

# GUID Partition Table

tsaimh (2022, CC BY-SA)

lwhsu (2021, CC BY-SA)

jnlin(2019-2021, CC BY-SA) 國立陽明交通大學資工系資訊中心


? (1996-2018)

# 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



Reference:  
<https://man.freebsd.org/uefi>

# 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)
	512 bytes	Total size: 446 + 64 + 2

Reference:

[https://en.wikipedia.org/wiki/Master\\_boot\\_record](https://en.wikipedia.org/wiki/Master_boot_record)

[https://wiki.osdev.org/MBR\\_\(x86\)](https://wiki.osdev.org/MBR_(x86))

# 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 initializing and self testing with the firmware called BIOS
2. BIOS loads the MBR of the boot device to memory (0000:7C00), then point CPU to start execute it (as 1st bootloader).
3. Bootloader reads the partition table and find the next boot program
  - Conventional Windows/DOS MBR bootloader searches for one active and primary partition
  - Different operating systems or boot manager have their own implementations
4. The final member in the boot chain loads the 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
- Also recognized on a BIOS system via the protective MBR (LBA 0)

LBA: Logical Block Address

Reference:

[https://en.wikipedia.org/wiki/Logical\\_block\\_addressing](https://en.wikipedia.org/wiki/Logical_block_addressing)

[https://en.wikipedia.org/wiki/GUID\\_Partition\\_Table](https://en.wikipedia.org/wiki/GUID_Partition_Table)

# 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

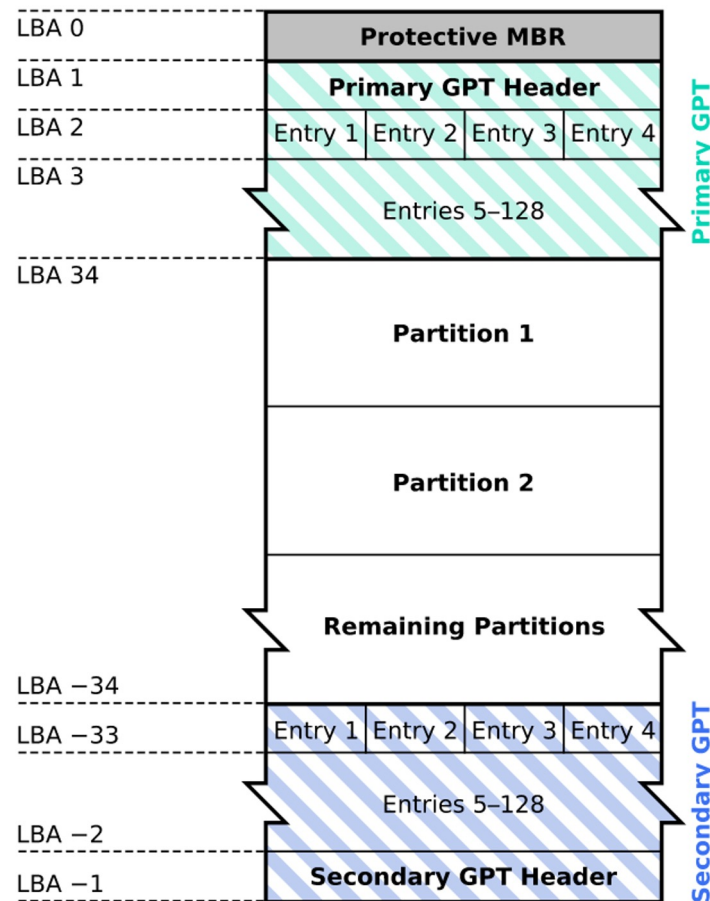


# GUID Partition Table (3/9)

- GPT Scheme

- LBA 0: Protective 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
    - Refuse to take actions unless further instructions (safety reasons)
  - For GPT-aware OSes: check the protective MBR
    - Check if the enclosed partition type is single 0xEE type
    - If not, refuse to take action

# 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
1+0 records in
1+0 records out
512 bytes transferred in 0.004644 secs (110259 bytes/sec)
00000000 45 46 49 20 50 41 52 54 00 00 01 00 5c 00 00 00 |EFI PART....\...|
00000010 b4 89 4d 11 00 00 00 00 01 00 00 00 00 00 00 00 |..M.....|
00000020 ff ff ff 01 00 00 00 00 28 00 00 00 00 00 00 00 |.....(.....|
00000030 d7 ff ff 01 00 00 00 00 fa 2e 89 f8 d5 c6 ea 11 |.....|
00000040 ad c5 08 00 27 9c b4 87 02 00 00 00 00 00 00 00 |....'.....|
00000050 80 00 00 00 80 00 00 00 0f 3e 88 1f 00 00 00 00 |.....>.....|
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	<b>Partition type</b> GUID
16	16 bytes	Unique partition GUID
32	8 bytes	<b>First LBA</b> (little-endian)
40	8 bytes	<b>Last LBA</b> (inclusive, usually odd)
48	8 bytes	Attribute flags (e.g. bit 60 denotes read-only)
56	72 bytes	<b>Partition name</b> (36 UTF-16LE code units)
	128 bytes	Total

# GUID Partition Table (8/9)

- Partition type GUID

efi	C12A7328-F81F-11D2-BA4B-00A0C93EC93B
freebsd-boot	83BD6B9D-7F41-11DC-BE0B-001560B84F0F
freebsd-swap	516E7CB5-6ECF-11D6-8FF8-00022D09712B
freebsd-ufs	516E7CB6-6ECF-11D6-8FF8-00022D09712B
freebsd-zfs	516E7CBA-6ECF-11D6-8FF8-00022D09712B

Reference:

gpart(8) <https://man.freebsd.org/gpart/8>

# GUID Partition Table (9/9)

- Partition entries (LBA 2)

```
# dd if=/dev/ada0 bs=512 count=1 skip=2 | hd
1+0 records in
1+0 records out
512 bytes transferred in 0.000425 secs (1205747 bytes/sec)
```

```
00000000  9d 6b bd 83 41 7f dc 11  be 0b 00 15 60 b8 4f 0f  |.k..A.....`.0.|
00000010  d1 0e 8a f8 d5 c6 ea 11  ad c5 08 00 27 9c b4 87  |.....'...|
00000020  28 00 00 00 00 00 00 00  27 04 00 00 00 00 00 00  |(.....'.....|
00000030  00 00 00 00 00 00 00 00  67 00 70 00 74 00 62 00  |.....g.p.t.b.|
00000040  6f 00 6f 00 74 00 30 00  00 00 00 00 00 00 00 00  |o.o.t.0.....|
00000050  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |.....|
```

**freebsd-boot**

```
*
00000080  b5 7c 6e 51 cf 6e d6 11  8f f8 00 02 2d 09 71 2b  |.|nQ.n.....-.q+|
00000090  f8 60 8f f8 d5 c6 ea 11  ad c5 08 00 27 9c b4 87  |.`.....'...|
000000a0  28 04 00 00 00 00 00 00  27 04 40 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  30 00 00 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  |.....|
```

**freebsd-swap**

```
*
00000100  ba 7c 6e 51 cf 6e d6 11  8f f8 00 02 2d 09 71 2b  |.|nQ.n.....-.q+|
00000110  ca 38 94 f8 d5 c6 ea 11  ad c5 08 00 27 9c b4 87  |.8.....'...|
00000120  28 04 40 00 00 00 00 00  d7 ff ff 01 00 00 00 00  |(.@.....|
00000130  00 00 00 00 00 00 00 00  7a 00 66 00 73 00 30 00  |.....z.f.s.0.|
00000140  00 00 00 00 00 00 00 00  00 00 00 00 00 00 00 00  |.....|
```

**freebsd-zfs**

```
*
00000200
```

# gpart(8)

```
$ ls /dev/nvd0*
/dev/nvd0      /dev/nvd0p1  /dev/nvd0p2  /dev/nvd0p3
$ gpart show
=>          40  976773088  nvd0  GPT  (466G)
           40    532480    1  efi  (260M)
           532520    2008    - free -  (1.0M)
           534528   4194304    2  freebsd-swap  (2.0G)
           4728832  972044288    3  freebsd-zfs  (464G)
           976773120      8    - free -  (4.0K)
```