

# Multiboot Specification version 0.6.96

---

Next: [Overview](#), Up: [\(dir\)](#)

## Multiboot Specification

This file documents Multiboot Specification, the proposal for the boot sequence standard. This edition documents version 0.6.96.

Copyright © 1995,96 Bryan Ford <baford@cs.utah.edu>

Copyright © 1995,96 Erich Stefan Boleyn <erich@uruk.org>

Copyright © 1999,2000,2001,2002,2005,2006,2009,2010 Free Software Foundation, Inc.

Permission is granted to make and distribute verbatim copies of this manual provided the copyright notice and this permission notice are preserved on all copies.

Permission is granted to copy and distribute modified versions of this manual under the conditions for verbatim copying, provided also that the entire resulting derived work is distributed under the terms of a permission notice identical to this one.

Permission is granted to copy and distribute translations of this manual into another language, under the above conditions for modified versions.

- [Overview](#)
- [Terminology](#)
- [Specification](#)
- [Examples](#)
- [History](#)
- [Invoking mbchk](#): How to use the Multiboot checker
- [Index](#)

---

Next: [Terminology](#), Previous: [Top](#), Up: [Top](#)

## 1 Introduction to Multiboot Specification

This chapter describes some rough information on the Multiboot Specification. Note that this is not a part of the specification itself.

- [Motivation](#)
- [Architecture](#)
- [Operating systems](#)
- [Boot sources](#)
- [Boot-time configuration](#)
- [Convenience to operating systems](#)
- [Boot modules](#)

---

Next: [Architecture](#), Up: [Overview](#)

### 1.1 The background of Multiboot Specification

Every operating system ever created tends to have its own boot loader. Installing a new operating system on a machine generally involves installing a whole new set of boot mechanisms, each with completely different install-time and boot-time user interfaces. Getting multiple operating systems to coexist reliably on one machine through typical *chaining* mechanisms can be a nightmare. There is little or no choice of boot loaders for a particular operating system — if the one that comes with the operating system doesn't do exactly what you want, or doesn't work on your machine, you're screwed.

While we may not be able to fix this problem in existing proprietary operating systems, it shouldn't be too difficult for a few people in the free operating system communities to put their heads together and solve this problem for the popular free operating systems. That's what this specification aims for. Basically, it specifies an interface between a boot loader and a operating system, such that any complying boot loader should be able to load any complying operating system. This specification does *not* specify how boot loaders should work — only how they must interface with the operating system being loaded.

---

Next: [Operating systems](#), Previous: [Motivation](#), Up: [Overview](#)

### 1.2 The target architecture

This specification is primarily targeted at PC, since they are the most common and have the largest variety of operating systems and boot loaders. However, to the extent that certain other architectures may need a boot specification and do not have one already, a variation of this specification, stripped of the x86-specific details, could be adopted for them as well.

---

Next: [Boot sources](#), Previous: [Architecture](#), Up: [Overview](#)

### 1.3 The target operating systems

This specification is targeted toward free 32-bit operating systems that can be fairly easily modified to support the specification without going through lots of bureaucratic rigmarole. The particular free operating systems that this specification is being primarily designed for are Linux, the kernels of FreeBSD and NetBSD, Mach, and VSTa. It is hoped that other emerging free operating systems will adopt it from the start, and thus immediately be able to take advantage of existing boot loaders. It would be nice if proprietary operating system vendors eventually adopted this specification as well, but that's probably a pipe dream.

---

## 1.4 Boot sources

It should be possible to write compliant boot loaders that load the OS image from a variety of sources, including floppy disk, hard disk, and across a network.

Disk-based boot loaders may use a variety of techniques to find the relevant OS image and boot module data on disk, such as by interpretation of specific file systems (e.g. the BSD/Mach boot loader), using precalculated *blocklists* (e.g. LILO), loading from a special *boot partition* (e.g. OS/2), or even loading from within another operating system (e.g. the VSTa boot code, which loads from DOS). Similarly, network-based boot loaders could use a variety of network hardware and protocols.

It is hoped that boot loaders will be created that support multiple loading mechanisms, increasing their portability, robustness, and user-friendliness.

## 1.5 Configure an operating system at boot-time

It is often necessary for one reason or another for the user to be able to provide some configuration information to an operating system dynamically at boot time. While this specification should not dictate how this configuration information is obtained by the boot loader, it should provide a standard means for the boot loader to pass such information to the operating system.

## 1.6 How to make OS development easier

OS images should be easy to generate. Ideally, an OS image should simply be an ordinary 32-bit executable file in whatever file format the operating system normally uses. It should be possible to nm or disassemble OS images just like normal executables. Specialized tools should not be required to create OS images in a *special* file format. If this means shifting some work from the operating system to a boot loader, that is probably appropriate, because all the memory consumed by the boot loader will typically be made available again after the boot process is created, whereas every bit of code in the OS image typically has to remain in memory forever. The operating system should not have to worry about getting into 32-bit mode initially, because mode switching code generally needs to be in the boot loader anyway in order to load operating system data above the 1MB boundary, and forcing the operating system to do this makes creation of OS images much more difficult.

Unfortunately, there is a horrendous variety of executable file formats even among free Unix-like PC-based operating systems — generally a different format for each operating system. Most of the relevant free operating systems use some variant of a.out format, but some are moving to ELF. It is highly desirable for boot loaders not to have to be able to interpret all the different types of executable file formats in existence in order to load the OS image — otherwise the boot loader effectively becomes operating system specific again.

This specification adopts a compromise solution to this problem. Multiboot-compliant OS images always contain a magic *Multiboot header* (see [OS image format](#)), which allows the boot loader to load the image without having to understand numerous a.out variants or other executable formats. This magic header does not need to be at the very beginning of the executable file, so kernel images can still conform to the local a.out format variant in addition to being Multiboot-compliant.

## 1.7 Boot modules

Many modern operating system kernels, such as Mach and the microkernel in VSTa, do not by themselves contain enough mechanism to get the system fully operational: they require the presence of additional software modules at boot time in order to access devices, mount file systems, etc. While these additional modules could be embedded in the main OS image along with the kernel itself, and the resulting image be split apart manually by the operating system when it receives control, it is often more flexible, more space-efficient, and more convenient to the operating system and user if the boot loader can load these additional modules independently in the first place.

Thus, this specification should provide a standard method for a boot loader to indicate to the operating system what auxiliary boot modules were loaded, and where they can be found. Boot loaders don't have to support multiple boot modules, but they are strongly encouraged to, because some operating systems will be unable to boot without them.

# 2 The definitions of terms used through the specification

### *must*

We use the term *must*, when any boot loader or OS image needs to follow a rule — otherwise, the boot loader or OS image is *not* Multiboot-compliant.

### *should*

We use the term *should*, when any boot loader or OS image is recommended to follow a rule, but it doesn't need to follow the rule.

### *may*

We use the term *may*, when any boot loader or OS image is allowed to follow a rule.

### *boot loader*

Whatever program or set of programs loads the image of the final operating system to be run on the machine. The boot loader may itself consist of several stages, but that is an implementation detail not relevant to this specification. Only the *final* stage of the boot loader — the stage that eventually transfers control to an operating system — must follow the rules specified in this document in order to be *Multiboot-compliant*; earlier boot loader stages may be designed in whatever way is most convenient.

### *OS image*

The initial binary image that a boot loader loads into memory and transfers control to start an operating system. The OS image is typically an executable containing the operating system kernel.

### *boot module*

Other auxiliary files that a boot loader loads into memory along with an OS image, but does not interpret in any way other than passing their locations to the operating system when it is invoked.

<i>Multiboot-compliant</i>	A boot loader or an OS image which follows the rules defined as <i>must</i> is Multiboot-compliant. When this specification specifies a rule as <i>should</i> or <i>may</i> , a Multiboot-complaint boot loader/OS image doesn't need to follow the rule.
<i>u8</i>	The type of unsigned 8-bit data.
<i>u16</i>	The type of unsigned 16-bit data. Because the target architecture is little-endian, u16 is coded in little-endian.
<i>u32</i>	The type of unsigned 32-bit data. Because the target architecture is little-endian, u32 is coded in little-endian.
<i>u64</i>	The type of unsigned 64-bit data. Because the target architecture is little-endian, u64 is coded in little-endian.

Next: [Examples](#), Previous: [Terminology](#), Up: [Top](#)

### 3 The exact definitions of Multiboot Specification

There are three main aspects of a boot loader/OS image interface:

1. The format of an OS image as seen by a boot loader.
2. The state of a machine when a boot loader starts an operating system.
3. The format of information passed by a boot loader to an operating system.

- [OS image format](#)
- [Machine state](#)
- [Boot information format](#)

Next: [Machine state](#), Up: [Specification](#)

#### 3.1 OS image format

An OS image may be an ordinary 32-bit executable file in the standard format for that particular operating system, except that it may be linked at a non-default load address to avoid loading on top of the pc's I/O region or other reserved areas, and of course it should not use shared libraries or other fancy features.

An OS image must contain an additional header called *Multiboot header*, besides the headers of the format used by the OS image. The Multiboot header must be contained completely within the first 8192 bytes of the OS image, and must be longword (32-bit) aligned. In general, it should come *as early as possible*, and may be embedded in the beginning of the text segment after the *real* executable header.

- [Header layout](#): The layout of Multiboot header
- [Header magic fields](#): The magic fields of Multiboot header
- [Header address fields](#)
- [Header graphics fields](#)

Next: [Header magic fields](#), Up: [OS image format](#)

##### 3.1.1 The layout of Multiboot header

The layout of the Multiboot header must be as follows:

Offset	Type	Field Name	Note
0	u32	magic	required
4	u32	flags	required
8	u32	checksum	required
12	u32	header_addr	if flags[16] is set
16	u32	load_addr	if flags[16] is set
20	u32	load_end_addr	if flags[16] is set
24	u32	bss_end_addr	if flags[16] is set
28	u32	entry_addr	if flags[16] is set
32	u32	mode_type	if flags[2] is set
36	u32	width	if flags[2] is set
40	u32	height	if flags[2] is set
44	u32	depth	if flags[2] is set

The fields ‘magic’, ‘flags’ and ‘checksum’ are defined in [Header magic fields](#), the fields ‘header\_addr’, ‘load\_addr’, ‘load\_end\_addr’, ‘bss\_end\_addr’ and ‘entry\_addr’ are defined in [Header address fields](#), and the fields ‘mode\_type’, ‘width’, ‘height’ and ‘depth’ are defined in [Header graphics fields](#).

Next: [Header address fields](#), Previous: [Header layout](#), Up: [OS image format](#)

##### 3.1.2 The magic fields of Multiboot header

‘magic’  
The field ‘magic’ is the magic number identifying the header, which must be the hexadecimal value 0x1BADB002.

‘flags’

The field ‘flags’ specifies features that the OS image requests or requires of an boot loader. Bits 0-15 indicate requirements; if the boot loader sees any of these bits set but doesn't understand the flag or can't fulfill the requirements it indicates for some reason, it must notify the user and fail to load the OS image. Bits 16-31 indicate optional features; if any bits in this range are set but the boot loader doesn't understand them, it may simply ignore them and



proceed as usual. Naturally, all as-yet-undefined bits in the ‘ <code>flags</code> ’ word must be set to zero in OS images. This way, the ‘ <code>flags</code> ’ fields serves for version control as well as simple feature selection.	
If bit 0 in the ‘ <code>flags</code> ’ word is set, then all boot modules loaded along with the operating system must be aligned on page (4KB) boundaries. Some operating systems expect to be able to map the pages containing boot modules directly into a paged address space during startup, and thus need the boot modules to be page-aligned.	
If bit 1 in the ‘ <code>flags</code> ’ word is set, then information on available memory via at least the ‘ <code>mem_*</code> ’ fields of the Multiboot information structure (see <a href="#">Boot information format</a> ) must be included. If the boot loader is capable of passing a memory map (the ‘ <code>mmap_*</code> ’ fields) and one exists, then it may be included as well.	
If bit 2 in the ‘ <code>flags</code> ’ word is set, information about the video mode table (see <a href="#">Boot information format</a> ) must be available to the kernel.	
If bit 16 in the ‘ <code>flags</code> ’ word is set, then the fields at offsets 12-28 in the Multiboot header are valid, and the boot loader should use them instead of the fields in the actual executable header to calculate where to load the OS image. This information does not need to be provided if the kernel image is in ELF format, but it <i>must</i> be provided if the images is in a.out format or in some other format. Compliant boot loaders must be able to load images that either are in ELF format or contain the load address information embedded in the Multiboot header; they may also directly support other executable formats, such as particular a.out variants, but are not required to.	
‘ <code>checksum</code> ’ The field ‘ <code>checksum</code> ’ is a 32-bit unsigned value which, when added to the other magic fields (i.e. ‘ <code>magic</code> ’ and ‘ <code>flags</code> ’), must have a 32-bit unsigned sum of zero.	

Next: [Header graphics fields](#), Previous: [Header magic fields](#), Up: [OS image format](#)

### 3.1.3 The address fields of Multiboot header

All of the address fields enabled by flag bit 16 are physical addresses. The meaning of each is as follows:

header_addr	Contains the address corresponding to the beginning of the Multiboot header — the physical memory location at which the magic value is supposed to be loaded. This field serves to <i>synchronize</i> the mapping between OS image offsets and physical memory addresses.
load_addr	Contains the physical address of the beginning of the text segment. The offset in the OS image file at which to start loading is defined by the offset at which the header was found, minus (header_addr - load_addr). load_addr must be less than or equal to header_addr.
load_end_addr	Contains the physical address of the end of the data segment. (load_end_addr - load_addr) specifies how much data to load. This implies that the text and data segments must be consecutive in the OS image; this is true for existing a.out executable formats. If this field is zero, the boot loader assumes that the text and data segments occupy the whole OS image file.
bss_end_addr	Contains the physical address of the end of the bss segment. The boot loader initializes this area to zero, and reserves the memory it occupies to avoid placing boot modules and other data relevant to the operating system in that area. If this field is zero, the boot loader assumes that no bss segment is present.
entry_addr	The physical address to which the boot loader should jump in order to start running the operating system.

Previous: [Header address fields](#), Up: [OS image format](#)

### 3.1.4 The graphics fields of Multiboot header

All of the graphics fields are enabled by flag bit 2. They specify the preferred graphics mode. Note that that is only a *recommended* mode by the OS image. Boot loader may choose a different mode if it sees fit.

The meaning of each is as follows:

mode_type	Contains ‘0’ for linear graphics mode or ‘1’ for EGA-standard text mode. Everything else is reserved for future expansion. Note that the boot loader may set a text mode even if this field contains ‘0’, or set a video mode even if this field contains ‘1’.
width	Contains the number of the columns. This is specified in pixels in a graphics mode, and in characters in a text mode. The value zero indicates that the OS image has no preference.
height	Contains the number of the lines. This is specified in pixels in a graphics mode, and in characters in a text mode. The value zero indicates that the OS image has no preference.
depth	Contains the number of bits per pixel in a graphics mode, and zero in a text mode. The value zero indicates that the OS image has no preference.

Next: [Boot information format](#), Previous: [OS image format](#), Up: [Specification](#)

## 3.2 Machine state

When the boot loader invokes the 32-bit operating system, the machine must have the following state:

‘EAX’	Must contain the magic value ‘0x2BADB002’; the presence of this value indicates to the operating system that it was loaded by a Multiboot-compliant boot loader (e.g. as opposed to another type of boot loader that the operating system can also be loaded from).
‘EBX’	Must contain the 32-bit physical address of the Multiboot information structure provided by the boot loader (see <a href="#">Boot information format</a> ).
‘CS’	Must be a 32-bit read/execute code segment with an offset of ‘0’ and a limit of ‘0xFFFFFFFF’. The exact value is undefined.
‘DS’	
‘ES’	

FS  
‘GS’  
‘SS’  
Must be a 32-bit read/write data segment with an offset of ‘0’ and a limit of ‘0xFFFFFFFF’. The exact values are all undefined.

‘A20 gate’  
Must be enabled.

‘CR0’  
Bit 31 (PG) must be cleared. Bit 0 (PE) must be set. Other bits are all undefined.

‘EFLAGS’  
Bit 17 (VM) must be cleared. Bit 9 (IF) must be cleared. Other bits are all undefined.

All other processor registers and flag bits are undefined. This includes, in particular:

‘ESP’  
The OS image must create its own stack as soon as it needs one.

‘GDTR’  
Even though the segment registers are set up as described above, the ‘GDTR’ may be invalid, so the OS image must not load any segment registers (even just reloading the same values!) until it sets up its own ‘GDT’.

‘IDTR’  
The OS image must leave interrupts disabled until it sets up its own IDT.

However, other machine state should be left by the boot loader in *normal working order*, i.e. as initialized by the BIOS (or DOS, if that's what the boot loader runs from). In other words, the operating system should be able to make BIOS calls and such after being loaded, as long as it does not overwrite the BIOS data structures before doing so. Also, the boot loader must leave the PIC programmed with the normal BIOS/DOS values, even if it changed them during the switch to 32-bit mode.

Previous: [Machine state](#), Up: [Specification](#)

### 3.3 Boot information format

FIXME: Split this chapter like the chapter “OS image format”.

Upon entry to the operating system, the EBX register contains the physical address of a *Multiboot information* data structure, through which the boot loader communicates vital information to the operating system. The operating system can use or ignore any parts of the structure as it chooses; all information passed by the boot loader is advisory only.

The Multiboot information structure and its related substructures may be placed anywhere in memory by the boot loader (with the exception of the memory reserved for the kernel and boot modules, of course). It is the operating system's responsibility to avoid overwriting this memory until it is done using it.

The format of the Multiboot information structure (as defined so far) follows:


0	+-----+   flags   +-----+	(required)
4	mem_lower   8   mem_upper   +-----+	(present if flags[0] is set) (present if flags[0] is set)
12	boot_device   +-----+	(present if flags[1] is set)
16	cmdline   +-----+	(present if flags[2] is set)
20	mods_count   24   mods_addr   +-----+	(present if flags[3] is set) (present if flags[3] is set)
28 – 40	syms   +-----+	(present if flags[4] or flags[5] is set)
44	mmap_length   48   mmap_addr   +-----+	(present if flags[6] is set) (present if flags[6] is set)
52	drives_length   56   drives_addr   +-----+	(present if flags[7] is set) (present if flags[7] is set)
60	config_table   +-----+	(present if flags[8] is set)
64	boot_loader_name   +-----+	(present if flags[9] is set)
68	apm_table   +-----+	(present if flags[10] is set)
72	vbe_control_info   76   vbe_mode_info   80   vbe_mode   82   vbe_interface_seg   84   vbe_interface_off   86   vbe_interface_len   +-----+	(present if flags[11] is set)
88	framebuffer_addr   96   framebuffer_pitch   100   framebuffer_width   104   framebuffer_height   108   framebuffer_bpp   109   framebuffer_type   110–115   color_info   +-----+	(present if flags[12] is set)

The first longword indicates the presence and validity of other fields in the Multiboot information structure. All as-yet-undefined bits must be set to zero by the boot loader. Any set bits that the operating system does not understand should be ignored. Thus, the ‘flags’ field also functions as a version indicator, allowing the Multiboot information structure to be expanded in the future without breaking anything.

If bit 0 in the ‘flags’ word is set, then the ‘mem\_’ fields are valid. ‘mem\_lower’ and ‘mem\_upper’ indicate the amount of lower and upper memory, respectively, in

kilobytes. Lower memory starts at address 0, and upper memory starts at address 1 megabyte. The maximum possible value for lower memory is 640 kilobytes. The value returned for upper memory is maximally the address of the first upper memory hole minus 1 megabyte. It is not guaranteed to be this value.

If bit 1 in the ‘flags’ word is set, then the ‘boot\_device’ field is valid, and indicates which BIOS disk device the boot loader loaded the OS image from. If the OS image was not loaded from a BIOS disk, then this field must not be present (bit 3 must be clear). The operating system may use this field as a hint for determining its own *root* device, but is not required to. The ‘boot\_device’ field is laid out in four one-byte subfields as follows:




The most significant byte contains the BIOS drive number as understood by the BIOS INT 0x13 low-level disk interface: e.g. 0x00 for the first floppy disk or 0x80 for the first hard disk.

The three remaining bytes specify the boot partition. ‘part1’ specifies the *top-level* partition number, ‘part2’ specifies a *sub-partition* in the top-level partition, etc. Partition numbers always start from zero. Unused partition bytes must be set to 0xFF. For example, if the disk is partitioned using a simple one-level DOS partitioning scheme, then ‘part1’ contains the DOS partition number, and ‘part2’ and ‘part3’ are both 0xFF. As another example, if a disk is partitioned first into DOS partitions, and then one of those DOS partitions is subdivided into several BSD partitions using BSD's *disklabel* strategy, then ‘part1’ contains the DOS partition number, ‘part2’ contains the BSD sub-partition within that DOS partition, and ‘part3’ is 0xFF.

DOS extended partitions are indicated as partition numbers starting from 4 and increasing, rather than as nested sub-partitions, even though the underlying disk layout of extended partitions is hierarchical in nature. For example, if the boot loader boots from the second extended partition on a disk partitioned in conventional DOS style, then ‘part1’ will be 5, and ‘part2’ and ‘part3’ will both be 0xFF.

If bit 2 of the ‘flags’ longword is set, the ‘cmdline’ field is valid, and contains the physical address of the command line to be passed to the kernel. The command line is a normal C-style zero-terminated string. The exact format of command line is left to OS developpers. General-purpose boot loaders should allow user a complete control on command line independently of other factors like image name. Boot loaders with specific payload in mind may completely or partially generate it algorithmically.


If bit 3 of the ‘flags’ is set, then the ‘mods’ fields indicate to the kernel what boot modules were loaded along with the kernel image, and where they can be found. ‘mods\_count’ contains the number of modules loaded; ‘mods\_addr’ contains the physical address of the first module structure. ‘mods\_count’ may be zero, indicating no boot modules were loaded, even if bit 3 of ‘flags’ is set. Each module structure is formatted as follows:



The first two fields contain the start and end addresses of the boot module itself. The ‘string’ field provides an arbitrary string to be associated with that particular boot module; it is a zero-terminated ASCII string, just like the kernel command line. The ‘string’ field may be 0 if there is no string associated with the module. Typically the string might be a command line (e.g. if the operating system treats boot modules as executable programs), or a pathname (e.g. if the operating system treats boot modules as files in a file system), but its exact use is specific to the operating system. The ‘reserved’ field must be set to 0 by the boot loader and ignored by the operating system.


**Caution:** Bits 4 & 5 are mutually exclusive.

If bit 4 in the ‘flags’ word is set, then the following fields in the Multiboot information structure starting at byte 28 are valid:




These indicate where the symbol table from an a.out kernel image can be found. ‘addr’ is the physical address of the size (4-byte unsigned long) of an array of a.out format *nlist* structures, followed immediately by the array itself, then the size (4-byte unsigned long) of a set of zero-terminated ASCII strings (plus sizeof(unsigned long) in this case), and finally the set of strings itself. ‘tabsize’ is equal to its size parameter (found at the beginning of the symbol section), and ‘strsize’ is equal to its size parameter (found at the beginning of the string section) of the following string table to which the symbol table refers. Note that ‘tabsize’ may be 0, indicating no symbols, even if bit 4 in the ‘flags’ word is set.

If bit 5 in the ‘flags’ word is set, then the following fields in the Multiboot information structure starting at byte 28 are valid:



These indicate where the section header table from an ELF kernel is, the size of each entry, number of entries, and the string table used as the index of names. They correspond to the ‘shdr\_\*’ entries (‘shdr\_num’, etc.) in the Executable and Linkable Format (ELF) specification in the program header. All sections are loaded, and the physical address fields of the ELF section header then refer to where the sections are in memory (refer to the i386 ELF documentation for details as to how to read the section header(s)). Note that ‘shdr\_num’ may be 0, indicating no symbols, even if bit 5 in the ‘flags’ word is set.

If bit 6 in the ‘flags’ word is set, then the ‘mmap\_\*’ fields are valid, and indicate the address and length of a buffer containing a memory map of the machine provided by the BIOS. ‘mmap\_addr’ is the address, and ‘mmap\_length’ is the total size of the buffer. The buffer consists of one or more of the following size/structure pairs (‘size’ is really used for skipping to the next pair):





8		length	
16		type	
+-----+			

where ‘size’ is the size of the associated structure in bytes, which can be greater than the minimum of 20 bytes. ‘base\_addr’ is the starting address. ‘length’ is the size of the memory region in bytes. ‘type’ is the variety of address range represented, where a value of 1 indicates available RAM, value of 3 indicates usable memory holding ACPI information, value of 4 indicates reserved memory which needs to be preserved on hibernation, value of 5 indicates a memory which is occupied by defective RAM modules and all other values currently indicated a reserved area.

The map provided is guaranteed to list all standard RAM that should be available for normal use.

If bit 7 in the ‘flags’ is set, then the ‘drives\_’ fields are valid, and indicate the address of the physical address of the first drive structure and the size of drive structures. ‘drives\_addr’ is the address, and ‘drives\_length’ is the total size of drive structures. Note that ‘drives\_length’ may be zero. Each drive structure is formatted as follows:

		+-----+
0		size
		+
4		drive_number
		+
5		drive_mode
		+
6		drive_cylinders
8		drive_heads
9		drive_sectors
		+
10 - xx		drive_ports
		+

The ‘size’ field specifies the size of this structure. The size varies, depending on the number of ports. Note that the size may not be equal to (10 + 2 \* the number of ports), because of an alignment.

The ‘drive\_number’ field contains the BIOS drive number. The ‘drive\_mode’ field represents the access mode used by the boot loader. Currently, the following modes are defined:

- ‘0’ CHS mode (traditional cylinder/head/sector addressing mode).
- ‘1’ LBA mode (Logical Block Addressing mode).

The three fields, ‘drive\_cylinders’, ‘drive\_heads’ and ‘drive\_sectors’, indicate the geometry of the drive detected by the BIOS. ‘drive\_cylinders’ contains the number of the cylinders. ‘drive\_heads’ contains the number of the heads. ‘drive\_sectors’ contains the number of the sectors per track.

The ‘drive\_ports’ field contains the array of the I/O ports used for the drive in the BIOS code. The array consists of zero or more unsigned two-bytes integers, and is terminated with zero. Note that the array may contain any number of I/O ports that are not related to the drive actually (such as DMA controller's ports).

If bit 8 in the ‘flags’ is set, then the ‘config\_table’ field is valid, and indicates the address of the ROM configuration table returned by the *GET CONFIGURATION* BIOS call. If the BIOS call fails, then the size of the table must be *zero*.

If bit 9 in the ‘flags’ is set, the ‘boot\_loader\_name’ field is valid, and contains the physical address of the name of a boot loader booting the kernel. The name is a normal C-style zero-terminated string.

If bit 10 in the ‘flags’ is set, the ‘apm\_table’ field is valid, and contains the physical address of an APM table defined as below:

	+-----+
0	version
2	cseg
4	offset
8	cseg_16
10	dseg
12	flags
14	cseg_len
16	cseg_16_len
18	dseg_len
	+

The fields ‘version’, ‘cseg’, ‘offset’, ‘cseg\_16’, ‘dseg’, ‘flags’, ‘cseg\_len’, ‘cseg\_16\_len’, ‘dseg\_len’ indicate the version number, the protected mode 32-bit code segment, the offset of the entry point, the protected mode 16-bit code segment, the protected mode 16-bit data segment, the flags, the length of the protected mode 32-bit code segment, the length of the protected mode 16-bit code segment, and the length of the protected mode 16-bit data segment, respectively. Only the field ‘offset’ is 4 bytes, and the others are 2 bytes. See [Advanced Power Management \(APM\) BIOS Interface Specification](#), for more information.

If bit 11 in the ‘flags’ is set, the VBE table is available.

The fields ‘vbe\_control\_info’ and ‘vbe\_mode\_info’ contain the physical addresses of VBE control information returned by the VBE Function 00h and VBE mode information returned by the VBE Function 01h, respectively.

The field ‘vbe\_mode’ indicates current video mode in the format specified in VBE 3.0.

The rest fields ‘vbe\_interface\_seg’, ‘vbe\_interface\_off’, and ‘vbe\_interface\_len’ contain the table of a protected mode interface defined in VBE 2.0+. If this information is not available, those fields contain zero. Note that VBE 3.0 defines another protected mode interface which is incompatible with the old one. If you want to use the new protected mode interface, you will have to find the table yourself.

The fields for the graphics table are designed for VBE, but Multiboot boot loaders may simulate VBE on non-VBE modes, as if they were VBE modes.

If bit 12 in the ‘flags’ is set, the FRAMEBUFFER table is available.

The field ‘framebuffer\_addr’ contains framebuffer physical address. This field is 64-bit wide but bootloader *should* set it under 4 GiB if possible for

compatibility with kernels which aren't aware of PAE or AMD64. The field ‘framebuffer\_pitch’ contains the framebuffer pitch in bytes. The fields ‘framebuffer\_width’, ‘framebuffer\_height’ contain the framebuffer dimensions in pixels. The field ‘framebuffer\_bpp’ contains the number of bits per pixel. If ‘framebuffer\_type’ is set to ‘0’ it means indexed color will be used. In this case color\_info is defined as follows:

	+-----+
110	framebuffer_palette_addr
114	framebuffer_palette_num_colors
	+-----+

‘framebuffer\_palette\_addr’ contains the address of the color palette, which is an array of color descriptors. Each color descriptor has the following structure:

	+-----+
0	red_value
1	green_value
2	blue_value
	+-----+

If ‘framebuffer\_type’ is set to ‘1’ it means direct RGB color will be used. Then color\_type is defined as follows:

	+-----+
110	framebuffer_red_field_position
111	framebuffer_red_mask_size
112	framebuffer_green_field_position
113	framebuffer_green_mask_size
114	framebuffer_blue_field_position
115	framebuffer_blue_mask_size
	+-----+

If ‘framebuffer\_type’ is set to ‘2’ it means EGA-standard text mode will be used. In this case ‘framebuffer\_width’ and ‘framebuffer\_height’ are expressed in characters instead of pixels. ‘framebuffer\_bpp’ is equal to 16 (bits per character) and ‘framebuffer\_pitch’ is expressed in bytes per text line. All further values of ‘framebuffer\_type’ are reserved for future expansion.

Next: [History](#), Previous: [Specification](#), Up: [Top](#)

## 4 Examples

**Caution:** The following items are not part of the specification document, but are included for prospective operating system and boot loader writers.

- [Notes on PC](#)
- [BIOS device mapping techniques](#)
- [Example OS code](#)
- [Example boot loader code](#)

Next: [BIOS device mapping techniques](#), Up: [Examples](#)

### 4.1 Notes on PC

In reference to bit 0 of the ‘flags’ parameter in the Multiboot information structure, if the bootloader in question uses older BIOS interfaces, or the newest ones are not available (see description about bit 6), then a maximum of either 15 or 63 megabytes of memory may be reported. It is *highly* recommended that boot loaders perform a thorough memory probe.

In reference to bit 1 of the ‘flags’ parameter in the Multiboot information structure, it is recognized that determination of which BIOS drive maps to which device driver in an operating system is non-trivial, at best. Many kludges have been made to various operating systems instead of solving this problem, most of them breaking under many conditions. To encourage the use of general-purpose solutions to this problem, there are 2 BIOS device mapping techniques (see [BIOS device mapping techniques](#)).

In reference to bit 6 of the ‘flags’ parameter in the Multiboot information structure, it is important to note that the data structure used there (starting with ‘BaseAddrLow’) is the data returned by the INT 15h, AX=E820h — Query System Address Map call. See See [Query System Address Map](#), for more information. The interface here is meant to allow a boot loader to work unmodified with any reasonable extensions of the BIOS interface, passing along any extra data to be interpreted by the operating system as desired.

Next: [Example OS code](#), Previous: [Notes on PC](#), Up: [Examples](#)

### 4.2 BIOS device mapping techniques

Both of these techniques should be usable from any PC operating system, and neither require any special support in the drivers themselves. This section will be flushed out into detailed explanations, particularly for the I/O restriction technique.

The general rule is that the data comparison technique is the quick and dirty solution. It works most of the time, but doesn't cover all the bases, and is relatively simple.

The I/O restriction technique is much more complex, but it has potential to solve the problem under all conditions, plus allow access of the remaining BIOS devices when not all of them have operating system drivers.

- [Data comparison technique](#)
- [I/O restriction technique](#)

Next: [I/O restriction technique](#), Up: [BIOS device mapping techniques](#)

#### 4.2.1 Data comparison technique

Before activating *any* of the device drivers, gather enough data from similar sectors on each of the disks such that each one can be uniquely identified.



After activating the device drivers, compare data from the drives using the operating system drivers. This should hopefully be sufficient to provide such a mapping.

Problems:

1. The data on some BIOS devices might be identical (so the part reading the drives from the BIOS should have some mechanism to give up).
2. There might be extra drives not accessible from the BIOS which are identical to some drive used by the BIOS (so it should be capable of giving up there as well).

---

Previous: [Data comparison technique](#), Up: [BIOS device mapping techniques](#)

#### 4.2.2 I/O restriction technique

This first step may be unnecessary, but first create copy-on-write mappings for the device drivers writing into PC RAM. Keep the original copies for the *clean* BIOS *virtual machine* to be created later.

For each device driver brought online, determine which BIOS devices become inaccessible by:

1. Create a *clean* BIOS *virtual machine*.
2. Set the I/O permission map for the I/O area claimed by the device driver to no permissions (neither read nor write).
3. Access each device.
4. Record which devices succeed, and those which try to access the *restricted* I/O areas (hopefully, this will be an *xor* situation).

For each device driver, given how many of the BIOS devices were subsumed by it (there should be no gaps in this list), it should be easy to determine which devices on the controller these are.

In general, you have at most 2 disks from each controller given BIOS numbers, but they pretty much always count from the lowest logically numbered devices on the controller.

---

Next: [Example boot loader code](#), Previous: [BIOS device mapping techniques](#), Up: [Examples](#)

### 4.3 Example OS code

In this distribution, the example Multiboot kernel `kernel` is included. The kernel just prints out the Multiboot information structure on the screen, so you can make use of the kernel to test a Multiboot-compliant boot loader and for reference to how to implement a Multiboot kernel. The source files can be found under the directory `doc` in the Multiboot source distribution.

The kernel `kernel` consists of only three files: `boot.s`, `kernel.c` and `multiboot.h`. The assembly source `boot.s` is written in GAS (see [GNU assembler](#)), and contains the Multiboot information structure to comply with the specification. When a Multiboot-compliant boot loader loads and execute it, it initialize the stack pointer and `EFLAGS`, and then call the function `cmain` defined in `kernel.c`. If `cmain` returns to the callee, then it shows a message to inform the user of the halt state and stops forever until you push the reset key. The file `kernel.c` contains the function `cmain`, which checks if the magic number passed by the boot loader is valid and so on, and some functions to print messages on the screen. The file `multiboot.h` defines some macros, such as the magic number for the Multiboot header, the Multiboot header structure and the Multiboot information structure.

- [multiboot.h](#)
- [boot.S](#)
- [kernel.c](#)
- [Other Multiboot kernels](#)

---

Next: [boot.S](#), Up: [Example OS code](#)

#### 4.3.1 multiboot.h

This is the source code in the file `multiboot.h`:

```
/* multiboot.h - Multiboot header file. */
/* Copyright (C) 1999,2003,2007,2008,2009,2010 Free Software Foundation, Inc.
*
* Permission is hereby granted, free of charge, to any person obtaining a copy
* of this software and associated documentation files (the "Software"), to
* deal in the Software without restriction, including without limitation the
* rights to use, copy, modify, merge, publish, distribute, sublicense, and/or
* sell copies of the Software, and to permit persons to whom the Software is
* furnished to do so, subject to the following conditions:
*
* The above copyright notice and this permission notice shall be included in
* all copies or substantial portions of the Software.
*
* THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR
* IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY,
* FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL ANY
* DEVELOPER OR DISTRIBUTOR BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY,
* WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR
* IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.
*/

#ifndef MULTIBOOT_HEADER
#define MULTIBOOT_HEADER 1

/* How many bytes from the start of the file we search for the header. */
#define MULTIBOOT_SEARCH 8192
```

```

#define MULTIBOOT_HEADER_ALIGN 4

/* The magic field should contain this. */
#define MULTIBOOT_HEADER_MAGIC 0x1BADB002

/* This should be in %eax. */
#define MULTIBOOT_BOOTLOADER_MAGIC 0x2BADB002

/* Alignment of multiboot modules. */
#define MULTIBOOT_MOD_ALIGN 0x00001000

/* Alignment of the multiboot info structure. */
#define MULTIBOOT_INFO_ALIGN 0x00000004

/* Flags set in the 'flags' member of the multiboot header. */

/* Align all boot modules on i386 page (4KB) boundaries. */
#define MULTIBOOT_PAGE_ALIGN 0x00000001

/* Must pass memory information to OS. */
#define MULTIBOOT_MEMORY_INFO 0x00000002

/* Must pass video information to OS. */
#define MULTIBOOT_VIDEO_MODE 0x00000004

/* This flag indicates the use of the address fields in the header. */
#define MULTIBOOT_AOUT_KLUDGE 0x00010000

/* Flags to be set in the 'flags' member of the multiboot info structure. */

/* is there basic lower/upper memory information? */
#define MULTIBOOT_INFO_MEMORY 0x00000001
/* is there a boot device set? */
#define MULTIBOOT_INFO_BOOTDEV 0x00000002
/* is the command-line defined? */
#define MULTIBOOT_INFO_CMDLINE 0x00000004
/* are there modules to do something with? */
#define MULTIBOOT_INFO_MODS 0x00000008

/* These next two are mutually exclusive */

/* is there a symbol table loaded? */
#define MULTIBOOT_INFO_AOUT_SYMS 0x00000010
/* is there an ELF section header table? */
#define MULTIBOOT_INFO_ELF_SHDR 0X00000020

/* is there a full memory map? */
#define MULTIBOOT_INFO_MEM_MAP 0x00000040

/* Is there drive info? */
#define MULTIBOOT_INFO_DRIVE_INFO 0x00000080

/* Is there a config table? */
#define MULTIBOOT_INFO_CONFIG_TABLE 0x00000100

/* Is there a boot loader name? */
#define MULTIBOOT_INFO_BOOT_LOADER_NAME 0x00000200

/* Is there a APM table? */
#define MULTIBOOT_INFO_APM_TABLE 0x00000400

/* Is there video information? */
#define MULTIBOOT_INFO_VBE_INFO 0x00000800
#define MULTIBOOT_INFO_FRAMEBUFFER_INFO 0x00001000

#ifndef ASM_FILE

typedef unsigned char multiboot_uint8_t;
typedef unsigned short multiboot_uint16_t;
typedef unsigned int multiboot_uint32_t;
typedef unsigned long long multiboot_uint64_t;

struct multiboot_header
{
    /* Must be MULTIBOOT_MAGIC - see above. */
    multiboot_uint32_t magic;

    /* Feature flags. */
    multiboot_uint32_t flags;

    /* The above fields plus this one must equal 0 mod 2^32. */
    multiboot_uint32_t checksum;

    /* These are only valid if MULTIBOOT_AOUT_KLUDGE is set. */
    multiboot_uint32_t header_addr;
    multiboot_uint32_t load_addr;
    multiboot_uint32_t load_end_addr;
    multiboot_uint32_t bss_end_addr;
    multiboot_uint32_t entry_addr;

    /* These are only valid if MULTIBOOT_VIDEO_MODE is set. */
    multiboot_uint32_t mode_type;

```

```

multiboot_uint32_t width;
multiboot_uint32_t height;
multiboot_uint32_t depth;
};

/* The symbol table for a.out. */
struct multiboot_aout_symbol_table
{
    multiboot_uint32_t tabsize;
    multiboot_uint32_t strsize;
    multiboot_uint32_t addr;
    multiboot_uint32_t reserved;
};
typedef struct multiboot_aout_symbol_table multiboot_aout_symbol_table_t;

/* The section header table for ELF. */
struct multiboot_elf_section_header_table
{
    multiboot_uint32_t num;
    multiboot_uint32_t size;
    multiboot_uint32_t addr;
    multiboot_uint32_t shndx;
};
typedef struct multiboot_elf_section_header_table multiboot_elf_section_header_table_t;

struct multiboot_info
{
    /* Multiboot info version number */
    multiboot_uint32_t flags;

    /* Available memory from BIOS */
    multiboot_uint32_t mem_lower;
    multiboot_uint32_t mem_upper;

    /* "root" partition */
    multiboot_uint32_t boot_device;

    /* Kernel command line */
    multiboot_uint32_t cmdline;

    /* Boot-Module list */
    multiboot_uint32_t mods_count;
    multiboot_uint32_t mods_addr;

    union
    {
        multiboot_aout_symbol_table_t aout_sym;
        multiboot_elf_section_header_table_t elf_sec;
    } u;

    /* Memory Mapping buffer */
    multiboot_uint32_t mmap_length;
    multiboot_uint32_t mmap_addr;

    /* Drive Info buffer */
    multiboot_uint32_t drives_length;
    multiboot_uint32_t drives_addr;

    /* ROM configuration table */
    multiboot_uint32_t config_table;

    /* Boot Loader Name */
    multiboot_uint32_t boot_loader_name;

    /* APM table */
    multiboot_uint32_t apm_table;

    /* Video */
    multiboot_uint32_t vbe_control_info;
    multiboot_uint32_t vbe_mode_info;
    multiboot_uint16_t vbe_mode;
    multiboot_uint16_t vbe_interface_seg;
    multiboot_uint16_t vbe_interface_off;
    multiboot_uint16_t vbe_interface_len;

    multiboot_uint64_t framebuffer_addr;
    multiboot_uint32_t framebuffer_pitch;
    multiboot_uint32_t framebuffer_width;
    multiboot_uint32_t framebuffer_height;
    multiboot_uint8_t framebuffer_bpp;
#define MULTIBOOT_FRAMEBUFFER_TYPE_INDEXED 0
#define MULTIBOOT_FRAMEBUFFER_TYPE_RGB 1
#define MULTIBOOT_FRAMEBUFFER_TYPE_EGA_TEXT 2
    multiboot_uint8_t framebuffer_type;
    union
    {
        struct
        {
            multiboot_uint32_t framebuffer_palette_addr;
            multiboot_uint16_t framebuffer_palette_num_colors;
        };
        struct
        {
            multiboot_uint8_t framebuffer_red_field_position;
            multiboot_uint8_t framebuffer_red_mask_size;

```



```

        multiboot_uint8_t framebuffer_green_field_position;
        multiboot_uint8_t framebuffer_green_mask_size;
        multiboot_uint8_t framebuffer_blue_field_position;
        multiboot_uint8_t framebuffer_blue_mask_size;
    };
};
};
typedef struct multiboot_info multiboot_info_t;

struct multiboot_color
{
    multiboot_uint8_t red;
    multiboot_uint8_t green;
    multiboot_uint8_t blue;
};

struct multiboot_mmap_entry
{
    multiboot_uint32_t size;
    multiboot_uint64_t addr;
    multiboot_uint64_t len;
#define MULTIBOOT_MEMORY_AVAILABLE          1
#define MULTIBOOT_MEMORY_RESERVED          2
#define MULTIBOOT_MEMORY_ACPI_RECLAIMABLE  3
#define MULTIBOOT_MEMORY_NVS                4
#define MULTIBOOT_MEMORY_BADRAM             5
    multiboot_uint32_t type;
} __attribute__((packed));
typedef struct multiboot_mmap_entry multiboot_memory_map_t;

struct multiboot_mod_list
{
    /* the memory used goes from bytes 'mod_start' to 'mod_end-1' inclusive */
    multiboot_uint32_t mod_start;
    multiboot_uint32_t mod_end;

    /* Module command line */
    multiboot_uint32_t cmdline;

    /* padding to take it to 16 bytes (must be zero) */
    multiboot_uint32_t pad;
};
typedef struct multiboot_mod_list multiboot_module_t;

/* APM BIOS info. */
struct multiboot_apm_info
{
    multiboot_uint16_t version;
    multiboot_uint16_t cseg;
    multiboot_uint32_t offset;
    multiboot_uint16_t cseg_16;
    multiboot_uint16_t dseg;
    multiboot_uint16_t flags;
    multiboot_uint16_t cseg_len;
    multiboot_uint16_t cseg_16_len;
    multiboot_uint16_t dseg_len;
};

#endif /* !ASM_FILE */

#endif /* !MULTIBOOT_HEADER */

```

---

Next: [kernel.c](#), Previous: [multiboot.h](#), Up: [Example OS code](#)

### 4.3.2 boot.S

In the file `boot.S`:

```

/* boot.S - bootstrap the kernel */
/* Copyright (C) 1999, 2001, 2010 Free Software Foundation, Inc.
*
* This program is free software: you can redistribute it and/or modify
* it under the terms of the GNU General Public License as published by
* the Free Software Foundation, either version 3 of the License, or
* (at your option) any later version.
*
* This program is distributed in the hope that it will be useful,
* but WITHOUT ANY WARRANTY; without even the implied warranty of
* MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
* GNU General Public License for more details.
*
* You should have received a copy of the GNU General Public License
* along with this program. If not, see <http://www.gnu.org/licenses/>.
*/

#define ASM_FILE          1
#include <multiboot.h>

/* C symbol format. HAVE_ASM_USCORE is defined by configure. */
#ifdef HAVE_ASM_USCORE
# define EXT_C(sym)          _ ## sym

```

```
#else
# define EXT_C(sym)                sym
#endif

/* The size of our stack (16KB). */
#define STACK_SIZE                0x4000

/* The flags for the Multiboot header. */
#ifdef __ELF__
# define AOUT_KLUDGE 0
#else
# define AOUT_KLUDGE MULTIBOOT_AOUT_KLUDGE
#endif
#define MULTIBOOT_HEADER_FLAGS    MULTIBOOT_PAGE_ALIGN | MULTIBOOT_MEMORY_INFO | MULTIBOOT_VIDEO_MODE | AOUT_KLUDGE

    .text

    .globl start, _start
start:
_start:
    jmp     multiboot_entry

/* Align 32 bits boundary. */
    .align 4

/* Multiboot header. */
multiboot_header:
    /* magic */
    .long   MULTIBOOT_HEADER_MAGIC
    /* flags */
    .long   MULTIBOOT_HEADER_FLAGS
    /* checksum */
    .long   -(MULTIBOOT_HEADER_MAGIC + MULTIBOOT_HEADER_FLAGS)
#ifdef __ELF__
    /* header_addr */
    .long   multiboot_header
    /* load_addr */
    .long   _start
    /* load_end_addr */
    .long   _edata
    /* bss_end_addr */
    .long   _end
    /* entry_addr */
    .long   multiboot_entry
#else /* !__ELF__ */
    .long   0
    .long   0
    .long   0
    .long   0
    .long   0
#endif /* __ELF__ */
    .long   0
    .long   1024
    .long   768
    .long   32

multiboot_entry:
    /* Initialize the stack pointer. */
    movl    $(stack + STACK_SIZE), %esp

    /* Reset EFLAGS. */
    pushl    $0
    popf

    /* Push the pointer to the Multiboot information structure. */
    pushl    %ebx
    /* Push the magic value. */
    pushl    %eax

    /* Now enter the C main function... */
    call     EXT_C(cmain)

    /* Halt. */
    pushl    $halt_message
    call     EXT_C(sprintf)

loop:    hlt
        jmp     loop

halt_message:
    .asciz   "Halted."

/* Our stack area. */
    .comm    stack, STACK_SIZE
```

4.3.3 kernel.c

And, in the file kernel.c:

```

/* kernel.c - the C part of the kernel */
/* Copyright (C) 1999, 2010 Free Software Foundation, Inc.
*
* This program is free software: you can redistribute it and/or modify
* it under the terms of the GNU General Public License as published by
* the Free Software Foundation, either version 3 of the License, or
* (at your option) any later version.
*
* This program is distributed in the hope that it will be useful,
* but WITHOUT ANY WARRANTY; without even the implied warranty of
* MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
* GNU General Public License for more details.
*
* You should have received a copy of the GNU General Public License
* along with this program. If not, see <http://www.gnu.org/licenses/>.
*/

#include <multiboot.h>

/* Macros. */

/* Check if the bit BIT in FLAGS is set. */
#define CHECK_FLAG(flags, bit) ((flags) & (1 << (bit)))

/* Some screen stuff. */
/* The number of columns. */
#define COLUMNS 80
/* The number of lines. */
#define LINES 24
/* The attribute of an character. */
#define ATTRIBUTE 7
/* The video memory address. */
#define VIDEO 0xB8000

/* Variables. */
/* Save the X position. */
static int xpos;
/* Save the Y position. */
static int ypos;
/* Point to the video memory. */
static volatile unsigned char *video;

/* Forward declarations. */
void cmain (unsigned long magic, unsigned long addr);
static void cls (void);
static void itoa (char *buf, int base, int d);
static void putchar (int c);
void printf (const char *format, ...);

/* Check if MAGIC is valid and print the Multiboot information structure
pointed by ADDR. */
void
cmain (unsigned long magic, unsigned long addr)
{
    multiboot_info_t *mbi;

    /* Clear the screen. */
    cls ();

    /* Am I booted by a Multiboot-compliant boot loader? */
    if (magic != MULTIBOOT_BOOTLOADER_MAGIC)
    {
        printf ("Invalid magic number: 0x%x\n", (unsigned) magic);
        return;
    }

    /* Set MBI to the address of the Multiboot information structure. */
    mbi = (multiboot_info_t *) addr;

    /* Print out the flags. */
    printf ("flags = 0x%x\n", (unsigned) mbi->flags);

    /* Are mem_* valid? */
    if (CHECK_FLAG (mbi->flags, 0))
        printf ("mem_lower = %uKB, mem_upper = %uKB\n",
            (unsigned) mbi->mem_lower, (unsigned) mbi->mem_upper);

    /* Is boot_device valid? */
    if (CHECK_FLAG (mbi->flags, 1))
        printf ("boot_device = 0x%x\n", (unsigned) mbi->boot_device);

    /* Is the command line passed? */
    if (CHECK_FLAG (mbi->flags, 2))
        printf ("cmdline = %s\n", (char *) mbi->cmdline);

    /* Are mods_* valid? */
    if (CHECK_FLAG (mbi->flags, 3))
    {
        multiboot_module_t *mod;
        int i;

```



```

printf ("mods_count = %d, mods_addr = 0x%x\n",
        (int) mbi->mods_count, (int) mbi->mods_addr);
for (i = 0, mod = (multiboot_module_t *) mbi->mods_addr;
     i < mbi->mods_count;
     i++, mod++)
    printf (" mod_start = 0x%x, mod_end = 0x%x, cmdline = %s\n",
            (unsigned) mod->mod_start,
            (unsigned) mod->mod_end,
            (char *) mod->cmdline);
}

/* Bits 4 and 5 are mutually exclusive! */
if (CHECK_FLAG (mbi->flags, 4) && CHECK_FLAG (mbi->flags, 5))
{
    printf ("Both bits 4 and 5 are set.\n");
    return;
}

/* Is the symbol table of a.out valid? */
if (CHECK_FLAG (mbi->flags, 4))
{
    multiboot_aout_symbol_table_t *multiboot_aout_sym = &(mbi->u.aout_sym);

    printf ("multiboot_aout_symbol_table: tabsize = 0x%0x, "
            "strsize = 0x%x, addr = 0x%x\n",
            (unsigned) multiboot_aout_sym->tabsize,
            (unsigned) multiboot_aout_sym->strsize,
            (unsigned) multiboot_aout_sym->addr);
}

/* Is the section header table of ELF valid? */
if (CHECK_FLAG (mbi->flags, 5))
{
    multiboot_elf_section_header_table_t *multiboot_elf_sec = &(mbi->u.elf_sec);

    printf ("multiboot_elf_sec: num = %u, size = 0x%x, "
            " addr = 0x%x, shndx = 0x%x\n",
            (unsigned) multiboot_elf_sec->num, (unsigned) multiboot_elf_sec->size,
            (unsigned) multiboot_elf_sec->addr, (unsigned) multiboot_elf_sec->shndx);
}

/* Are mmap_* valid? */
if (CHECK_FLAG (mbi->flags, 6))
{
    multiboot_memory_map_t *mmap;

    printf ("mmap_addr = 0x%x, mmap_length = 0x%x\n",
            (unsigned) mbi->mmap_addr, (unsigned) mbi->mmap_length);
    for (mmap = (multiboot_memory_map_t *) mbi->mmap_addr;
         (unsigned long) mmap < mbi->mmap_addr + mbi->mmap_length;
         mmap = (multiboot_memory_map_t *) ((unsigned long) mmap
                                             + mmap->size + sizeof (mmap->size)))
        printf (" size = 0x%x, base_addr = 0x%x%08x, "
                " length = 0x%x%08x, type = 0x%x\n",
                (unsigned) mmap->size,
                (unsigned) (mmap->addr >> 32),
                (unsigned) (mmap->addr & 0xffffffff),
                (unsigned) (mmap->len >> 32),
                (unsigned) (mmap->len & 0xffffffff),
                (unsigned) mmap->type);
}

/* Draw diagonal blue line. */
if (CHECK_FLAG (mbi->flags, 12))
{
    multiboot_uint32_t color;
    unsigned i;
    void *fb = (void *) (unsigned long) mbi->framebuffer_addr;

    switch (mbi->framebuffer_type)
    {
        case MULTIBOOT_FRAMEBUFFER_TYPE_INDEXED:
        {
            unsigned best_distance, distance;
            struct multiboot_color *palette;

            palette = (struct multiboot_color *) mbi->framebuffer_palette_addr;

            color = 0;
            best_distance = 4*256*256;

            for (i = 0; i < mbi->framebuffer_palette_num_colors; i++)
            {
                distance = (0xff - palette[i].blue) * (0xff - palette[i].blue)
                    + palette[i].red * palette[i].red
                    + palette[i].green * palette[i].green;
                if (distance < best_distance)
                {
                    color = i;
                    best_distance = distance;
                }
            }
        }
        break;

        case MULTIBOOT_FRAMEBUFFER_TYPE_RGB:

```

```

        color = ((1 << mbi->framebuffer_blue_mask_size) - 1)
        << mbi->framebuffer_blue_field_position;
        break;

    case MULTIBOOT_FRAMEBUFFER_TYPE_EGA_TEXT:
        color = '\\\\' | 0x0100;
        break;

    default:
        color = 0xffffffff;
        break;
}
for (i = 0; i < mbi->framebuffer_width
    && i < mbi->framebuffer_height; i++)
{
    switch (mbi->framebuffer_bpp)
    {
        case 8:
        {
            multiboot_uint8_t *pixel = fb + mbi->framebuffer_pitch * i + i;
            *pixel = color;
        }
        break;
        case 15:
        case 16:
        {
            multiboot_uint16_t *pixel
                = fb + mbi->framebuffer_pitch * i + 2 * i;
            *pixel = color;
        }
        break;
        case 24:
        {
            multiboot_uint32_t *pixel
                = fb + mbi->framebuffer_pitch * i + 3 * i;
            *pixel = (color & 0xffffffff) | (*pixel & 0xff000000);
        }
        break;

        case 32:
        {
            multiboot_uint32_t *pixel
                = fb + mbi->framebuffer_pitch * i + 4 * i;
            *pixel = color;
        }
        break;
    }
}
}

/* Clear the screen and initialize VIDEO, XPOS and YPOS. */
static void
cls (void)
{
    int i;

    video = (unsigned char *) VIDEO;

    for (i = 0; i < COLUMNS * LINES * 2; i++)
        *(video + i) = 0;

    xpos = 0;
    ypos = 0;
}

/* Convert the integer D to a string and save the string in BUF. If
BASE is equal to 'd', interpret that D is decimal, and if BASE is
equal to 'x', interpret that D is hexadecimal. */
static void
itoa (char *buf, int base, int d)
{
    char *p = buf;
    char *p1, *p2;
    unsigned long ud = d;
    int divisor = 10;

    /* If %d is specified and D is minus, put '-' in the head. */
    if (base == 'd' && d < 0)
    {
        *p++ = '-';
        buf++;
        ud = -d;
    }
    else if (base == 'x')
        divisor = 16;

    /* Divide UD by DIVISOR until UD == 0. */
    do
    {
        int remainder = ud % divisor;

        *p++ = (remainder < 10) ? remainder + '0' : remainder + 'a' - 10;
    }
    while (ud /= divisor);

```

```

/* Terminate BUF. */
*p = 0;

/* Reverse BUF. */
p1 = buf;
p2 = p - 1;
while (p1 < p2)
{
    char tmp = *p1;
    *p1 = *p2;
    *p2 = tmp;
    p1++;
    p2--;
}

/* Put the character C on the screen. */
static void
putchar (int c)
{
    if (c == '\n' || c == '\r')
    {
        newline:
        xpos = 0;
        ypos++;
        if (ypos >= LINES)
            ypos = 0;
        return;
    }

    *(video + (xpos + ypos * COLUMNS) * 2) = c & 0xFF;
    *(video + (xpos + ypos * COLUMNS) * 2 + 1) = ATTRIBUTE;

    xpos++;
    if (xpos >= COLUMNS)
        goto newline;
}

/* Format a string and print it on the screen, just like the libc
function printf. */
void
printf (const char *format, ...)
{
    char **arg = (char **) &format;
    int c;
    char buf[20];

    arg++;

    while ((c = *format++) != 0)
    {
        if (c != '%')
            putchar (c);
        else
        {
            char *p, *p2;
            int pad0 = 0, pad = 0;

            c = *format++;
            if (c == '0')
            {
                pad0 = 1;
                c = *format++;
            }

            if (c >= '0' && c <= '9')
            {
                pad = c - '0';
                c = *format++;
            }

            switch (c)
            {
                case 'd':
                case 'u':
                case 'x':
                    itoa (buf, c, *((int *) arg++));
                    p = buf;
                    goto string;
                    break;

                case 's':
                    p = *arg++;
                    if (! p)
                        p = "(null)";

            string:
                for (p2 = p; *p2; p2++);
                for (; p2 < p + pad; p2++)
                    putchar (pad0 ? '0' : ' ');
                while (*p)
                    putchar (*p++);
                break;

            default:
                putchar (*((int *) arg++));
                break;
            }
        }
    }
}

```



```
    }
}
}
```

Previous: [kernel.c](#), Up: [Example OS code](#)

### 4.3.4 Other Multiboot kernels

Other useful information should be available in Multiboot kernels, such as GNU Mach and Fiasco <http://os.inf.tu-dresden.de/fiasco/>. And, it is worth mentioning the OSKit <http://www.cs.utah.edu/projects/flux/oskit/>, which provides a library supporting the specification.

Previous: [Example OS code](#), Up: [Examples](#)

## 4.4 Example boot loader code

The GNU GRUB (see [GRUB](#)) project is a Multiboot-compliant boot loader, supporting all required and many optional features present in this specification. A public release has not been made, but the test release is available from:

<ftp://alpha.gnu.org/gnu/grub>

See the webpage <http://www.gnu.org/software/grub/grub.html>, for more information.

Next: [Invoking mbchk](#), Previous: [Examples](#), Up: [Top](#)

# 5 The change log of this specification

0.7

- *Multiboot Standard* is renamed to *Multiboot Specification*.
- Graphics fields are added to Multiboot header.
- BIOS drive information, BIOS configuration table, the name of a boot loader, APM information, and graphics information are added to Multiboot information.
- Rewritten in Texinfo format.
- Rewritten, using more strict words.
- The maintainer changes to the GNU GRUB maintainer team [bug-grub@gnu.org](mailto:bug-grub@gnu.org), from Bryan Ford and Erich Stefan Boleyn.
- The byte order of the ‘boot\_device’ in Multiboot information is reversed. This was a mistake.
- The offset of the address fields were wrong.
- The format is adapted to a newer Texinfo, and the version number is specified more explicitly in the title.

0.6

- A few wording changes.
- Header checksum.
- Classification of machine state passed to an operating system.

0.5

- Name change.

0.4

- Major changes plus HTMLification.

Next: [Index](#), Previous: [History](#), Up: [Top](#)

# 6 Invoking mbchk

The program `mbchk` checks for the format of a Multiboot kernel. We recommend using this program before booting your own kernel.

`mbchk` accepts the following options:

- `--help`  
Print a summary of the command-line options and exit.
- `--version`  
Print the version number of GRUB and exit.
- `--quiet`  
Suppress all normal output.

Previous: [Invoking mbchk](#), Up: [Top](#)

# Index

# Table of Contents

- [Multiboot Specification](#)
- [1 Introduction to Multiboot Specification](#)
  - [1.1 The background of Multiboot Specification](#)
  - [1.2 The target architecture](#)
  - [1.3 The target operating systems](#)
  - [1.4 Boot sources](#)
  - [1.5 Configure an operating system at boot-time](#)
  - [1.6 How to make OS development easier](#)
  - [1.7 Boot modules](#)
- [2 The definitions of terms used through the specification](#)
- [3 The exact definitions of Multiboot Specification](#)
  - [3.1 OS image format](#)
    - [3.1.1 The layout of Multiboot header](#)
    - [3.1.2 The magic fields of Multiboot header](#)
    - [3.1.3 The address fields of Multiboot header](#)
    - [3.1.4 The graphics fields of Multiboot header](#)
  - [3.2 Machine state](#)
  - [3.3 Boot information format](#)
- [4 Examples](#)
  - [4.1 Notes on PC](#)
  - [4.2 BIOS device mapping techniques](#)
    - [4.2.1 Data comparison technique](#)
    - [4.2.2 I/O restriction technique](#)
  - [4.3 Example OS code](#)
    - [4.3.1 multiboot.h](#)
    - [4.3.2 boot.S](#)
    - [4.3.3 kernel.c](#)
    - [4.3.4 Other Multiboot kernels](#)
  - [4.4 Example boot loader code](#)
- [5 The change log of this specification](#)
- [6 Invoking mbchk](#)
- [Index](#)