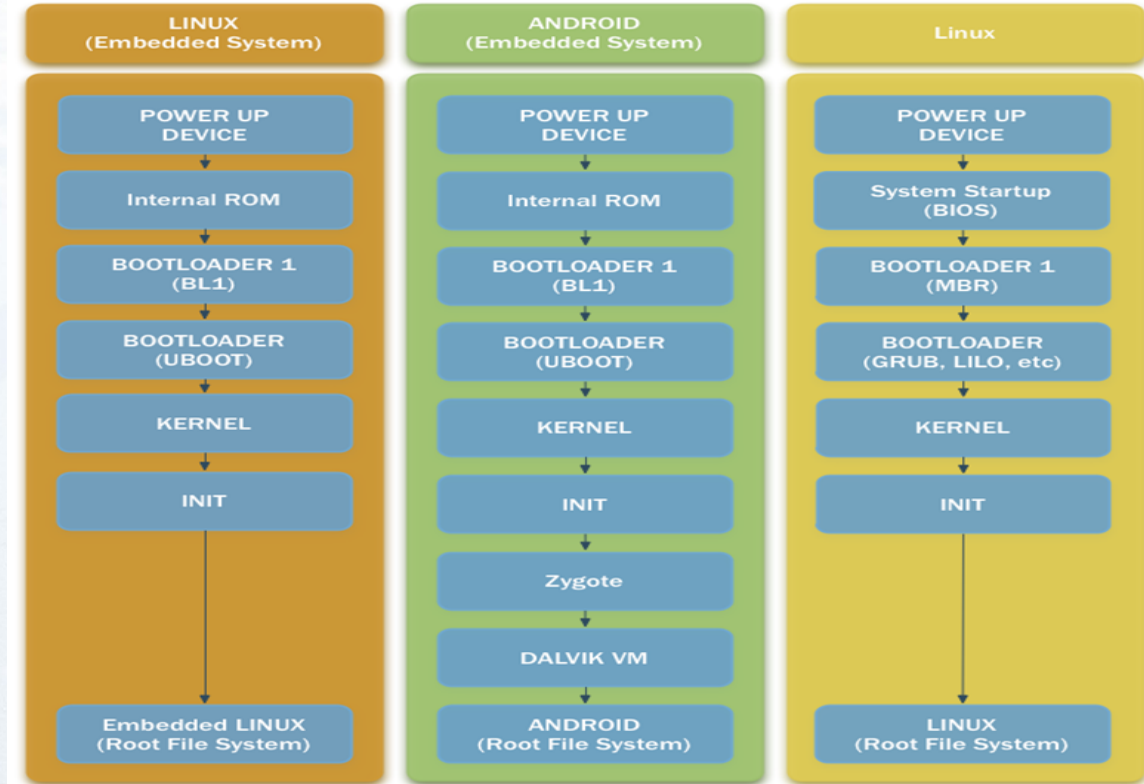




Type of OS

- Personal OS : Linux, Unix, MacOS , etc
- Mobile OS : Android, iOS, WinPhone, UIQ, etc
- Real-Time OS : VxWork (NASA), QNX, RTLinux, etc
- Network OS : Router OS, Switch OS, etc
- Distributed OS : Internet, Telephone networks

Linux Kernel

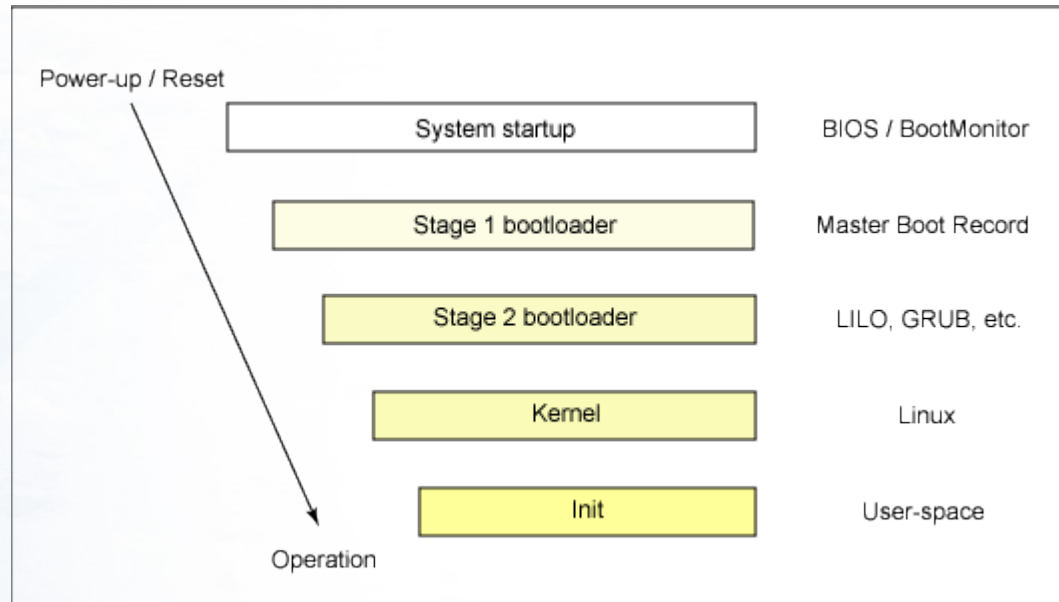




Android Device Linux Kernels Source

- Google: <https://android.googlesource.com/kernel/msm/>
- HTC: <https://www.htcdev.com/devcenter/downloads>
- OnePluseOS : https://github.com/OnePlusOSS/android_kernel_oneplus_msm8996
- Moto X : <https://github.com/MotorolaMobilityLLC/kernel-msm>
- Sony : <https://github.com/sonyxperiadev/kernel>

The Linux Boot Process





The Linux Boot Process

- **BIOS (Basic Input/Output System)**

The processor executes code at a well-known location in a personal computer (PC), which is stored in flash memory on the motherboard.

booting Linux begins in the BIOS at address 0xFFFF0.

When your computer **boots**—and after the **POST** finishes—the BIOS looks for a **Master Boot Record**, or **MBR**, stored on the **boot device** and uses it to launch the **bootloader (GRUB)**.

BIOS Tools

- apt-cache search bios | grep -i bios
- apt-get source phnxdeco (phonix tech)
- **Security advisory (Intel):**
<https://edk2-docs.gitbooks.io/security-advisory/content/>
- **Security Tool:**
Platform Security Assessment Framework
<https://github.com/chipsec/chipsec>
- **Attacking and Defending BIOS in 2015**
- <http://c7zero.info/stuff/AttackingAndDefendingBIOS-RECon2015.pdf>



Bootkit Malware

What is bootkit?

Malware that executes before the operating system boots.

Malwares:

FIN1 : Network protocols and communication channels for command and control (C2C).

Including: file transfer (http DLL web shell), screen capture, keystroke logging, process injection

BOOTRASH :

File Content : Core.sys, vfs.sys and etc

Including : Services, Run keys, Scheduled tasks, Startup folders

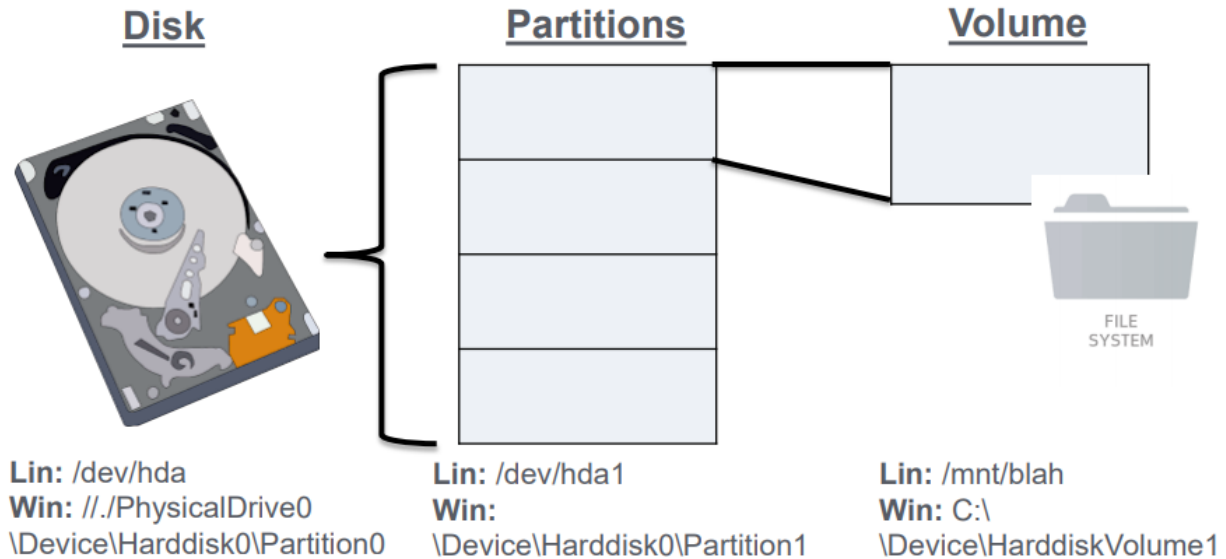
Article :

[https://](https://www.fireeye.com/blog/threat-research/2015/12/fin1-targets-boot-record.html)

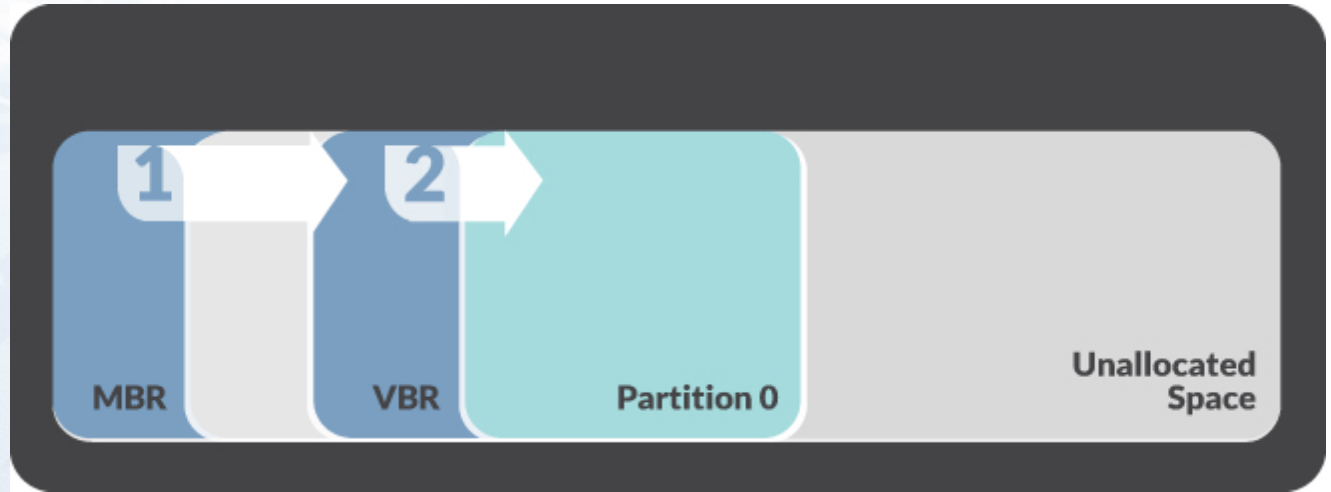
www.fireeye.com/blog/threat-research/2015/12/fin1-targets-boot-record.html

DISKS, PARTITIONS, VOLUMES

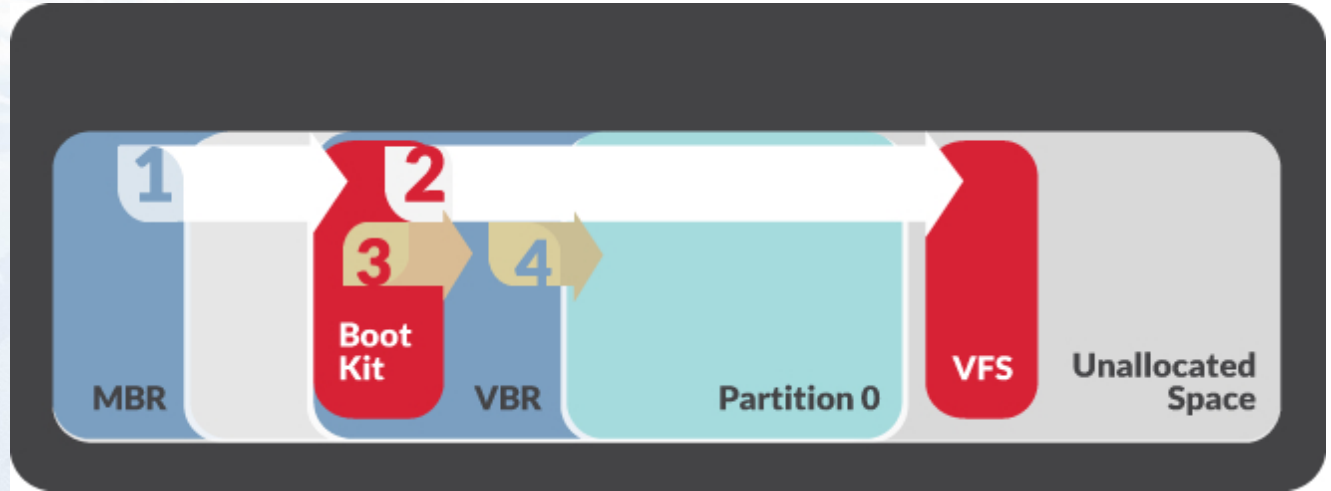
DISKS, PARTITIONS, VOLUMES.



Normal boot process windows



Hijacked boot process





What is UEFI?

- You need to buy new hardware that supports and includes UEFI.
- UEFI firmware can boot from drives of 2.2 TB or larger.
- UEFI can run in 32-bit or 64-bit mode.
- Your boot process is faster.
- UEFI screens can be slicker than BIOS settings screens, including graphics and mouse cursor support.
- UEFI supports Secure Boot, which means the operating system can be checked for validity to ensure no malware has tampered with the boot process.
- UEFI support networking features.
- UEFI is modular.

- Analyzing UEFI BIOSes from attacking [BH 2014] (<https://youtu.be/CGBpil0S5NI>)

UEFI

ASUS EFI BIOS Utility - EZ Mode Exit/Advanced Mode

06:06 P8P67 DELUXE English
BIOS Version : 0304 Build Date : 10/21/2010
CPU Type : Intel(R) Core(TM) i5-2400 CPU @ 3.10GHz Speed : 3100 MHz
Tuesday[11/16/2010] Total Memory : 2048 MB (DDR3 1333MHz)

Temperature
CPU +116.6°F/+47.0°C
MB +96.8°F/+36.0°C

Voltage
CPU 1.200V 5V 5.120V
3.3V 3.408V 12V 12.288V

Fan Speed
CPU_FAN 1834RPM PWR_FAN1 N/A
CHA_FAN1 N/A CHA_FAN2 N/A

System Performance
Quiet Performance Energy Saving Normal

Boot Priority
UEFI UEFI Hard Drive Floppy

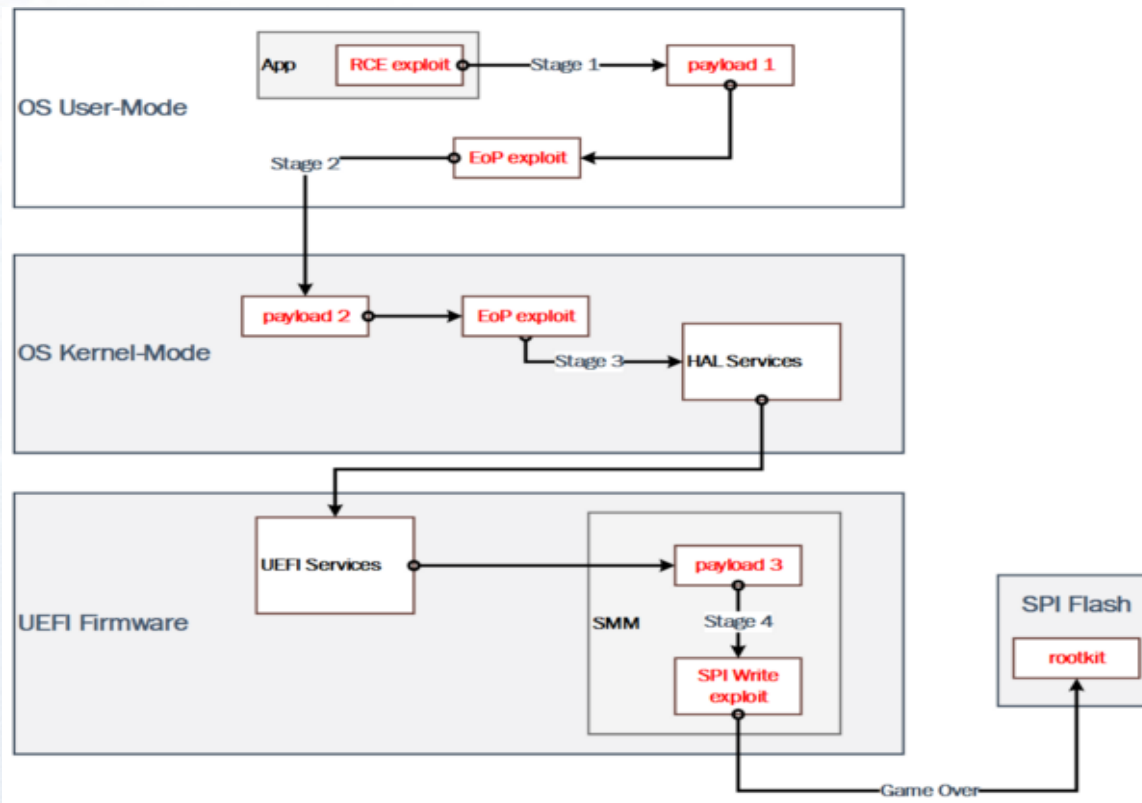
Use the mouse to drag or keyboard to navigate to decide the boot priority.

Boot Menu(F8) Default(F5)


UEFI Rootkit Malware (LoJax)

- How malware works:
 - Information dumping tool:
 - *RwDrv.sys* driver, attackers reads the information on your UEFI BIOS. This information is then saved to a text file. This step helps the malware understand the victim system.
 - System Firmware Image Creation :
 - **Firmware** Image of SPI Flash Memory where the UEFI/BIOS is located. This image is then again saved to a file.
 - Rootkit Installation:
 - The firmware image is infected. This infected Firmware image is then installed onto the **SPI Flash Memory**

OS drivers dangerous for BIOS



Create malware is important for country





ICS (Industrial Control Systems) malware

- **2010 Stuxnet** : This cyber weapon was created to target Iranian centrifuges.
- **2013 Havex** : Targeted energy grids, electricity firms, and many others.
- **2015 BlackEnergy** : It targeted critical infrastructure and destroyed files stored on workstations and servers. In Ukraine.
- **2015 IronGate** : It targeted Siemens control systems and had functionalities similar to Stuxnet's.
- **2016 Industroyer** : The attack caused a second shutdown of Ukraine's power grid.
- **2017 Triton** : The attack did not succeed



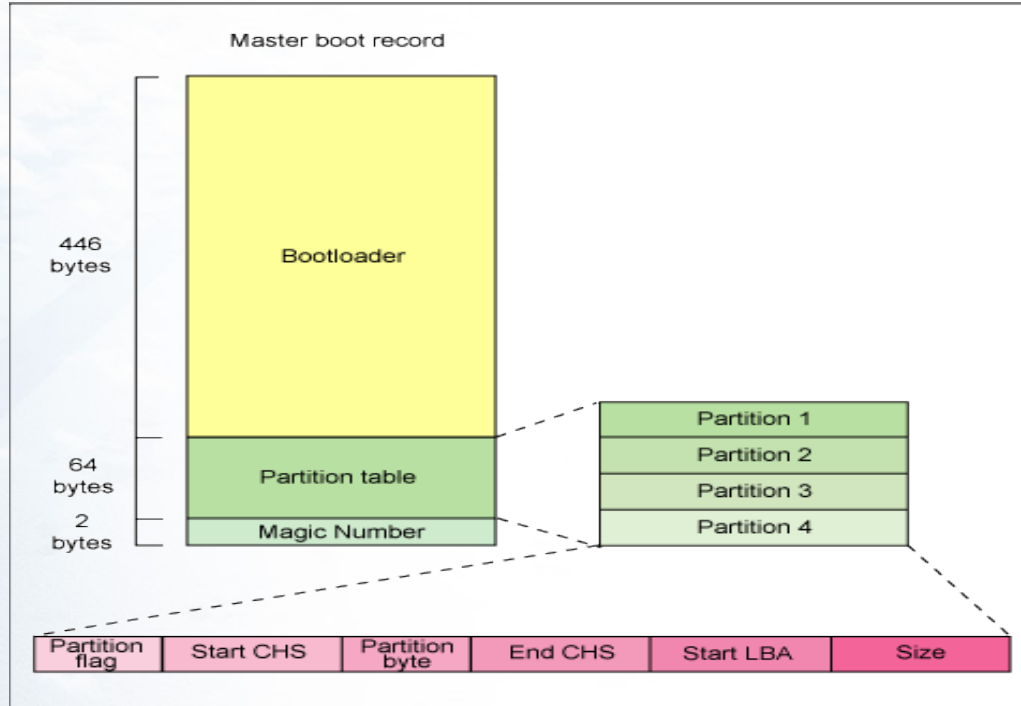
What Is BootLoader?

- It is the part that starts the system up and loads the operating system kernel
- **Bootloader has two main jobs:**
- **[1]** Initialize the system to a basic level (MBR) and to **[2]** Load the kernel.

Type Bootloader

Name	Main architectures supported
Das U-Boot	ARC, ARM, Blackfin, Microblaze, MIPS, Nios2, OpenRisc, PowerPC, SH
Barebox	ARM, Blackfin, MIPS, Nios2, PowerPC
GRUB 2	X86, X86_64
Little Kernel	ARM
RedBoot	ARM, MIPS, PowerPC, SH
CFE	Broadcom MIPS
YAMON	MIPS

First stage bootloader = MBR





MBR (Master Boot Record)

- **MBR :**

The first 446 bytes are the primary boot loader, which contains both executable code and error message text

The next sixty-four bytes are the partition table, which contains a record for each of four partitions The MBR ends with two bytes that are defined as the magic number (**0xAA55**)



Partition table information of MBR

- # file mbr.bin

mbr.bin: x86 boot sector; **partition 1**: ID=0x83, **active**,
starthead 32, **startsector** 2048, **19451904 sectors**;
partition 2: ID=0x5, starthead 254, **startsector** 19455998,
2093058 sectors, code offset 0x63

Second stage bootloader

Splash screen is commonly displayed, and Linux and an optional initial RAM disk (temporary root file system) are loaded into memory.

second-stage, boot loader called the kernel loader. The task at this stage is to load the Linux kernel and optional initial RAM disk.



```
GNU GRUB version 2.00

setparams 'openSUSE 12.3'
load_video
set gfxpayload=keep
insmod gzio
insmod part_msdos

Minimum Emacs-like screen editing is
supported. TAB lists completions.
Press Ctrl-x or F10 to boot, Ctrl-c or
F2 for a command-line or ESC to discard
edits and return to the GRUB menu.
```



Second stage = GRUB and etc

Type of bootloader :

1- Grub

2- LILO

3- GRand

4- ...

GRUB

Good knowledge of Linux file system. Instead of using raw sectors on the disk, as LILO.

GRUB can load a Linux kernel from an ext2 or ext3 file system

```
GNU GRUB  version 0.97  (640K lower / 3072K upper memory)

[ Minimal BASH-like line editing is supported.  For
  the first word, TAB lists possible command
  completions.  Anywhere else TAB lists the possible
  completions of a device/filename. ]

grub> root (hd0,0)
Filesystem type is ext2fs, partition type 0x83

grub> setup (hd0)
Checking if "/boot/grub/stage1" exists... yes
Checking if "/boot/grub/stage2" exists... yes
Checking if "/boot/grub/e2fs_stage1_5" exists... yes
Running "embed /boot/grub/e2fs_stage1_5 (hd0)"...  16 sectors are embedded.
succeeded
Running "install /boot/grub/stage1 (hd0) (hd0)1+16 p (hd0,0)/boot/grub/stage2 /bo
ot/grub/menu.lst"...  succeeded
Done.

grub> quit
```



Stage 1 (MBR) + Stage 1.5 + Stage 2 (GRUB)

What is stage 1.5?

stage 1.5 boot loader that understands the particular file system containing the Linux kernel image.

Examples :

CR-ROMs use the **iso9660_stage_1_5**

Ext2 or ext3 file system use the **e2fs_stage1_5**

GRUB *.cfg (Stage 1.5 □ Stage 2 loaded)

```
1127 else
1128     set linux_gfx_mode=keep
1129 fi
1130 else
1131     set linux_gfx_mode=text
1132 fi
1133 export linux_gfx_mode
1134 menuentry 'Ubuntu' --class ubuntu --class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-simple-e339680f-111b-48c3-a09f-e68023009a24' {
1135     recordrat
1136     load_video
1137     gfxmode $linux_gfx_mode
1138     insmod gzio
1139     if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
1140     insmod part_gpt
1141     insmod ext2
1142     set root='hd0,gpt2'
1143     if [ x$feature_platform_search_hint = xy ]; then
1144         search --no-floppy --fs-uuid --set=root --hint-bios=hd0,gpt2 --hint-efi=hd0,gpt2 --hint-baremetal=ahci0,gpt2  e339680f-111b-48c3-a09f-e68023009a24
1145     else
1146         search --no-floppy --fs-uuid --set=root e339680f-111b-48c3-a09f-e68023009a24
1147     fi
1148     linux /boot/vmlinuz-4.15.0-041500rc8-generic root=UUID=e339680f-111b-48c3-a09f-e68023009a24 ro quiet splash $vt_handoff
1149     initrd /boot/initrd.img-4.15.0-041500rc8-generic
1150 }
1151 submenu 'Advanced options for Ubuntu' $menuentry_id_option 'gnulinux-advanced-e339680f-111b-48c3-a09f-e68023009a24' {
1152     menuentry 'Ubuntu, with Linux 4.15.0-041500rc8-generic' --class ubuntu --class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-4.15.0-041500rc8-generi
1153         recordrat
1154         load_video
1155         gfxmode $linux_gfx_mode
1156         insmod gzio
1157         if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
1158         insmod part_gpt
1159         insmod ext2
1160         set root='hd0,gpt2'
1161         if [ x$feature_platform_search_hint = xy ]; then
1162             search --no-floppy --fs-uuid --set=root --hint-bios=hd0,gpt2 --hint-efi=hd0,gpt2 --hint-baremetal=ahci0,gpt2  e339680f-111b-48c3-a09f-e68023009a24
1163         else
1164             search --no-floppy --fs-uuid --set=root e339680f-111b-48c3-a09f-e68023009a24
1165         fi
1166         echo 'Loading Linux 4.15.0-041500rc8-generic ...'
1167         linux /boot/vmlinuz-4.15.0-041500rc8-generic root=UUID=e339680f-111b-48c3-a09f-e68023009a24 ro quiet splash $vt_handoff
1168         echo 'Loading initial ramdisk ...'
1169         initrd /boot/initrd.img-4.15.0-041500rc8-generic
1170     }
1171     menuentry 'Ubuntu, with Linux 4.15.0-041500rc8-generic (recovery mode)' --class ubuntu --class gnu-linux --class gnu --class os $menuentry_id_option 'gnulinux-4.15.
1172         recordrat
1173         load_video
1174         insmod xzio
1175         if [ x$grub_platform = xxen ]; then insmod xzio; insmod lzopio; fi
1176         insmod part_gpt
1177         insmod ext2
1178         set root='hd0,gpt2'
```



GRUB GUIDE

Refs:

<https://thestarman.pcministry.com/asm/mbr/GRUB.htm>

<http://people.ds.cam.ac.uk/fanf2/hermes/src/grub-e1000/>

<https://>

www.gnu.org/software/grub/manual/grub/grub.html#General-bootstrap-methods

<https://>

github.com/coreos/grub/blob/93fb3dac4ae7a97c080d51d951d0e5a3109aaac7/grub-core/kern/main.c

Understanding the Various Grub Modules

```
$ ls /boot/grub/x86_64-efi/
```

```
/boot/grub/x86_64-efi/915resolution.mod
```

```
...
```

Grub module :

<https://github.com/coreos/grub/tree/2.02-coreos/grub-core>

Grub module error: file `/boot/grub/*/*.mod` not found.

```
# GRUB loading.
# Welcome to GRUB!
#
# error: file `/boot/grub/i386-pc/normal.mod` not found.
# Entering rescue mode...
grub rescue> ls
hd(0) (hd0,msdos1)
grub rescue> set
cmdpath=(hd0)
prefix=(hd0,msdos1)/boot/grub
root=hd0,msdos1
grub rescue> set prefix=(hd0,msdos1)/usr/lib/grub
grub rescue> insmod normal
grub rescue> normal
grub> ls (hd0,msdos1) # Display UUID
grub> linux /boot/vmlinuz-linux root=UUID=xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx # Enter UUID
grub> initrd /boot/initramfs-linux.img
grub> boot
> pacman -Syuu
> grub-mkconfig -o /boot/grub/grub.cfg
> grub-install --recheck /dev/sda
> reboot
```



Load kernel image with GRUB

```
grub> kernel /bzImage-<version>
```

```
[Linux-bzImage, setup=0x1400, size=0x29672e]
```

```
grub> initrd /initrd-<version>.img
```

```
[Linux-initrd @ 0x5f13000, 0xcc199 bytes]
```

```
grub> boot
```

```
Uncompressing Linux... Ok, booting the kernel.
```

Refs install manual : http://tinycorelinux.net/install_manual.html



Count of Line Code Ubuntu Kernel

```
~/Desktop/kernel/linux-4.13.0$ cloc .
62037 text files.
61240 unique files.
Unescaped left brace in regex is deprecated here (and will be fatal in Perl 5.30), passed through in regex;
11993 files ignored.

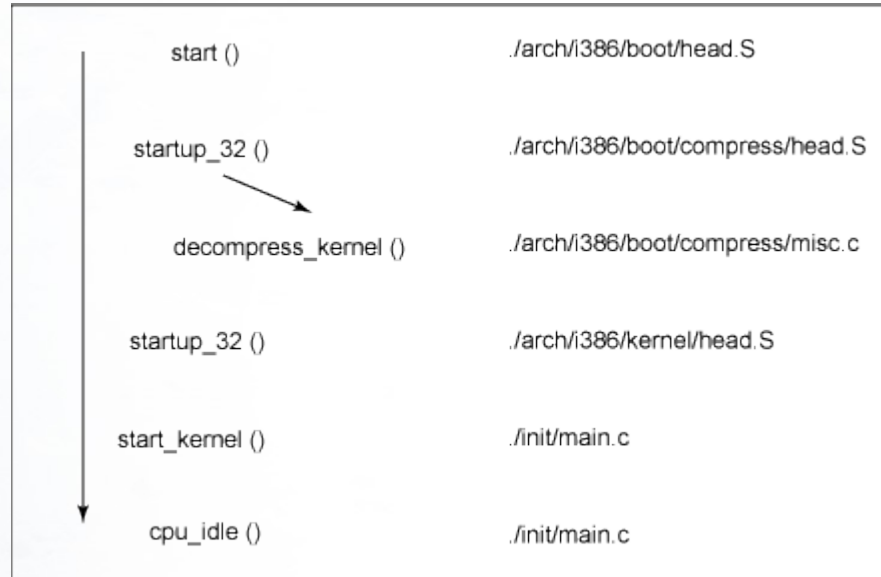
github.com/AlDanial/cloc v 1.70 T=369.78 s (135.4 files/s, 63323.0 lines/s)
-----
Language           files          blank          comment          code
-----
C                   25558          2502431         2271988          12698359
C/C++ Header       19873          495917          953724           3732465
Assembly           1434           49270           114785           248172
JSON                138            0                0                88336
Bourne Shell       268            20145           11000            69947
make                2424           8995            8517            39497
Perl                66             5481            4104            29010
m4                  105            1572            1167            15862
Python              80             2466            2852            13827
HTML                3              565             0                4730
yacc                9              682             357            4530
Bourne Again Shell 61             616             513            2940
lex                 8              302             300            1907
C++                 7              287             71            1838
awk                 13             186             179            1515
Markdown            1              220             0                1077
TeX                 1              108             3                904
NAnt script         2              158             0                602
Pascal              3              49              0                231
Objective C++       1              55              0                189
XSLT                5              13              26                61
CSS                 1              14              23                35
YAML                1              1                0                30
vim script          1              3                12                27
Windows Module Definition 1              0                0                8
sed                 1              2                5                5
-----
SUM:                50065          3089538         3369626          16956104
-----
```

Count of Line Code Main Kernel

```
~/Desktop/kernel$ cloc linux-4.19.tar.xz
61693 text files.
61264 unique files.
12228 files ignored.

github.com/AlDanial/cloc v 1.70 T=401.07 s (123.4 files/s, 59120.6 lines/s)
-----
Language          files          blank          comment          code
-----
C                  26082          2586769          2268941          13131964
C/C++ Header      18856          491495           902716           3655384
Assembly          1323           47331            106333           232938
JSON              194            0                0                105175
make              2389           8763             9449             38120
Bourne Shell      377            6590             5904             28255
Perl              55             5426             4004             27407
Python            108            3093             3341             17645
HTML              5              670              0                5497
yacc              9              701              375             4648
lex               8              326              314             2007
C++               7              286              77              1844
Bourne Again Shell 51             352              318             1722
awk               11             170              155             1386
Markdown          1              220              0                1077
TeX               1              108              3                915
NAnt script       2              155              0                588
Windows Module Definition 2              14               0                102
m4                1              15               1                95
XSLT              5              13               26              61
CSS               1              18               27              44
vim script        1              3                12              27
Ruby              1              4                0                25
INI               1              1                0                6
sed               1              2                5                5
-----
SUM:              49492          3152525          3302001          17256937
-----
```

Kernel Boot Road Map





Understand of Kernel Image

vmlinux: Plain linux ELF file just the way it was created by the linker, including symbols and everything.

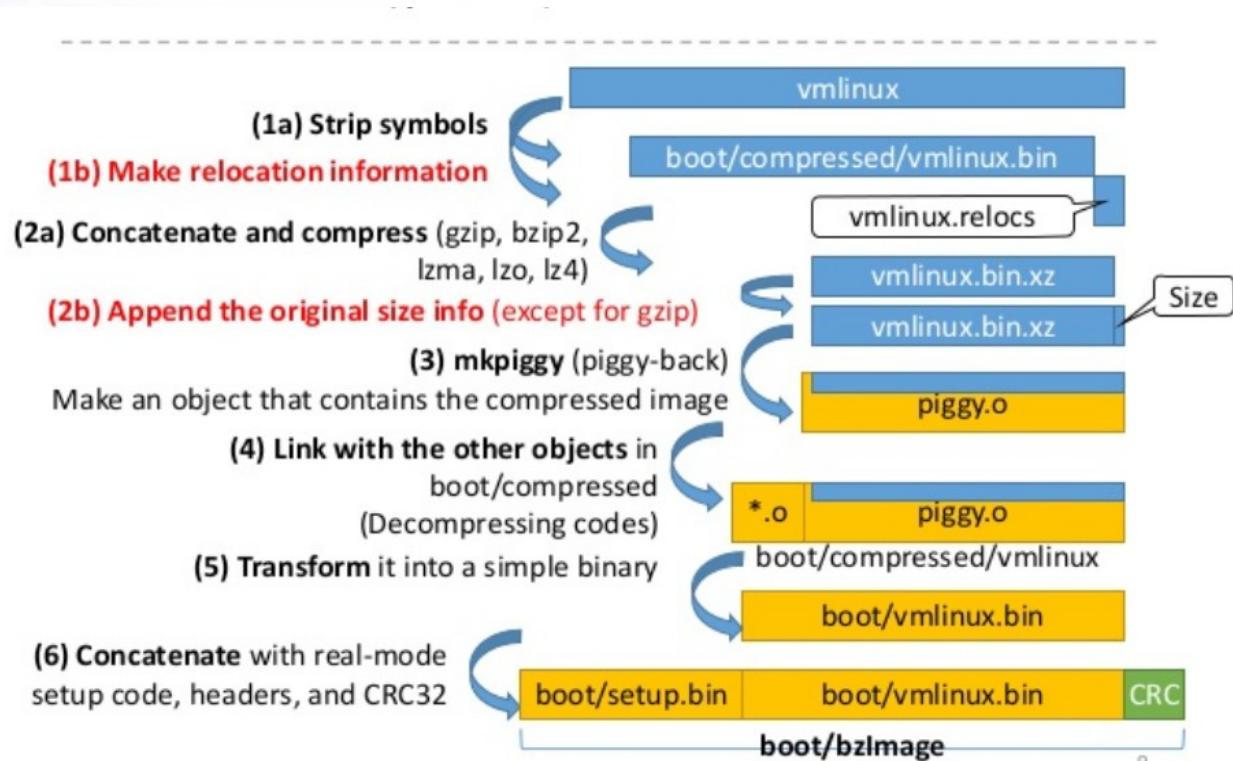
vmlinuz: Gzipped vmlinux file which got stripped of all its symbols

zImage: `bootsect.o` + `setup.o` + `misc.o` + `piggy.o` (`piggy.o` contains the piggy-backed **vmlinuz**).

zImage : is bootable because it can decompress and run the kernel it contains.

bzImage: Same as zImage except that it is built slightly differently which enables it to carry bigger kernels.

Vmlinux to Vmlinuz (make bzImage)





Vmlinux

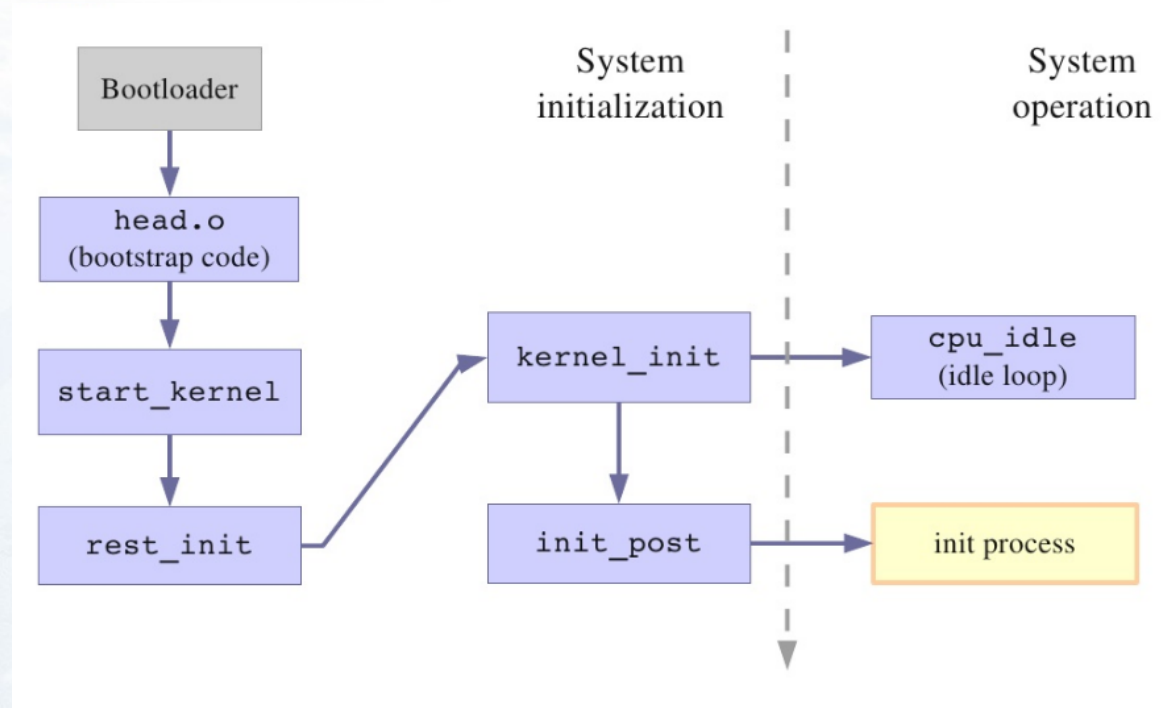
- Vmlinux is a ELF format,
- How to get ELF format?
- Download kernel source of <https://kernel.com> and compiled, or use \$ `apt-get source linux`



Vmlinuz

- `$ sudo file /boot/vmlinuz-4.15.0-041500rc8-generic`
- `/boot/vmlinuz-4.15.0-041500rc8-generic: Linux kernel x86 boot executable bzImage, version 4.15.0-041500rc8-generic (kernel@gloin) #201801142030 SMP Mon Jan 15 01:31:43 UTC 2018, RO-rootFS, swap_dev 0x7, Normal VGA`

Kernel Start Up





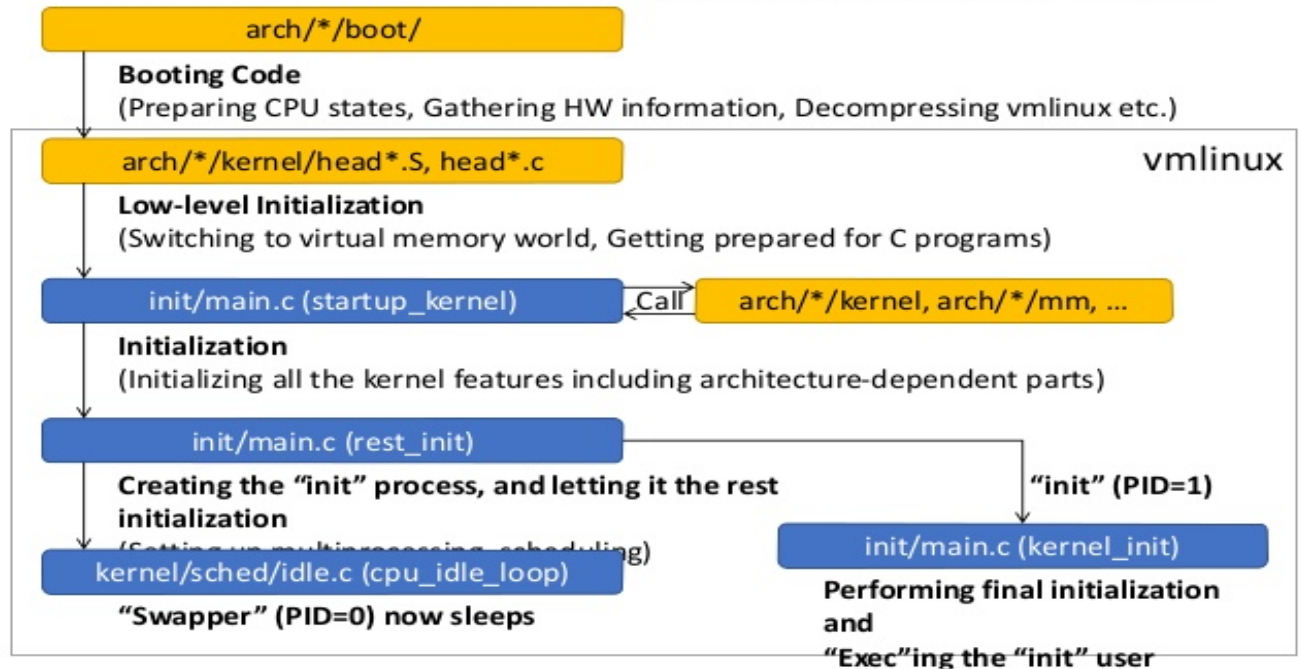
Kernel Image Process

- Kernel image (bzImage) load to Memory and kernel stage started ...
- Typically zImage compressed image, less than 512KB a bzImage (big compressed image, greater than 512KB)
- When the bzImage (for an i386 image) is invoked, you begin at `./arch/i386/boot/head.S` in the `start` assembly routine
- The kernel is then decompressed (`./arch/i386/boot/compressed/misc.c`) through a call to a C function called `decompress_kernel` function
- When the kernel is decompressed into memory, it is called. This is yet another `startup_32` function, but this function is in `./arch/i386/kernel/head.S`.
- More info : <https://www.slideshare.net/itembedded/linux-kernel-image>

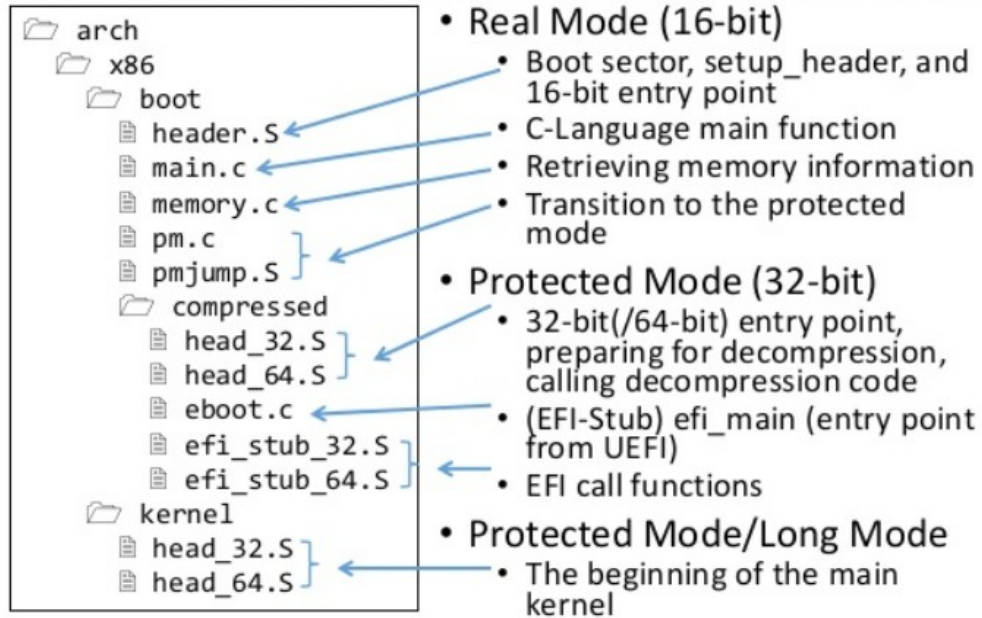
Kernel Overview

4

Initialization Overview

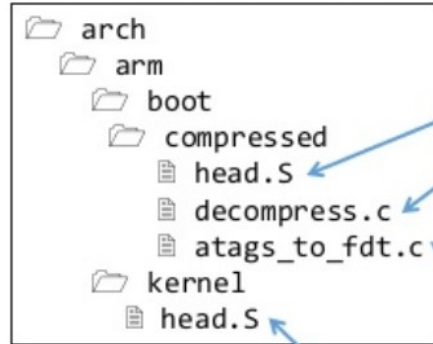


Kernel Source x86



* The ..._32.S files are used in the 32-bit kernel, and ..._64.S files are not. Vice versa.

Kernel Source ARM



- Compressed
 - Entry point
 - Decompressing function
 - Actual decompressing algorithm is in lib/decompress_*.c
 - Building a FDT from ATAGS for compatibility (CONFIG_ARM_ATAG_DTB_COMPAT)
- Decompressed
 - The beginning of the main kernel



Kernel Source Tree

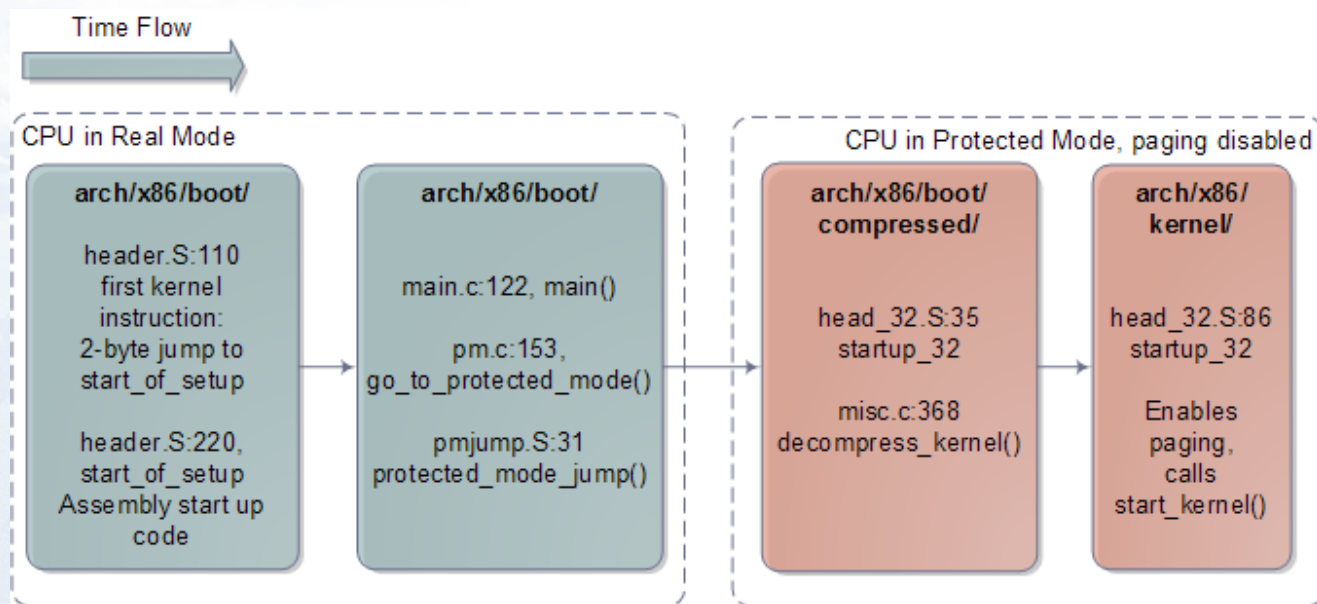
- Github linux kernel source tree :
 - <https://github.com/torvalds/linux>
- Bootlin kernel source tree :
 - <https://elixir.bootlin.com/linux/latest/source>



Kernel Type Mode

- Real Mode
- Protected Mode
- Long Mode

Architecture Linux Kernel Initialization





Kernel Boot Paging (Virtual Memory)

Before the kernel starts...

- x86 (32-bit)
 - Paging is disabled
 - kernel/head_32.S creates a page table and turns on paging
- x86 (64-bit)
 - compressed/head_64.S creates an identical (virtual = physical) page table for the first 4G
 - Long mode requires paging enabled.
 - kernel/head_64.S creates better page table
- ARM
 - kernel/head.S creates a page table and turns on paging

Real Mode Kernel

- header.S
 - **Boot sector code which is no longer used**
 - Contains setup_header
 - Prepares stack and BSS to run C programs
 - Jumps into the C program (main.c)
- main.c
 - Copies setup_header into *“zeropage”*
 - Setups early console
 - Initializes heap
 - Checks the CPUs (64-bit capable for 64-bit kernel?)
 - Collect HW information by querying to BIOS, and stores the results in *“zeropage”*
 - Finally transits to protected-mode, and jumps into the *“protected-mode kernel”*

Real Mode Boot sector (header.S)

Boot sector (Useless)

```
.global bootsect_start
bootsect_start:
#ifdef CONFIG_EFI_STUB
# "MZ", MS-DOS header
.byte 0x4d
.byte 0x5a
#endif

# Normalize the start address
ljmp     $BOOTSEG, $start2

start2:
movw    %cs, %ax
movw    %ax, %ds
movw    %ax, %es
movw    %ax, %ss
xorw    %sp, %sp
sti
cld

movw    $bugger_off_msg, %si

jmp     msg_loop
```

Normalize CS to
BOOTSEG (0x7c0).

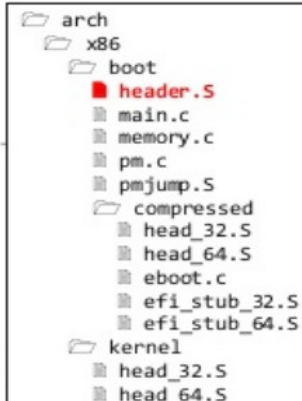
movw %ds, %cs is not allowed.

stack starts at 0x17c00

Enable interrupts *cf.* cli

Reset directions for string instructions
(Clear DF Flag) *cf.* std

Show the message "Direct floppy boot
is not supported."





Kernel Real Mode (setup_header)

- header.S

- Boot sector code which is no longer used
- Contains setup_header
- Prepares stack and BSS to run C programs
- Jumps into the C program (main.c)

- main.c

- Copies setup_header into “zeropage”
- Setups early console
- Initializes heap
- Checks the CPUs (64-bit capable for 64-bit kernel?)
- Collect HW information by querying to BIOS, and stores the results in “zeropage”
- Finally transits to protected-mode, and jumps into the “protected-mode kernel”


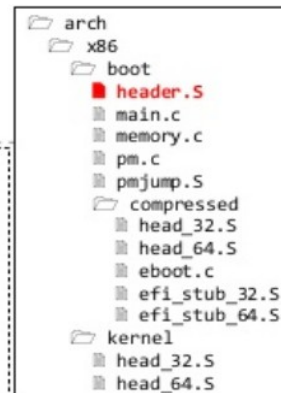
Kernel Real Mode (setup_header)

```
.section ".header", "a"
.globl sentinel
sentinel:.byte 0xff, 0xff /* Used to detect broken loaders */

.globl hdr
hdr:
setup_sects: .byte 0 /* Filled in by build.c */
root_flags: .word ROOT_RDONLY
syssize: .long 0 /* Filled in by build.c */
ram_size: .word 0 /* Obsolete */
vid_mode: .word SVGA_MODE
root_dev: .word 0 /* Filled in by build.c */
boot_flag: .word 0xAA55

# offset 512, entry point
.globl _start
_start:
.byte 0xeb # short (2-byte) jump
.byte start_of_setup-1f

1:
.ascii "HdrS" # header signature
.word 0x020d # header version number (>= 0x0105)
```

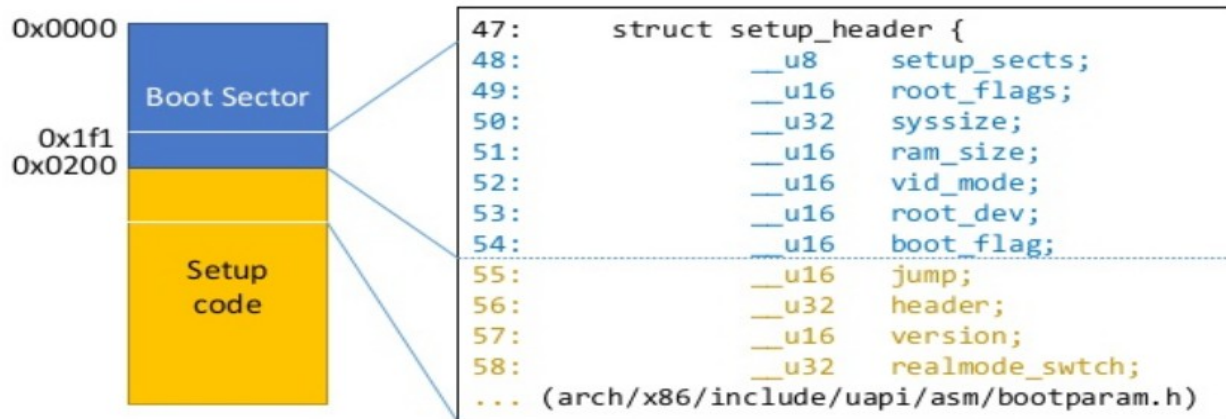


To prevent the compiler from accidentally producing a 3-byte jump

Struct setup_header (header.S)

Setup_header

- “.header” section starts at 495
 - 2-byte sentinel is located at the beginning.
 - Struct setup_header begins at 497 (=0x1f1)



Kernel Real Mode Stack (header.S)

Get prepared to C (stack)

```
.section ".entrytext", "ax"
start_of_setup:
# Force %es = %ds
movw    %ds, %ax
movw    %ax, %es
cld
movw    %ss, %dx
cmpw    %ax, %dx # %ds == %ss?
movw    %sp, %dx
je      2f      # -> assume %sp is reasonably set

# Invalid %ss, make up a new stack
movw    $_end, %dx
testb   $CAN_USE_HEAP, loadflags
jz      1f
movw    heap_end_ptr, %dx
1:      addw   $STACK_SIZE, %dx
jnc     2f
xorw    %dx, %dx # Prevent wraparound
2:      # Now %dx should point to the end of our stack
andw    $~3, %dx # dword align (might as well)
jnz     3f
movw    $0xffffc, %dx # Make sure we're not zero
3:      movw   %ax, %ss
movzwl  %dx, %esp # Clear upper half of %esp
```

```
arch
├── x86
│   └── boot
│       ├── header.S
│       ├── main.c
│       ├── memory.c
│       ├── pm.c
│       ├── pmjump.S
│       └── compressed
│           ├── head_32.S
│           ├── head_64.S
│           ├── eboot.c
│           ├── efi_stub_32.S
│           └── efi_stub_64.S
└── kernel
    ├── head_32.S
    └── head_64.S
```

If %ds == %ss, %sp is assumed to be properly set by the loader

If not, sets up a new stack. The address is `_end + STACK_SIZE` (512 byte) or `heap_end_ptr + STACK_SIZE` (if `CAN_USE_HEAP` is set)

Kernel Real Mode to C (header.S -> main.c)

Get prepared to C

```
        sti        # Now we should have a working stack

# We will have entered with %cs = %ds+0x20, normalize %cs so
# it is on par with the other segments.
        pushw    %ds
        pushw    $6f
        lretw

6:

# Check signature at end of setup
        cmpl    $0x5a5aaa55, setup_sig
        jne     setup_bad

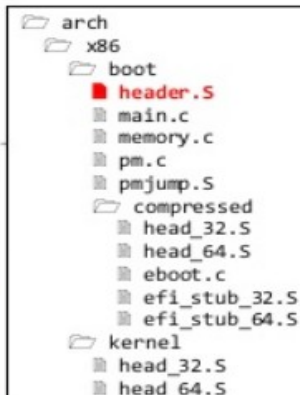
# Zero the bss
        movw    $__bss_start, %di
        movw    $_end+3, %cx
        xorl    %eax, %eax
        subw   %di, %cx
        shrw   $2, %cx
        rep; stosl

# Jump to C code (should not return)
        calll   main
```

\$6f is the address of the 6f, which is the offset from the boot sector.

Signature check

Fill the bss by zero. "rep; stosl" (string instruction) fills the memory from %es:%di for %cx DWORDs with %eax.



Kernel Real Mode Main.c

```
void main(void)
{
    /* First, copy the boot header into the "zeropage" */
    copy_boot_params();

    /* Initialize the early-boot console */
    console_init();
    if (cmdline_find_option_bool("debug"))
        puts("early console in setup code\n");

    /* End of heap check */
    init_heap();

    /* Make sure we have all the proper CPU support */
    if (validate_cpu()) {
        puts("Unable to boot - please use a kernel appropriate "
            "for your CPU.\n");
        die();
    }

    /* Tell the BIOS what CPU mode we intend to run in. */
    set_bios_mode();

    /* Detect memory layout */
    detect_memory();

    /* Set keyboard repeat rate (why?) and query the lock flags */
    keyboard_init();

    /* Query Intel SpeedStep (IST) information */
    query_ist();

    /* Query APM information */
    #if defined(CONFIG_APM) || defined(CONFIG_APM_MODULE)
        query_apm_bios();
    #endif
}
```

Copy header to zeropage

Secure | <https://elixir.bootlin.com/linux/v4.19-rc2/source/arch/x86/boot/main.c#L30>

```
 / arch / x86 / boot / main.c
24  /*
25  * Copy the header into the boot parameter block. Since this
26  * screws up the old-style command line protocol, adjust by
27  * filling in the new-style command line pointer instead.
28  */
29
30  static void copy_boot_params(void)
31  {
32      struct old_cmdline {
33          u16 cl_magic;
34          u16 cl_offset;
35      };
36      const struct old_cmdline * const oldcmd =
37          (const struct old_cmdline *)OLD_CL_ADDRESS;
38
39      BUILD_BUG_ON(sizeof boot_params != 4096);
40      memcpy(&boot_params.hdr, &hdr, sizeof hdr);
41
42      if (!boot_params.hdr.cmd_line_ptr &&
43          oldcmd->cl_magic == OLD_CL_MAGIC) {
44          /* Old-style command line protocol. */
45          u16 cmdline_seg;
46
47          /* Figure out if the command line falls in the region
48             of memory that an old kernel would have copied up
49             to 0x90000... */
50          if (oldcmd->cl_offset < boot_params.hdr.setup_move_size)
51              cmdline_seg = ds();
52          else
53              cmdline_seg = 0x9000;
54
55          boot_params.hdr.cmd_line_ptr =
56              (cmdline_seg << 4) + oldcmd->cl_offset;
57      }
58  }
59
```

Start_kernel Initialization

