

Windows GPT FAQ

(Microsoft Corporation)

This article provides answers to frequently asked questions about the GUID Partition Table (GPT) as implemented under Microsoft Windows XP x64 edition, Windows Server 2003 (64-bit), Windows Server 2003 SP1 (all versions), and the upcoming Windows Vista and Windows Server 2008 operating systems.

About the GPT

1. What is a GPT disk?

The GUID Partition Table (GPT) was introduced as part of the Extensible Firmware Interface (EFI) initiative. GPT provides a more flexible mechanism for partitioning disks than the older Master Boot Record (MBR) partitioning scheme that has been common to PCs.

A partition is a contiguous space of storage on a physical or logical disk that functions as though it were a physically separate disk. Partitions are visible to the system firmware and the installed operating systems. Access to a partition is controlled by the system firmware before the system boots the operating system, and then by the operating system after it starts.

2. What is wrong with MBR partitioning?

MBR disks support only four partition table entries. If more partitions are wanted, a secondary structure known as an extended partition is necessary. Extended partitions can then be subdivided into one or more logical disks. By convention, Windows creates MBR disk partitions and logical drives on cylinder boundaries based on the reported geometry, although this information no longer has any relationship to the physical characteristics of the hardware (disk driver or RAID controller). For Windows Vista and Windows Server 2008, more logical boundaries will be selected when the hardware provides better hints at the true cache or physical alignment. Because this partition information is stored on the drive itself, the operating system is not dependent on the alignment. MBR partitioning rules are complex and poorly specified. For example, does cylinder alignment mean that each partition must be at least one cylinder in length? An MBR partition is identified by a two-byte field, and coordination is necessary to avoid collision. IBM originally provided that coordination; today there is no single authoritative list of partition identifiers.

Another common practice is to use partitioned or "hidden" sectors to hold specific information. That practice is undocumented and results in severe system problems that are difficult to debug. Over the years, vendor-specific implementations and tools have been released to the public, making support difficult.

3. Why do we need GPT?

GPT disks can grow to a very large size. The number of partitions on a GPT disk is not constrained by temporary schemes such as container partitions as defined by the MBR Extended Boot Record (EBR).

The GPT disk partition format is well defined and fully self-identifying. Data critical to platform operation is located in partitions and not in unpartitioned or "hidden" sectors. GPT disks use primary and backup partition tables for redundancy and CRC32 fields for improved partition data structure integrity. The GPT partition format uses version number and size fields for future expansion.

Each GPT partition has a unique identification GUID and a partition content type, so no coordination is necessary to prevent partition identifier collision. Each GPT partition has a 36-character Unicode name, which means that any software can present a human-readable name for the partition without any additional understanding of the partition.

Windows GPT FAQ (Microsoft Corporation)

4. **Where can I find the specification for GPT disk partitioning?**

Chapter 16 of the Extensible Firmware Interface (EFI) specification defines the GPT format. This document is available at <http://www.Intel.com/technology/efi/>.

5. **Is EFI required for a GPT disk?**

No. GPT disks are self-identifying. All the information needed to interpret the partitioning scheme of a GPT disk is completely contained in structures in specified locations on the physical media.

6. **How big can a GPT disk be?**

In theory, a GPT disk can be up to 2^{64} logical blocks in length. Logical blocks are commonly 512 bytes in size.

The maximum partition (and disk) size is a function of the operating system version. Windows XP and the original release of Windows Server 2003 have a limit of 2TB per physical disk, including all partitions. For Windows Server 2003 SP1 Windows XP x64 edition, and later versions, the maximum raw partition of 18 exabytes can be supported. (Windows file systems currently are limited to 256 terabytes each.)

7. **How many partitions can a GPT disk have?**

The specification allows an almost unlimited number of partitions, but the Windows implementation restricts this to 128 partitions. The number of partitions is limited by the amount of space reserved for partition entries in the GPT.

8. **Can a disk be both GPT and MBR?**

No. However, all GPT disks contain a Protective MBR.

9. **What is a Protective MBR?**

The Protective MBR, beginning in sector 0, precedes the GPT partition table on the disk. The MBR contains one type 0xEE partition that spans the disk.

10. **Why does the GPT have a Protective MBR?**

The Protective MBR protects GPT disks from previously released MBR disk tools such as Microsoft MS-DOS FDISK or Microsoft Windows NT Disk Administrator. These tools are not aware of GPT and do not know how to properly access a GPT disk. Legacy software that does not know about GPT interprets only the Protected MBR when it accesses a GPT disk. These tools will view a GPT disk as having a single encompassing (possibly unrecognized) partition by interpreting the Protected MBR, rather than mistaking the disk for one that is unpartitioned.

11. **Why would a GPT-partitioned disk appear to have an MBR on it?**

If this occurred, you must have used an MBR-only-aware disk tool to access the GPT disk. See answers #8 through #10, earlier.

12. **If the disk is larger than the maximum size an MBR can report, will the entire disk contents be protected?**

The EE partition in the Protective MBR is specified to be the maximum size allowable in an MBR.

Windows XP Disk Support

13. **Can Windows XP x64 read, write, and boot from GPT disks?**

Windows XP x64 edition can use GPT disks for data only. Only Windows for Itanium-based systems can boot from GPT partitions.

Windows GPT FAQ (Microsoft Corporation)

14. Can the 32-bit version of Windows XP read, write, and boot from GPT disks?

No. The 32-bit version will see only the Protective MBR. The EE partition will not be mounted or otherwise exposed to application software.

15. Can the 32-bit versions of Windows Server 2003 read, write, and boot from GPT disks?

All versions of Windows 2003 since Server Pack 1 can use GPT partitioned disks for data. Booting is only supported for Itanium-based systems.

16. Can Windows Vista and Windows Server 2008 read, write, and boot from GPT disks?

Yes, all versions can use GPT partitioned disks for data. Booting is only supported for EFI-based systems.

17. Can Windows 2000, Windows NT 4, or Windows 95/98 read, write, and boot from GPT?

No. Again, legacy software will see only the Protective MBR.

18. What about mixing and matching GPT and MBR disks on the same system?

GPT and MBR disks can be mixed on systems that support GPT, as described earlier the following restrictions apply:

- Systems that support EFI require that the boot partition must reside on a GPT disk. Other hard disks can be either MBR or GPT.
- Both MBR and GPT disks can be present in a single dynamic disk group. Volume sets can span both MBR and GPT disks. However, the MBR cylinder alignment restriction might cause some difficulties mirroring or striping MBR and GPT disks.

19. What about removable media?

Removable media must be MBR or "superfloppy."

20. What is a superfloppy?

Removable media without either GPT or MBR formatting is considered a "superfloppy." The entire media is treated as a single partition.

The media manufacturer performs any MBR partitioning of removable media. If the media does have an MBR, only one partition is supported. There is little user-discernible difference between MBR-partitioned media and superfloppies.

Examples of removable media include floppy disk drives, JAZ disk cartridges, magneto-optical media, DVD-ROM, and CD-ROM. Hard disk drives on external buses such as SCSI or IEEE 1394 are not considered removable.

21. What is the default behavior of Windows XP 64-Bit Edition Version 2003 when partitioning media?

For Windows XP 64-Bit Edition Version 2003 only (for Itanium-based systems), fixed disks are partitioned using GPT partitioning. GPT disks can be converted to MBR disks only if all existing partitioning is first deleted, with associated loss of data.

22. What is the default behavior of the 32-bit version of Windows XP, Windows Server 2003 and Windows XP x64 when partitioning media?

Only MBR disks can be used. MBR disks cannot be converted to GPT disks.

Windows GPT FAQ (Microsoft Corporation)

23. How can a drive letter in the operating system be mapped to a partition in EFI firmware?

There is no inherent mapping between drive letter and partition that can be used to determine one from the other. A basic data partition must be identified by its partition GUID.

24. How can an ESP partition be created?

ESP partitions can be created using the EFI firmware utility Diskpart.efi or the Windows command line utility Diskpart.exe, or they can be created programmatically by using IOCTL_DISK_SET_DRIVE_LAYOUT.

25. What can be changed on a partition?

You should not change any partition header entry directly. Do not use disk tools or utilities to make alterations or changes.

26. What partitioning does Windows XP support on detachable disks?

Detachables disks are commonly expected to migrate between machines or simply to be unavailable to the operating system at times. Examples of detachable disks are IEEE 1394 disks, which can be easily disconnected by the end user, or Microsoft Cluster Services (MSCS) shared disks, which move between nodes in a cluster.

Windows XP supports only MBR partitioning on detachable disks. For more about removable media, see questions 19-21. Later versions of Windows support GPT partitions on detachable disks.

Windows GPT Required Partitions

EFI System Partition

27. What is the EFI System Partition (ESP)?

The ESP contains the NTLDR, HAL, Boot.txt, and other files necessary to boot the system, such as drivers. The Partition GUID defines the ESP:

```
DEFINE_GUID (PARTITION_SYSTEM_GUID, 0xC12A7328L, 0xF81F,  
0x11D2, 0xBA, 0x4B, 0x00, 0xA0, 0xC9, 0x3E, 0xC9, 0x3B)
```

28. Do only GPT Disks have ESPs?

No, MBR disks can also have ESPs. EFI specifies booting from either GPT or MBR. The ESP on an MBR disk is identified by partition type 0xEF. However, Windows does not support booting EFI from MBR disks or 0xEF partitions.

29. How big is the ESP?

The ESP is approximately 100 MB.

30. Can there be two ESPs on a single disk?

Such a configuration should not be created and will not be supported.

31. What about two ESPs on two different disks?

ESP partitions can be replicated for high-availability configurations. Replication must be done manually and the contents must be synchronized manually when using software volumes. Hardware vendors may provide additional solutions for high availability. ESP partitions cannot be mirrored.

32. What does Microsoft place in the ESP?

Windows GPT FAQ (Microsoft Corporation)

Microsoft places the HAL, loader, and other files necessary to boot the operating system in the ESP.

33. Where should the ESP be placed on the disk?

The ESP should be first on the disk. While there is no architectural requirement, there are numerous reasons why it is beneficial to place the ESP first. The primary reason for this is that it is impossible to span volumes when the ESP is logically between the two data partitions you are attempting to span.

34. What should a system or device manufacturer place in the ESP?

The ESP should only include files required for booting an operating system, platform tools that run before operating system boot, or files that must be accessed before operating system boot, for example in performing pre-boot system maintenance. Other value-add files or diagnostics used while the operating system is running should not be placed in the ESP. It is important to note that the space in the ESP is a limited system resource; its primary purpose is to provide storage for the files necessary to boot the operating system.

35. Where should a system manufacturer place files such as platform diagnostics or other value-added files?

The preferred option is for system manufacturers to place value-add contents in an OEM-specific partition. Just like MBR OEM partitions, the contents of GPT OEM (or other unrecognized) partitions are not exposed (given drive letters or returned in volume lists). Users are warned that deleting the partition can cause the system to fail to operate. An OEM-specific partition should be placed before the MSR and after any ESP on the disk. Although not architectural, this placement has the same benefits as placing the ESP first. For example, it is also impossible to span volumes when an OEM-specific partition is logically between the two data partitions you are attempting to span.

Placement in the ESP is an option for applications or files that execute in the pre-operating system boot environment. However, the ESP is architecturally shared space and represents a limited resource. Consuming space in the ESP should be considered carefully. Files that are not relevant to the pre-operating system boot environment should not be placed in the ESP.

Microsoft Reserved Partition

36. What is a Microsoft Reserved Partition (MSR)?

The Microsoft Reserved Partition (MSR) reserves space on each disk drive for subsequent use by operating system software. GPT disks do not allow hidden sectors. Software components that formerly used hidden sectors now allocate portions of the MSR for component-specific partitions. For example, converting a basic disk to a dynamic disk causes the MSR on that disk to be reduced in size and a newly created partition holds the dynamic disk database. The MSR has the Partition GUID:

```
DEFINE_GUID (PARTITION_MSFT_RESERVED_GUID, 0xE3C9E316L, 0x0B5C,  
0x4DB8, 0x81, 0x7D, 0xF9, 0x2D, 0xF0, 0x02, 0x15, 0xAE)
```

37. What disks require an MSR?

Every GPT disk *must* contain an MSR. The order of partitions on the disk should be ESP (if any), OEM (if any), and MSR followed by primary data partition(s). It is particularly important that the MSR be created before other primary data partitions.

Windows GPT FAQ (Microsoft Corporation)

38. Who creates the MSR?

The MSR must be created when disk-partitioning information is first written to the drive. If the manufacturer partitions the disk, the manufacturer must create the MSR at the same time. If Windows partitions the disk during setup, creates the MSR.

39. Why must the MSR be created when the disk is first partitioned?

After the disk is partitioned, there will be no free space left to create an MSR.

40. How big is the MSR?

When initially created, the size of the MSR depends on the size of the disk drive:

- On drives less than 16GB in size, the MSR is 32 MB.
- On drives greater than or equal two 16 GB, the MSR is 128 MB.

As the MSR is divided into other partitions, it becomes smaller.

Windows GPT ESP Implementation

41. What partitions are required by Windows XP?

For EFI systems, each bootable drive must contain an ESP, an MSR, and at least one basic data partition that contains the operating system. Each data drive must contain at least an MSR and one basic data partition.

All basic data partitions on the drive should be contiguous. As noted above, placing an OEM-specific or other unrecognized partition between data partitions imposes limitations on later volume spanning.

42. What is a basic data partition?

Basic data partitions correspond to primary MBR partitions 0x6 (FAT), 0x7 (NTFS), or 0xB (FAT32). Each basic partition can be mounted using a drive letter or mount point, other volume device object, or both. Each basic data partition is represented in Windows as a volume device object, and optionally as a mount point or a drive letter.

43. How is a basic data partition identified?

It has the following partition type GUID:

```
DEFINE_GUID (PARTITION_BASIC_DATA_GUID, 0xEBD0A0A2L, 0xB9E5,  
0x4433, 0x87, 0xC0, 0x68, 0xB6, 0xB7, 0x26, 0x99, 0xC7);
```

44. Will end users see the ESP, MSR, and OEM-specific partitions?

The user won't see these partitions exposed in Windows Explorer, nor is any recognized file system exposed to legacy programs such as Context Indexing. The ESP, OEM-specific, and other unrecognized partitions will be visible only in the Disk Management MMC snap-in since they will not have a recognizable file system.

45. What partitions are mounted by default by Windows?

Windows XP exposes only basic data partitions. Other partitions with FAT file systems may be mounted, but not exposed only programmatically. Only basic data partitions are assigned drive letters or mount points.

Windows GPT FAQ (Microsoft Corporation)

The ESP FAT file system is mounted, but not exposed. This allows programs running under Windows to update the contents of the ESP. Assigning a drive letter to the ESP using 'mountvol /s' will allow access to the partition. Access to the ESP requires admin privilege.

Although the MSR, and any partitions created from the MSR, could have recognizable file systems, none are exposed.

Any OEM-specific partitions or partitions associated with other operating systems are not recognized by Windows. Unrecognized partitions with recognizable file systems are treated like the ESP. They will be mounted, but not exposed. Unlike MBR disks, there is no practical difference between OEM-specific partitions and other operating system partitions; all are "unrecognized."

46. **How can the user see the ESP, OEM, and other unrecognized partitions?**

The user can use disk management tools such as the Disk Management utility or the diskpart.exe Windows command line. The MSR and any partitions created from the MSR are only visible from the command line.

47. **What about dynamic disks?**

Dynamic disks use two different GPT partitions:

A data container partition corresponding to the MBR partition 0x42, with the following GUID: DEFINE_GUID (PARTITION_LDM_DATA_GUID, 0xAF9B60A0L, 0x1431, 0x4F62, 0xBC, 0x68, 0x33, 0x11, 0x71, 0x4A, 0x69, 0xAD);

A partition to contain the dynamic configuration database, with the following GUID: DEFINE_GUID(PARTITION_LDM_METADATA_GUID, 0x5808C8AAL, 0x7E8F, 0x42E0, 0x85, 0xD2, 0xE1, 0xE9, 0x04, 0x34, 0xCF, 0xB3);

Volumes are created in the data container and mounted by default. Again, this is exactly the same as the contents of 0x42 MBR partitions.

48. **What happens when a basic disk is converted to dynamic?**

For a drive to be eligible for conversion to dynamic, all basic data partitions on the drive must be contiguous. If other unrecognized partitions separate basic data partitions, the disk cannot be converted. This is one of the reasons that the MSR must be created before any basic data partitions.

The first step in conversion is to separate a portion of the MSR to create the configuration database partition. All non-bootable basic partitions are then combined into a single data container partition. Boot partitions are retained as separate data container partitions. This is analogous to conversion of primary partitions.

Windows XP and later versions of the Windows operating system differs from Windows 2000 in that basic and extended partitions are preferentially converted to a single 0x42 partition, rather than being retained as multiple distinct 0x42 partitions as on Windows 2000.

49. **Can a system contain a mix of GPT and MBR dynamic disks?**

Yes. See answer #18 earlier.

50. **How can a specific partition be mounted?**

Windows GPT FAQ (Microsoft Corporation)

You can access the GPT disk partitions of different types using the tools listed in the following table:

	Windows	Firmware
Diskpart.efi Disk Partition Tool		ESP MSR Data
Diskpart.exe Disk Partition Tool	ESP MSR Data	
Diskmgmt.msc Logical Disk Manager	ESP Data	
Explorer.exe File Explorer	Data	

You can also develop your own tools, using the Microsoft Platform SDK APIs, to access the GPT disk partitions at their primitive levels.

51. How are GPT disks managed in Windows ?

GPT and MBR disks are managed the same way. Disks can be formatted as GPT or MBR using the Diskpart.exe command-line utility or by using the Disk Administrator snap-in. Volumes can be created on both GPT and MBR disks, and both kinds of disks can be mixed in the same dynamic disk group.

52. What about FTdisk sets?

There is no FTdisk set support on Windows XP or later for MBR or GPT disks. The only support for logical volumes is through dynamic disks.

53. Can a disk be converted from GPT to MBR, and vice versa?

Yes, but only if the disk contains no partitions or volumes. Any data on the disk will be destroyed.

54. What file systems are supported on GPT disks?

NTFS is recommended on all basic data partitions and all dynamic volumes. Windows Setup and the Disk Management snap-in offer only NTFS. To circumvent that, the partition or volume must be formatted explicitly via the Format command-line tool.

Manipulating GPT Disks and Their Contents

55. Is it possible to make a sector-by-sector copy of a GPT disk?

No. The Disk and Partition GUIDs will no longer be unique. This must never happen. You can make a sector-by-sector copy of the contents of ESP or basic data partitions.

56. Is there any way to copy a whole GPT disk using the OPK imaging tools?

Yes; however, there are some key caveats. The OEM Preinstallation Kit (OPK) initializes the Disk and Partition GUIDs to zero. On first boot of Windows, the operating system generates unique GUIDs. The OPK only supports generation of ESP, MSR, and basic data partitions.

If an application has recorded any Disk or Partition GUIDs it may break. Any applications, drivers, utilities, or firmware implementations supplied by system manufacturers or application vendors that rely on GUIDs should be capable of handling GUIDs that change from the OPK initialization values to those generated by the operating system.

57. What is the Diskpart.efi MAKE command?

Windows GPT FAQ (Microsoft Corporation)

It is a way for OEMs to simplify operating system pre-installation and system recovery. This command can easily be extended to create a "default" disk configuration for the platform. For example, the system manufacturer could extend the MAKE command to automatically partition the boot drive with an ESP, MSR, an OEM-specific partition, and one basic data partition.

For example, consider a possible disk configuration called BOOT_DISK. In the event of business failure recovery, MAKE BOOT_DISK would allow the customer to completely repartition a boot disk to the original factory defaults.

58. **What happens if a duplicate Disk or Partition GUID is detected?**

Windows will generate new GUIDs for any duplicate Disk GUID, MSR Partition GUID, or MSR basic data GUID upon detection. This is similar to the duplicate MBR signature handling in Windows 2000. Duplicate GUIDs on a dynamic container or database partition cause unpredictable results.