

GUID Partition Table

jnlin

Topics

- Unified Extensible Firmware Interface (UEFI)
- Master Boot Record (MBR)
- GUID Partition Table (GPT)

Unified Extensible Firmware Interface

❑ Legacy BIOS limitations


- 16-bit processor mode
- 1 MB addressable space

❑ Advantages

- 32-bit/64-bit processor mode
- Ability to boot from larger disk with a GPT
- Flexible pre-OS environment, including network capability
- Modular design

❑ Compatibility Support Module (CSM)

- BIOS-MBR
- BIOS-GPT



Master Boot Record (1/2)

- ❑ The Master Boot Record (MBR) is the first 512 bytes of a storage device

Offset	Length	Contents
0	446 bytes	Boot code area
446	64 bytes	Partition tables, each has 16 bytes
510	2 bytes	Boot signature (0x55AA)
128	Total	

Master Boot Record (2/2)

❑ Drawbacks

- (4 primary partitions) or (3 primary + 1 extended partitions)
 - Arbitrary number of logical partitions within the extended partition
- The logical partition meta-data is stored in a linked-list structure
- One byte partition type codes which leads to many collisions
- Maximum addressable size is 2 TiB, i.e. any space beyond 2 TiB cannot be defined as a partition
 - MBR stores partition sector information using 32-bit LBA values
 - 512 bytes per sector
 - $2^{32} * 512 \text{ bytes} = 2 \text{ TiB}$

Booting Process

1. System initialization with firmware called BIOS
2. The BIOS looks for the bootloader on the MBR, then executes it
3. Bootloader reads the partition table
 - Conventional Windows/DOS MBR bootloader search for one active and primary partition
 - GRUB safely ignores this
4. Loading operating system

GUID Partition Table (1/9)

- ❑ GUID stands for Globally Unique Identifier
 - Ex: 3F2504E0-4F89-41D3-9A0C-0305E82C3301
- ❑ Part of the UEFI specification
- ❑ Solves some legacy problems with MBR but also may have compatibility issues
- ❑ Can be used also on BIOS system via a protective MBR

GUID Partition Table (2/9)

□ Advantages

- Filesystem-independent
- No partition type collision because of GUIDs
- 8 ZiB
 - GPT uses 64-bit LBA
 - 512 bytes per sector
 - $2^{64} * 512 \text{ bytes} = 8 \text{ ZiB}$
- Backup header and partition table at the end of the disk
- CRC32 checksums for header and partition table


LBA: Logical Block Address

GUID Partition Table (3/9)

❑ GPT Scheme

- LBA 0: Legacy MBR
- LBA 1: GPT header
- LBA 2~33: Partition entries
 - Up to 128 partitions
- LBA 34~: Partitions
- LBA -34~-1: Secondary GPT data

GUID Partition Table Scheme



GUID Partition Table (4/9)

❑ Legacy MBR (LBA 0)

- A single partition type of 0xEE
- For OSes cannot read GPT disks: Unknown type, no empty space
- For GPT-aware OSes: check the protective MBR

GUID Partition Table (5/9)

❑ GPT header (LBA 1)

Offset	Length	Contents
0	8 bytes	Signature (" EFI PART ", 45 46 49 20 50 41 52 54)
8	4 bytes	Revision (For GPT version 1.0 (through at least UEFI version 2.3.1), the value is 00 00 01 00)
12	4 bytes	Header size in little endian (in bytes, usually 5C 00 00 00 meaning 92 bytes)
16	4 bytes	CRC32 of header (0 to header size), with this field zeroed during calculation
20	4 bytes	Reserved; must be zero
24	8 bytes	Current LBA (location of this header copy)
32	8 bytes	Backup LBA (location of the other header copy)
40	8 bytes	First usable LBA for partitions (primary partition table last LBA + 1)
48	8 bytes	Last usable LBA (secondary partition table first LBA - 1)
56	16 bytes	Disk GUID (also referred as UUID on UNIXes)
72	8 bytes	Partition entries starting LBA (always 2 in primary copy)
80	4 bytes	Number of partition entries
84	4 bytes	Size of a partition entry (usually 128)
88	4 bytes	CRC32 of partition array
92	*	Reserved; must be zeroes for the rest of the block (420 bytes for a 512-byte LBA)

GUID Partition Table (6/9)

❑ GPT header (LBA 1)

- `dd if=/dev/ada0 bs=512 count=1 skip=1 | hd`

```
# dd if=/dev/ada0 bs=512 count=1 skip=1 | hd
00000000  45 46 49 20 50 41 52 54  00 00 01 00  5c 00 00 00  EFI PART ... \... |
00000010  ad 09 1d 1d 00 00 00 00  01 00 00 00  00 00 00 00  |.....|
00000020  ff ff 7f 02 00 00 00 00  22 00 00 00  00 00 00 00  |.....".....|
00000030  de ff 7f 02 00 00 00 00  65 67 3c f3  ea 40 e4 11  |.....eg<..@..|
00000040  a2 27 55 0b 19 3d b4 a4  02 00 00 00  00 00 00 00  |.'U..=.....|
00000050  80 00 00 00 80 00 00 00  82 f4 3d 77  00 00 00 00  |.....=w.....|
00000060  00 00 00 00 00 00 00 00  00 00 00 00  00 00 00 00  |.....|
*
00000200
```

GUID Partition Table (7/9)

❑ Partition entries (LBA 2)

Offset	Length	Contents
0	16 bytes	Partition type GUID
16	16 bytes	Unique partition GUID
32	8 bytes	First LBA (little-endian)
40	8 bytes	Last LBA (inclusive, usually odd)
48	8 bytes	Attribute flags (e.g. bit 60 denotes read-only)
56	72 bytes	Partition name (36 UTF-16LE code units)
	128 bytes	Total

GUID Partition Table (8/9)

❑ Partition type GUID

freebsd-boot	83BD6B9D-7F41-11DC-BE0B-001560B84F0F
freebsd	516E7CB4-6ECF-11D6-8FF8-00022D09712B
freebsd-swap	516E7CB5-6ECF-11D6-8FF8-00022D09712B
freebsd-ufs	516E7CB6-6ECF-11D6-8FF8-00022D09712B
freebsd-vinum	516E7CB8-6ECF-11D6-8FF8-00022D09712B
freebsd-zfs	516E7CBA-6ECF-11D6-8FF8-00022D09712B

GUID Partition Table (9/9)

❑ Partition entries (LBA 2)

- `dd if=/dev/ada0 bs=512 count=1 skip=2 | hd`

```
# dd if=/dev/ada0 bs=512 count=1 skip=2 | hd
00000000  9d 6b bd 83 41 7f dc 11  be 0b 00 15 60 b8 4f 0f  |.k..A.....`.O.|
00000010  0e 99 e2 03 eb 40 e4 11  a2 27 55 0b 19 3d b4 a4  |.....@...'U..=..|
00000020  22 00 00 00 00 00 00 00  a1 00 00 00 00 00 00 00  |".....|
00000030  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |.....|
*
00000080  b5 7c 6e 51 cf 6e d6 11  8f f8 00 02 2d 09 71 2b  |.|nQ.n.....-.q+|
00000090  98 66 a7 0f eb 40 e4 11  a2 27 55 0b 19 3d b4 a4  |.f...@...'U..=..|
000000a0  a2 00 00 00 00 00 00 00  a1 00 20 00 00 00 00 00  |.....|
000000b0  00 00 00 00 00 00 00 00  73 00 77 00 61 00 70 00  |.....s.w.a.p.|
000000c0  2d 00 30 00 00 00 00 00  00 00 00 00 00 00 00 00  |-.0.....|
000000d0  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |.....|
*
00000100  ba 7c 6e 51 cf 6e d6 11  8f f8 00 02 2d 09 71 2b  |.|nQ.n.....-.q+|
00000110  f6 11 10 1b eb 40 e4 11  a2 27 55 0b 19 3d b4 a4  |.....@...'U..=..|
00000120  a2 00 20 00 00 00 00 00  de ff 7f 02 00 00 00 00  |..|
00000130  00 00 00 00 00 00 00 00  7a 00 66 00 73 00 2d 00  |.....z.f.s.-|
00000140  30 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |0.....|
00000150  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |.....|
*
00000200
```

freebsd-boot

freebsd-swap

freebsd-zfs