

RELSOFT TECHNOLOGIES

NTFS On-Disk Structures

Visual Basic NTFS
Programmer's
Guide

NTFS ON-DISK STRUCTURES

Visual Basic NTFS Programmer's Guide

© 2004 Alex Ionescu
Relsoft Technologies
<http://www.relsoft.net>
All Rights Reserved

1. BASIC CONCEPTS	5
1.1 INTRODUCTION	5
1.2 NTFS TERMINOLOGY	5
1.3 GENERAL TERMINOLOGY	6
1.4 NTFS VERSIONS	7
2. THE BOOT RECORD	7
2.1 DEFINITION	7
2.2 STRUCTURE	7
3. METAFILES	10
3.1 INTRODUCTION	10
3.1.1 <i>Description</i>	10
3.1.2 <i>Listing (NTFS 3.0+)</i>	10
3.2 \$MFT	11
3.2.1 <i>Description</i>	11
3.2.2 <i>Structure</i>	11
3.3 \$MFTMIRR	12
3.3.1 <i>Description</i>	12
3.3.2 <i>Structure</i>	13
3.4 \$LOGFILE	13
3.4.1 <i>Description</i>	13
3.4.2 <i>Structure</i>	14
3.5 \$VOLUME	14
3.5.1 <i>Description</i>	14
3.5.2 <i>Structure</i>	14
3.6 \$ATTRDEF	15
3.6.1 <i>Description</i>	15
3.6.2 <i>Structure</i>	15
3.7 . (DOT)	16
3.7.1 <i>Description</i>	16
3.7.2 <i>Structure</i>	17
3.8 \$BITMAP	17
3.8.1 <i>Description</i>	17
3.8.2 <i>Structure</i>	17
3.9 \$BOOT	18
3.9.1 <i>Description</i>	18
3.9.2 <i>Structure</i>	18
3.10 \$BADCLUS	18
3.10.1 <i>Description</i>	18
3.10.2 <i>Structure</i>	18
3.11 \$SECURE	19
3.11.1 <i>Description</i>	19
3.11.2 <i>Structure</i>	20
3.12 \$UPCASE	21
3.12.1 <i>Description</i>	21
3.12.2 <i>Structure</i>	22
3.13 \$EXTEND	22
3.13.1 <i>Description</i>	22
3.15 \$OBJID	22
3.15.1 <i>Description</i>	22
3.15.2 <i>Structure</i>	23

3.16 \$QUOTA.....	23
3.16.1 <i>Description</i>	23
3.16.2 <i>Structure</i>	24
3.17.1 <i>Description</i>	26
3.17.2 <i>Structure</i>	26
3.18 \$USNJRNL.....	27
3.18.1 <i>Description</i>	27
3.18.2 <i>Structure</i>	27
4. ATTRIBUTES.....	29
4.1 INTRODUCTION.....	29
4.1.1 <i>Definition</i>	29
4.1.2 <i>Listing (NTFS 3.0+)</i>	29
4.2 TYPES OF ATTRIBUTES.....	31
4.2.1 <i>Attribute Definition</i>	31
4.2.2 <i>Attribute Structure</i>	31
4.2.3 <i>Nonresident Attribute Definition</i>	32
4.2.4 <i>Nonresident Attribute Structure</i>	32
4.2.5 <i>Resident Attribute Definition</i>	33
4.2.6 <i>Resident Attribute Structure</i>	33
4.2.7 <i>Named and Unnamed Attributes</i>	34
4.3 \$STANDARD_INFORMATION.....	34
4.3.1 <i>Description</i>	34
4.3.2 <i>Structure</i>	35
4.4 \$ATTRIBUTE_LIST.....	36
4.4.1 <i>Description</i>	36
4.4.2 <i>Structure</i>	36
4.5 \$FILE_NAME.....	37
4.5.1 <i>Description</i>	37
4.6 \$OBJECT_ID.....	38
4.6.1 <i>Description</i>	38
4.6.2 <i>Structure</i>	38
4.7 \$SECURITY_DESCRIPTOR.....	39
4.7.1 <i>Description</i>	39
4.7.2 <i>Structure</i>	39
4.8 \$VOLUME_NAME.....	41
4.8.1 <i>Description</i>	41
4.9 \$VOLUME_INFORMATION.....	41
4.9.1 <i>Description</i>	41
4.9.2 <i>Structure</i>	41
4.10 \$DATA.....	42
4.10.1 <i>Description</i>	42
4.10.2 <i>Structure</i>	43
4.11 \$INDEX_ROOT.....	43
4.11.1 <i>Description</i>	43
4.11.2 <i>Structure</i>	44
4.12 \$INDEX_ALLOCATION.....	45
4.12.1 <i>Description</i>	45
4.12.2 <i>Structure</i>	46
4.13 \$BITMAP.....	47
4.13.1 <i>Description</i>	47
4.13.2 <i>Structure</i>	47
4.14 \$REPARSE_POINT.....	47

4.14.1 <i>Description</i>	47
4.14.2 <i>Structure</i>	48
4.15 \$EA_INFORMATION	49
4.15.1 <i>Description</i>	49
4.15.2 <i>Structure</i>	50
4.16 \$EA	50
4.16.1 <i>Description</i>	50
4.16.2 <i>Structure</i>	50
4.17 \$LOGGED_UTILITY_STREAM	51
4.17.1 <i>Description</i>	51
4.17.2 <i>Structure</i>	51
5.0 ADVANCED CONCEPTS	52
5.1 VCNS AND LCNS	52
5.2 DATA RUNS	52
5.2.1 <i>Definition</i>	52
5.2.2 <i>Structure</i>	53
5.3 SECURITY CONCEPTS	58
5.3.1 <i>SIDs</i>	58
5.3.2 <i>ACLs</i>	61
5.3.3 <i>ACEs</i>	61
5.4 INDEXES	61
5.5 SPARSE FILES	61
5.6 ENCRYPTION	61
5.7 COMPRESSION	61
5.8 USNS	61

1. Basic Concepts

1.1 Introduction

This article, based heavily on the NTFS Linux project, aims to provide the VB coder with an intricate knowledge of the filesystem, which could be used for custom-made applications, defragmenters, undelete utilities, or simply system information and curiosity. It also aims to be a bit more user friendly, since the original NTFS Linux Documentation, while highly technical and accurate, is not too easy to understand for the average programmer, much less use in VB. In addendum, a bas and tlb module containing all the structures is presented, and these structures will be presented along the text. I recommend you print this article out, as it is several pages long, and you will probably end up referencing it often.

1.2 NTFS Terminology

NTFS is rather an overwhelming filesystem. Its complexity, which goes hand in hand with its power is one of the reasons that it has been so hard to document. NTFS first appeared in Windows NT around 1993 (even earlier in beta builds), and now, 10 years later, only 90% of it has been documented. And documenting it is only part of the problem, as implementing actual drivers has been a challenge. Apart from the Linux-NTFS project and NTFSDOS, there are no other common read/write drivers for NTFS. You would expect that after such a long time, either Microsoft would've released some sort of specifications, or that the format would be a common knowledge.

I've tried to keep this article as simple as possible, but some keywords must be first defined, as I will often use them during the documentation. First of all, is it important to know that every file on your disk is stored in the *MFT*, or *Master File Table*. This table is analogous to the *File Allocation Table* on FAT disks. Unlike the simplistic FAT structure, the MFT uses a

variety of *records*, which are structures that define an *object*, most usually a file on disk. The record that describes files is called a *file record*. The peculiar thing about NTFS is that everything is a file. The MFT itself is not seen as a hidden structure, but simply as a file. Just like FAT disks, NTFS disks also have a *Boot Record*, which tells the BIOS information about the disk, and contains the NT Loader code. Once again, the Boot Record is not a hidden structure, but merely a file. The files that contain the MFT, Boot Record, and other information such as the *Volume Label* are called Metafiles. NTFS currently has 16 metafiles, which I'll describe later on. Finally, each File Record is composed mainly of Attributes. Attributes are clearly defined (more on them later), and each one of them is a structure that describes the object in the File Record. Depending on the kind of file, and most especially with Metafiles, some attributes may only be specific to a certain kind of Metafile.

1.3 General Terminology

There are three more important terms you should know when studying filesystems. Information on a disk is stored in what are called *sectors*. A sector contains a set amount of *bytes*, which varies on filesystems. NTFS lets you choose this size. 4KB is common on NT 4, and 512 bytes on Windows 2000/XP. Sectors are grouped in *clusters*. A typical Windows 2000/XP NTFS filesystem has 8 sectors per cluster, which means a total of 4096 bytes, or 4KB. On the physical side of things, hard disks have similar sizes called *cylinders*, *heads per cylinder*. And *sectors per head*. Note that these sectors are different from the virtual ones described above. These three definitions make up what is called the CHS, which is what you would usually set up in the BIOS to properly detect your hard drive. Most recent computers now use AUTO mode however.

1.4 NTFS Versions

One more important issue with NTFS remains the multiple versions. There are two major ones that differ greatly. The one used in NT 4, which corresponds to *NTFS 1.2+* syntax, and the one used in Windows 2000 and XP, which corresponds to *NTFS 3.0+* syntax. Note that these might not apply to NTFS 4.0, but they are common between 3.0 and 3.1.

2. The Boot Record

2.1 Definition

The Boot Record is a tightly packed structure that contains the basic information the BIOS needs to load your OS. It contains a common section, called the *Boot Parameter Block*, which all filesystems and OSes share. This section tell the BIOS information like the size of the disk, and where to look for the *Boot Loader Code*, which will execute the first code sequence of the OS and starting loading it. The second section, the *Enhanced Boot Parameter Block*, greatly differs between filesystems. It is used by the OS or filesystem driver in order to obtain important secondary information about the disk. For example, NTFS stores the location of the first MFT Cluster.

2.2 Structure

Here's what the structure looks like:

```
Public Type BOOT_RECORD
    Jump(2)           As Byte
    OEMID(7)         As Byte
    BPB              As BIOS_PARAMETERS_BLOCK
    EBPB            As EXTENDED_BIOS_PARAMETERS_BLOCK
    BootStrap(425)   As Byte
    BootSignature    As Integer
    Padding          As Long
    Padding          As Long
End Type
```

Structure 2.2.1 - Boot Record

The first three bytes are an ASM jump command for the boot loader, and the next eight bytes are the OEM ID. You might've seen it before, usually named MSWIN4.0 or simply MSDOS. On NTFS, this is aptly named NTFS. The next two elements are two structures, which I've explained the basic role above. Let's look at them:

```
Public Type BIOS_PARAMETERS_BLOCK
    BytesPerSector(1)      As Byte
    SectorsPerCluster      As Byte
    ReservedSectors        As Integer
    Padding(2)             As Byte
    Unknown                 As Integer
    MediaType              As Byte
    Padding1               As Integer
    SectorsPerTrack        As Integer
    NumberOfHeads          As Integer
    HiddenSectors          As Long
    Unknown1(1)            As Long
End Type
```

Structure 2.2.2 - BIOS Parameter Block

The first three elements are the virtual information that the filesystem can usually choose when formatting, while the last 3 elements (except for the Unknown) define the physical values that interface with the BIOS CHS information. Finally, the Media Type describes what kind of media this is. We will usually find 0xF8, which is the hexadecimal constant for hard disk.

```
Public Type EXTENDED_BIOS_PARAMETERS_BLOCK
    TotalSectors           As LARGE_INTEGER
    MFTStartLcn            As LARGE_INTEGER
    MFTMirrorStartLcn     As LARGE_INTEGER
    ClustersPerFileRecord As Long
    ClustersPerIndexBlock As Long
    VolumeSerialNumber     As Long
    VolumeSerialNumber64  As Long
    CheckSum               As Long
End Type
```

Structure 2.2.3 - Extended BIOS Parameter Block

The first element gives the number of total sectors on the hard disk. By finding out how many bytes are in a sector and multiplying, we get the total usable size of the disk in bytes. The MFTStartLcn is the logical cluster where the MFT Metafile starts. The NT Loader must of course read this before doing anything else, as it needs it to load drivers and other important system files. The MFIMirrorStartLcn is the same information for the MFT Mirror, which will be discussed later. The other information that follows is useful when reading the MFT, and we then finally have the Serial Number of the Volume, which is assigned by the partition utility (and can be changed). The checksum is always zero, for now.

Getting back to the Boot Record, the following element is the *Bootstrap Code*. This is the compiled code that will represent the NT Loader. Each OS writes its own Boot Loader in the Bootstrap Code area. For NT, this is the code that will initialize, among other things, the Boot Menu. Finally, a *Boot Signature* marks the end of the Boot Record.

3. Metafiles

3.1 Introduction

3.1.1 Description

Metafiles are an important concept when it comes to learning the way in which NTFS was designed. As said above, everything in NTFS is a file. This means that information such as the cluster location of a file (the FAT in FAT32 partitions), the Volume Name, the Boot Sector and even Directories have to be seen as files. Because it isn't convenient for the user (and quite dangerous) to show files like "C:\bootsector.bin", NTFS uses the concept of Metafiles. These are totally hidden system files (not hidden by a simple "hidden" attribute, but hidden by the NTFS Driver itself) that contain the core information about an NTFS partition. By reading them, one is able to decode the entire on-disk structure of any NTFS partition.

3.1.2 Listing (NTFS 3.0+)

<i>File Record</i>	<i>Filename</i>	<i>Description</i>
0	\$MFT	Master File Table
1	\$MFTMirr	A backup copy of the first 4 records of the MFT
2	\$LogFile	Log File for CHKDSK
3	\$Volume	Volume Name, Serial Number etc...
4	\$AttrDef	Definitions of every Attribute
5	.(dot)	Root directory of the disk
6	\$Bitmap	Map of used and unused clusters
7	\$Boot	Boot record of the volume
8	\$BadClus	List of bad clusters on the partition
9	\$Secure	Security Descriptors for each file
10	\$UpCase	Table of uppercase characters used for conversion
11	\$Extend	Directory for the last four Metafiles.
12-23	UNUSED	Marked in use, or not in use, but empty.
Any	\$ObjId	Unique Object IDs given to every file
Any	\$Quota	Disk space usage quota information
Any	\$Reparse	Reparse point information
Any	\$UsnJrnl	NTFS USN Journal (for encryption)

Table 3.1.1 - NTFS 3.0+ Metafiles

The Metafiles above are present on all NTFS disks formatted by Windows 2000 or higher and thus part of the NTFS 3.0+ format. These file records are always present, even on a newly formatted partition. As such, the MFT must contain at least a minimum of 16 file records (16 to 23 are marked as reserved). The \$MFT is always the first record, and thus points to itself. Some Metafiles worth mentioning in higher detail are described below. The remainder of this chapter will present each Metafile in detail, along with its description and relevant structure in reading it.

3.2 \$MFT

3.2.1 Description

File Record: 0. Relevant structures in: \$DATA

The \$MFT file corresponds to the single most important structure of an NTFS partition, the Master File Table. As said above, every file on the partition (including \$MFT itself) is described in here. Even if your data is on the disk, should the MFT get damaged, it will be very hard to recover without knowing the exact clusters where the pieces of data are kept (almost impossible actually, unless you have a 100% unfragmented volume). The MFT is composed of successive file records, each describing attributes for a given object/file.

3.2.2 Structure

The Master File Record contains successive File Records. These are preceded by the File Record, which in itself contains the general NTFS Record Header, since File Records are not the sole type of records available on NTFS (other records will be shown later). Their structure follows the definition below:

```

Public Type NTFS_RECORD_HEADER
    Type                As Long
    UsaOffset           As Integer
    UsaCount            As Integer
    Usn                 As LARGE_INTEGER
End Type

```

Structure 3.2.2.1 - NTFS Record Header

```

Public Type FILE_RECORD_HEADER
    Ntfs                As NTFS_RECORD_HEADER
    SequenceNumber      As Integer
    LinkCount           As Integer
    AttributesOffset    As Integer
    Flags               As Integer
    BytesInUse          As Long
    BytesAllocated      As Long
    BaseFileRecord      As LARGE_INTEGER
    NextAttributeNumber As Integer
End Type

```

Structure 3.2.2.2 - File Record Header

```

Public Enum FileRecordFlags
    InUse = &H1
    IsDirectory = &H2
End Enum

```

Enumeration 3.2.2.1 - File Record Flags

3.3 \$MFTMirr

3.3.1 Description

File Record: 1. Relevant structures in: \$DATA

To avoid catastrophic damage to the partition, the \$MFTMirr Metafile contains a backup of the first four records in the \$MFT. These are the Volume information, without which the partition wouldn't even be recognized, the Logfile, which might be used for recovery, and a link to itself as well as the MFT. If the file records in the MFT are still valid, but it is only the MFT file record present in the MFT itself that is damaged (which means that the MFT can't find

itself), the MFTMirr will allow the drive to function normally and auto-repair itself.

3.3.2 Structure

The MFT Mirror contains the same structure as the MFT. It only has the first four File Records however. A program has no need to read it unless it wants to check for integrity (however when such damage has occurred, it is rare that a VB program could be able to help).

3.4 \$LogFile

3.4.1 Description

File Record: 2. Relevant structures in: \$DATA

The \$LogFile Metafile contains the logfile that CHKDSK will use when fixing your disks. Basically, every change done on the filesystem will be recorded here, in a circular structure. After a couple of seconds, if these changes have been made and your computer is still functioning, they will be saved as “OK” changes, meaning that they have been permanently written to the partition and are not lost somewhere or non-referenced. If however, there is a sudden power failure before the OK is given, CHKDSK will see the difference in the structures, which also contain undo information, and will effectively undo the modification, or in some cases complete it if it’s possible. The structure of the file is not given here, because it is very complex, prone to change, and pretty much useless unless you’re developing a low-level NT Native Application similar to CHDKSK, in which case the Microsoft IFS Kit might provide more information.

3.4.2 Structure

The Logfile structure is still very unknown. It is implemented by a special NT API which can be extended to other file systems as well. Because of its limited use in a VB application, only brief details will be given. The Logfile is a sequence of 4KB *Log Records*, with the type “RCRD”, followed by a sequence of 4KB *Restore Records*, with type “RSTR”. Every modification is recorded, as well as with information on how to restore it if it wasn’t written to disk properly. If your computer crashes during the write, the Logfile won’t have marked the operation as “complete”, and chkdsk will read the RCRD/RSTR records in order to undo the modifications made to your disk.

3.5 \$Volume

3.5.1 Description

File Record: 3. Relevant structures in: \$VOLUME Attributes

The \$Volume Metafile, once again, is one of those pieces of information that wouldn’t normally be a “file”. As repeated many times however, in NTFS everything is a file, and so is the Volume’s Name. The Metafile has no data stream, but instead, like the dot Metafile which includes special Index Attributes, includes two special Attributes called \$VOLUME_NAME and \$VOLUME_INFORMATION in which the volume information is located.

3.5.2 Structure

Because the information of the Volume Metafile is organized and presented in Attributes, its structure will be discussed in detail in Chapter 4, under the respective attributes contained in it. Note however that only the Volume Metafile includes these attributes.

3.6 \$AttrDef

3.6.1 Description

File Record: 4. Relevant structures in: \$Data

This useful Metafile contains structure definitions for all the attributes used on the filesystem. It also gives the minimum and maximum size, if these are available. The Metafile contains a series of ATTRIBUTE_DEFINITION structures, each with a name, and describing the respective attribute. It allows developers to easily find out more about which attributes they can read on the filesystem.

3.6.2 Structure

The Attribute List follows these structures:

```
Public Type ATTRIBUTE_DEFINITION
    AttributeName(127)    As Byte
    AttributeNumber      As Long
    DisplayRule          As Long
    CollationRule        As Long
    Flags                As Long
    MinimumSize          As LARGE_INTEGER
    MaximumSize          As LARGE_INTEGER
End Type
```

Structure 3.6.2.1 - Attribute Definition

```
Public Enum CollationRules
    Binary = &H0
    FileName = &H1
    UnicodeString = &H2
    UnsignedLong = &H10
    SID = &H11
    SecurityHash = &H12
    MultipleUnsignedLongs = &H13
End Enum
```

Structure 3.6.2.2 - Collation Rules

```

Public Enum AttrDefFlags
    Indexed = &H2
    Resident = &H40
    Nonresident = &H80
End Enum

```

Structure 3.6.2.3 - Attribute Definition Flags

The end of the attribute list is marked by the number &HFFFFFFF, or -1 in decimal. As such, a loop should check whether the first four bytes of the AttributeName element are FF FF FF FF (255 255 255 255). Reading this Metafile isn't of much importance unless you are suspecting that the Filesystem is a modified or unknowng NTFS version.

3.7 . (Dot)

3.7.1 Description

File Record: 5. Relevant structures in: \$Index (\$I30) Attributes

The dot Metafile is the only user-visible metafile at all times. As a matter of fact, you work with it several times a day, as it is simply the root directory of your volume, also sometimes called \. It is simply structured as a normal directory which, under NTFS, is simply a file. This file contains an Index Tree containg files or other subfolders in its attributes. The structure of directories is in a format called B-Tree, or more specifically for NTFS a B*/B+ Tree (it is a bit a mix of both). I won't go deeper in the study of such a tree, but a quick search on Google will offer many interesting sites. More information about directories and indexes is available in the Chapter 4 and 5. The only difference between the dot Metafile and a normal folder is a data stream called \$MountMgrDatabase, which is present when the volume has Reparse Points (directories that point to other volumes or directories).

3.7.2 Structure

This metafile is like a normal directory file, except for the Mount Manager Database which is structured in a repeating array of the following structure:

```
Public Type MOUNTMGRDATABASE
    EntrySize           As Long
    Flags               As Long
    PathOffset          As Long
    PathSize            As Long
    DataOffset          As Long
    DataSize            As Long
End Type
```

Structure 3.6.2.1 - \$MountMgrDatabase Structure

3.8 \$Bitmap

3.8.1 Description

File Record: 6. Relevant structures in: \$Data

The \$Bitmap Metafile can be quite hard to understand and read by using Visual Basic. It is however a smart way that NTFS uses to know which clusters are free and which clusters are used. By using this Metafile, it is easy to get a volume map (like Scandisk or Defrag show) of clusters instead of actually scanning the disk cluster by cluster and seeing which ones are in use or not. It also lets NTFS know where to write a new file.

3.8.2 Structure

This Metafile looks like a bunch of random characters, or usually zeroes. Actually, every bit corresponds to a cluster. The first byte therefore corresponds to the first 8 clusters on the disk. If the byte is 1, the cluster is in use. If it's 0, then the cluster is free.

3.9 \$Boot

3.9.1 Description

File Record: 7. Relevant structures in \$Boot Attribute

The \$Boot Metafile contains the boot sector of the current NTFS partition. It is an exact copy of the raw Boot Sector that can be read using any disk editor.

3.9.2 Structure

The \$Boot Metafile follows the exact same structure as described in Chapter 2. Because finding the \$MFT requires a program to read the Boot Sector by using raw cluster reading, it is rather useless to read the \$Boot Metafile instead.

3.10 \$BadClus

3.10.1 Description

File Record: 8. Relevant structures in \$Data (\$Bad Stream)

Closely similar to the \$Bitmap Metafile, this is another way for NTFS to quickly determine which clusters are “Bad” (meaning corrupted, damaged or unreadable). The difference is that unlike \$Bitmap, the \$BadClus Metafile is a *sparse file*. More information about this kind of file is available in Chapter 5.

3.10.2 Structure

Because the Metafile is sparse, it is basically unreadable by using normal means. The file is as big as the whole volume, but is sparse if the cluster is OK (meaning the data isn’t really there). If a cluster is damaged, then that specific location has data written to it.

3.11 \$Secure

3.11.1 Description

File Record: 9. Relevant structures in \$Index Attributes & \$Data

The \$Secure Metafile is probably the most important aspect of the NTFS Security implementation. Although older versions of the filesystem (NTFS 1.2) would use a special Attribute to store the security information (more on attributes in the next chapter), NTFS 3.0+ stores all the security information into a single Metafile, \$Secure.

Until now, all the data contained in the Metafiles was usually present in the \$DATA attribute, that is in the file itself, and not in any headers (with the exception of \$Volume). Unfortunately, \$Secure is a much more complicated beast, and the data is also cross-linked with two Index Attributes. Once again, Attributes are only discussed in Chapter 4, so it might be worthwhile to read it first before reading the information on \$Secure.

The \$Secure Metafile has a single named data stream (if you remember the previous article) called \$SDS, or *Security Descriptor Stream*. It contains a list of all the *Security Descriptors* on the partition, which are defined in the \$SECURITY_DESCRIPTOR attribute in Chapter 4. Even though that attribute isn't used anymore, the structure is the same, and is now in a gigantic list inside \$Secure.

The Security ID element of the \$STANDARD_INFORMATION Attribute (see Chapter 4) contains a number, which is linked into one of the Indexes (\$SII, *Security Id Index*), which in turn is cross-referenced with \$SDH (*Security Descriptor Hash*). The number is also located in the \$SDS stream, where a header for each Security ID gives the offset of the Security Descriptor structure. Once the descriptor is located, it can be read like its corresponding attribute.

3.11.2 Structure

Unlike normal Index Attributes shown in Chapter 4, the \$SDH Indexes follow the particular structure show above, although they still include the normal Index Header depicted in the next chapter. The Security Descriptor Hashes are organized according to the following structure:

```
Public Type SECURITY_DESCRIPTOR_HASH_INDEX
    IndexHeader           As INDEX_ALLOCATION_ATTRIBUTE
    DataOffset            As Integer
    DataSize              As Integer
    Padding               As Long
    IndexEntrySize        As Integer
    IndexKeySize          As Integer
    Flags                 As Integer
    Padding               As Integer
    SecurityHashKey       As Long
    SecurityIDKey         As Long
    SecurityHashData      As Long
    SecurityIDData        As Long
    SecureDescriptorOffset As LARGE_INTEGER
    SecurityDescriptorSize As Long
    Padding               As LARGE_INTEGER
End Type
```

Structure 3.11.2.1 - Security Descriptor Hash Index Structure

Similarly to the \$SDH Index, the \$SII Index also has a specific structure, closely resembling the \$SDH structure. In both cases, these structures are repeated one after the other. It is possible to determine the last \$SDH index by looking at the Padding element, which should be “II” in Unicode once the last index has been reached. Of course, they can also be calculated by looking at the index header and offsets/sizes. The Security ID Indexes are structured as following:

```

Public Type SECURITY_ID_INDEX
    IndexHeader           As INDEX_ALLOCATION_ATTRIBUTE
    DataOffset            As Integer
    DataSize              As Integer
    Padding               As Long
    IndexEntrySize       As Integer
    IndexKeySize         As Integer
    Flags                As Integer
    Padding              As Integer
    SecurityIDKey        As Long
    SecurityHashData     As Long
    SecurityIDDData      As Long
    SecureDescriptorOffset As LARGE_INTEGER
    SecurityDescriptorSize As Long
End Type

```

Structure 3.11.2.2 - Security ID Index Structure

Finally, in the \$SDS data, all the Security Descriptors are one after the other. They are all preceded by the header structure shown above. For the Descriptors themselves, they are shown in Chapter 4, since they have the exact same structure as the \$SECURITY_DESCRIPTOR Attribute that used to be present as an Attribute in each file on NTFS 1.2

```

Public Type SECURITY_DESCRIPTOR_ENTRY
    Hash                 As Long
    SecurityID           As Long
    EntryOffset          As LARGE_INTEGER
    EntrySize            As Long
End Type

```

Structure 3.11.2.3 - Security ID Index Structure

3.12 \$UpCase

3.12.1 Description

File Record: 10. Relevant structures in \$Data

The \$UpCase Metafile can be pretty hard to explain since it's hard to understand its use. It basically allows NTFS to compare filenames without caring about the character set or codepage. Every Unicode character is mapped to an upper case character which makes it easy to compare everything equally.

3.12.2 Structure

The \$UpCase Metafile doesn't have a specific readable structure. It is 128KB in size and enumerates all the Unicode Characters allowing with an ASCII representation.

3.13 \$Extend

3.13.1 Description

File Record: 11. Relevant structures in \$Index Attributes

The \$Extend Metafile is simply a Directory Index that contains information on where to locate the last four Metafiles (\$ObjId, \$Quota, \$Reparse and \$UsnJrnl). It has the same structure as the dot Metafile or any other Direcorey.

3.15 \$ObjID

3.15.1 Description

File Record: ANY. Relevant structures in \$Index Attributes

Because NT is an object-oriented system in which the kernel and all the modules are viewed as "objects", NTFS files themselves also are objects, and each have their own unique GUID. The \$ObjID metafile contains an Index of all the \$OBJECT_ID Attributes on the paritiion called 0\$. This GUID is rarely used but it is possible to use Native API to open a file on an NTFS volume with its GUID.

3.15.2 Structure

Much like the Index Attributes in \$Secure, \$ObjID includes an array of \$OBJECT_ID Attributes that use the following organizational structure:

```
Public Type OBJECT_ID_INDEX
    IndexHeader           As INDEX_ALLOCATION_ATTRIBUTE
    DataOffset            As Integer
    DataSize              As Integer
    Padding               As Long
    IndexEntrySize        As Integer
    IndexKeySize          As Integer
    Flags                 As Integer
    Padding2              As Integer
    GUIDObjectIDKey       As String * 16
    MFTReferenceData      As LARGE_INTEGER
    GuidBirthVolumeIDData As String * 16
    GuidBirthDomainData   As String * 16
    GuidBirthDomainIDData As String * 16
End Type
```

Structure 3.15.2.1 - Object ID Index Structure

3.16 \$Quota

3.16.1 Description

File Record: ANY. Relevant structures in \$Index Attributes

Just like \$Secure, the \$Quota Metafile is another important piece of the security and features of NTFS 3.0+. This Metafile includes all the Quota information for each user, arranged by SID. Once again, the information is present in two indexes, in this case O\$ (pay attention to the fact this Index has the same name as the \$ObjID Indexes, but they have nothing in common) and Q\$. It would theoretically be possible to modify the Quotas of a certain user by modifying this on-disk structure, so admins should always be wary.

3.16.2 Structure

The first Index, \$Q, contains an entry for each UserID on the partition. The Index Key points to each UserID that owns the quota. The UserID of a file (the owner) can easily be found in the \$STANDARD_INFORMATION Attribute of every file. The \$Q Index can be read as described below, after having loaded the Index Attribute (see Chapter 4):

```
Public Type QUOTA_LIMIT_INDEX
    IndexHeader           As INDEX_ALLOCATION_ATTRIBUTE
    DataOffset           As Integer
    DataSize             As Integer
    Padding              As Long
    IndexEntrySize       As Integer
    IndexKeySize         As Integer
    Padding2             As Long
    KeyOwnerID           As Long
    DataVersion          As Long
    DataFlags            As Long
    DataBytesUsed        As LARGE_INTEGER
    DataChangeTime       As LARGE_INTEGER
    DataWarningLimit     As LARGE_INTEGER
    DataHardLimit        As LARGE_INTEGER
    DataExceedTime       As LARGE_INTEGER
    DataSID()            As Byte
    Padding3()           As Byte
End Type
```

Structure 3.16.2.1 - Quota Limits Index Structure

Do take note that DataSID is variable, so you must determine the size beforehand. As for the \$O Index, it contains an entry for each User or Group that has been assigned a quota on the partition. The Index Key contains the SID of the UserID to which the entry pertains to. Finally, the Index Entry Data contains the UserID corresponding to the current SID. This UserID can be used to read the proper Quota associated with the SID. The \$O Index is structured as follows:

```

Public Type OWNER_ID_INDEX
    IndexHeader           As INDEX_ALLOCATION_ATTRIBUTE
    DataOffset            As Integer
    DataSize              As Integer
    Padding               As Long
    IndexEntrySize        As Integer
    IndexKeySize          As Integer
    Flags                 As Integer
    Padding2              As Long
    KeySID()              As Byte
    DataOwnerID           As Long
    Padding3()            As Byte
End Type

```

Structure 3.16.2.2 - Owner ID Index Structure

Once again, please take note that KeySID is variable. You should only read the structure until IndexKeySize, after which you can re-dimension the array and read the entire structure properly. Additionally, the following Flags are defined for the \$Q Structure:

```

Public Enum QuotaFlags
    DefaultLimitNoSID = &H1
    LimitReached = &H2
    IDDeleted = &H4
    TrackingEnabled = &H10
    EnforcementEnabled = &H20
    TrackingRequested = &H40
    LogTreshold = &H80
    LogLimit = &H100
    OutOfDate = &H200
    Corrupt = &H400
    PendingDeletes = &H800
End Enum

```

Enumeration 3.16.2.1 - Quota Limits Flags Enumerations

A value of 1 indicates that there will be no SID and that only padding will be present. This should be taken into account before loading the structure. Also, in both structures, Padding3 is variable but will always align the structure to 8 bytes.

3.17 \$Reparse

3.17.1 Description

File Record: ANY. Relevant structures in \$Index Attributes

Reparse Points are an exciting feature of NTFS which allow it to mount another Volume under a directory. For example, your D:\CD-ROM drive could be mapped as c:\cdrom. To keep track of these Reparse Points, the \$Reparse Metafile contains an Index called \$R of all of them, pointing at the MFT File Record which contains the \$REPARSE_POINT Attribute for the file.

3.17.2 Structure

The \$R Index doesn't contain much information, as most of the really useful data pertaining to the Reparse Point is contained in the attribute (see Chapter 4). The \$R Index can be read as follows:

```
Public Type REPARSE_INDEX
    IndexHeader           As INDEX_HEADER
    DataOffset            As Integer
    DataSize              As Integer
    Padding               As Long
    IndexEntrySize        As Integer
    IndexKeySize          As Integer
    Flags                 As Integer
    Padding2              As Integer
    KeyReparseTagAndFlags As Long
    KeyMFTReference       As LARGE_INTEGER
    Padding3              As Long
End Type
```

Structure 3.17.2.1 - Reparse Points Index Structure

The Reparse Tag and Flags are documented in the \$REPARSE_POINT Attribute section in Chapter 4 and are not shown here.

3.18 \$UsnJrnl

3.18.1 Description

File Record: ANY. Relevant structures in \$DATA (\$J & \$Max).

The USN Change Journal made its appearance in NTFS 3.0 supposedly for allowing OEMs to double-check the warranty on customer's computers, as well as to allow companies or law enforcement to check a full, complete log of everything that happened on the partition, in the \$UsnJrnl Metafile. Windows never reads the file by itself or needs it. It can be normally accessed by using the *DeviceIoControl* API, which is documented along with the proper functions and flags. However, the On-Disk method is just as viable.

3.18.2 Structure

The Entries in the Metafile are present in a Data Stream called \$J. They are a repetition of the same structure below:

```
Public Type USN_JRNL_ENTRY
    EntrySize           As Long
    MajorVersion        As Integer
    MinorVersion        As Integer
    MFTFileReference    As LARGE_INTEGER
    MFTParentFileReference As LARGE_INTEGER
    EntryOffset         As LARGE_INTEGER
    TimeStamp           As LARGE_INTEGER
    ReasonFlags         As Long
    SourceInfoFlags     As Long
    SID                 As Long
    FileAttributes      As Long
    FilenameSize        As Integer
    FilenameOffset      As Integer
    Filename()          As Byte
    Padding()           As Byte
End Type
```

Structure 3.18.2.1 - USN Journal Entry Structure

The Reason Flags are very useful in determining exactly the nature of the change, and can include pointers to useful information:

```
Public Enum UsnReasonFlags
    DataOverWrite = &H1
    DataExtend = &H2
    DataTruncation = &H4
    NamedDataOverWrite = &H10
    NamedDataExtend = &H20
    NamedDataTruncation = &H40
    FileCreate = &H100
    FileDelete = &H200
    EAChange = &H400
    SecurityChange = &H800
    RenameOldName = &H1000
    RenameNewName = &H2000
    IndexableChange = &H4000
    BasicInfoChange = &H8000
    HardLinkChange = &H10000
    CompressionChange = &H20000
    EncryptionChange = &H40000
    ObjectIDChange = &H80000
    ReparsePointChange = &H100000
    StreamChange = &H200000
    FileClose = &H8000000
End Enum
```

Enumeration 3.18.2.1 - USN Journal Reason Flags Enumerations

Finally, the Source Flags will document who caused the modification. Unless a Replication/Backup Service is being used, this will always most likely be Data Management (normal use):

```
Public Enum UsnSourceFlags
    DataManagement = &H1
    AuxiliaryData = &H2
    ReplicationManagement = &H3
End Enum
```

Enumeration 3.18.2.2 - USN Journal Source Flags Enumerations

4. Attributes

4.1 Introduction

4.1.1 Definition

Before starting to describe each attribute, it is important to understand what they really represent, and their different headers. An attribute is basically a piece of information that defines a file. If you are more familiar with processes and threads, think of the file as a process, and of the attribute as threads. In other words, the file is merely a container for the attributes, and does not exist as an entity itself (except as a File Record). Many people think of a file as a piece of data, but for NTFS, even the data is a separate attribute. All these defining pieces of information, attributes, are pointed and described by the File Record of the file inside the MFT. To read a File Record itself, view the documentation in Chapter 3, \$MFT.

4.1.2 Listing (NTFS 3.0+)

<i>Type</i>	<i>Name</i>	<i>IRN</i>	<i>Min Size</i>	<i>Max Size</i>
&H10	\$STANDARD_INFORMATION	R	48	72
&H20	\$ATTRIBUTE_LIST	N	-	-
&H30	\$FILE_NAME	IR	68	578
&H40	\$OBJECT_ID	R	-	256
&H50	\$SECURITY_DESCRIPTOR	N	-	-
&H60	\$VOLUME_NAME	R	2	256
&H70	\$VOLUME_INFORMATION	R	12	12
&H80	\$DATA	-	-	-
&H90	\$INDEX_ROOT	R	-	-
&HA0	\$INDEX_ALLOCATION	N	-	-
&HB0	\$BITMAP	N	-	-
&HC0	\$REPARSE_POINT	N	-	16384
&HD0	\$EA_INFORMATION	R	8	8
&HE0	\$EA	-	-	65536
&HF0	\$PROPERTY_SET	-	-	-
&H100	\$LOGGED_UTILITY_STREAM	N	-	65536

Table 4.1.1 - NTFS 3.0+ Attributes

IRN: Indexed/Resident/Nonresident
Type: NTFS Hexadecimal Attribute Type

These Attributes are present on Windows 2000 and higher formatted partitions. They are based on the \$AttrDef Metafile. Some differences exist in regards to NTFS 1.2, but you should be able to read the important ones as well. Another important thing to note is that \$PROPRETY_SET is not available on NTFS3.0+ but was only used in NT4. The signification of the IRN value will be described below. Sizes are in decimal bytes, but the type is in hexadecimal because this is how Microsoft made them. Writing in decimal would make it harder to remember and less well organized. The remaining sections of this chapter will describe in detail most of the attributes, along with their structure.

4.1.3 Structure

The VB Enumeration of the table above is defined as follows and can be used in a Select Case loop easily:

```
Public Enum NtfsAttributes
    AttributeStandardInformation = &H10
    AttributeAttributeList = &H20
    AttributeFileName = &H30
    AttributeObjectId = &H40
    AttributeSecurityDescriptor = &H50
    AttributeVolumeName = &H60
    AttributeVolumeInformation = &H70
    AttributeData = &H80
    AttributeIndexRoot = &H90
    AttributeIndexAllocation = &HA0
    AttributeBitmap = &HB0
    AttributeReparsePoint = &HC0
    AttributeEAInformation = &HD0
    AttributeEA = &HE0
    AttributePropertySet = &HF0
    AttributeLoggedUtilityStream = &H100
End Enum
```

Enumeration 4.1.1 - NTFS 3.0+ Attributes

4.2 Types of Attributes

4.2.1 Attribute Definition

Before starting to describe each attribute, it is important to understand what they really represent, and their different headers. An attribute is basically a piece of information that defines a file. If you are more familiar with processes and threads, think of the file as a process, and of the attribute as threads. In other words, the file is merely a container for the attributes, and does not exist as an entity itself (except as a File Record). Many people think of a file as a piece of data, but for NTFS, even the data is a separate attribute. All these defining pieces of information, attributes, are pointed and described by the File Record of the file inside the MFT.

4.2.2 Attribute Structure

Every attribute, be it resident, non-resident, named or unnamed contains the same standard *Attribute Header*. This header gives the basic information about the attribute, which in turns tells NTFS how it should read it. It has the following format:

```
Public Type NTFS_ATTRIBUTE
    AttributeType          As NtfsAttributes
    Length                 As Long
    Nonresident            As Byte
    NameLength             As Byte
    NameOffset             As Integer
    Flags                  As Integer
    AttributeNumber        As Integer
End Type
```

Structure 4.2.2.1 - Attribute Header

There are three more flags which can represent in which way the attribute is present:

```

Public Enum AttributeFlags
    Compressed = &H1
    Encrypted = &H4000
    Sparse = &H8000
End Enum

```

Enumeration 4.2.2.1 - Attribute Flags

The only difference between the header of a resident and non-resident attribute is that the nonresident byte will be set to 1 if the attribute is nonresident. In both cases, the representation of “NameLength” and “NameOffset” also vary, depending if the attribute is named or not. On an unnamed attribute, NameLength will simply be 0, and the NameOffset will actually point to the Attribute Offset. (see below).

4.2.3 Nonresident Attribute Definition

As you’ve no doubt wondered, if the data is an attribute, and attributes are in the MFT, then what happens if the data is over 4KB, and can’t fit into the MFT? It turns out that not only data has this characteristic, but many other attributes contain huge amounts of information inside them. As such, these attributes must be put outside the File Record, but still be linked by it. This concept is called a *Non-Resident Attribute*.

4.2.4 Nonresident Attribute Structure

The Nonresident Attribute Header precedes a nonresident attribute, and as such contains much more information on where to read the attribute, containing VCNs (*Virtual Cluster Numbers*) and other data. Furthermore, these attributes can also be compressed, and information is available on that as well. The Nonresident Attribute can be read in the structure below:

```

Public Type NONRESIDENT_ATTRIBUTE
    Attribute           As NTFS_ATTRIBUTE
    LowVcn              As LARGE_INTEGER
    HighVcn             As LARGE_INTEGER
    RunArrayOffset     As Integer
    CompressionUnit    As Byte
    AlignmentOrReserved As Long
    AllocatedSize      As LARGE_INTEGER
    DataSize           As LARGE_INTEGER
    InitializedSize    As LARGE_INTEGER
    CompressedSize     As LARGE_INTEGER
End Type

```

Structure 4.2.4.1 - Nonresident Attribute Header

The Attribute data is contained in the RunArrayOffset, where the Data Runs are located. The attribute is then distributed according to LowVCN and HighVCN (the first and last VCNs of the attribute). Once again, the same distinction is made between named/unnamed attributes regarding the offsets (This time the offset that points to the data is RunArrayOffset). The CompressedSize information is present only if the attribute is compressed. If not, it is usually the beginning Data Runs, or of the attributes's name followed by the Data Runs.

4.2.5 Resident Attribute Definition

An attribute that is guaranteed to always be present in the MFT File Record because of a defined size is called a *Resident Attribute*.

4.2.6 Resident Attribute Structure

The Resident Attribute Header precedes a resident attribute, and is quite easy to read, since only offsets are present. It is defined as such:

```

Public Type RESIDENT_ATTRIBUTE
    Attribute           As NTFS_ATTRIBUTE
    ValueLength         As Byte
    ValueOffset         As Integer
    IndexedFlag         As Integer
End Type

```

Structure 4.2.6.1 - Resident Attribute Header

If the attribute is named, ValueOffset will point to the attribute data, located after the name. If the attribute is unnamed, ValueOffset will have the same value as NameOffset in the Attribute Header.

4.2.7 Named and Unnamed Attributes

There is one more distinction to be made considering attributes: *Named Attributes* and *Unnamed Attributes*. Take notice that this does not refer to the name of the Attribute as in “\$STANDARD_INFORMATION”, which is a pre-defined file system name, but the real name of the Attribute itself. For example, \$DATA is an attribute, but it can contain data that has different names. These are called Alternate Data Streams by the way (Discussed in my previous article). Another example are the different \$Secure Indexes, which as Chapter 3 mentioned, are named \$SDS and \$SDH.

4.3 \$STANDARD_INFORMATION

4.3.1 Description

Type:0x10. Resident. Min Size: 48 bytes. Max Size: 72 bytes.

The \$STANDARD_INFORMATION Attribute is most often the most useful attribute when it comes to simply reading file information off an NTFS volume. Apart from the Filename, it is usually the only information a basic filesystem would need.

4.3.2 Structure

The structure of this Attribute is easily readable by using the following structure. You can convert all times into complete dates and hours by using the *ConvertFileTimeToSystemTime* APIs.

```
Public Type STANDARD_INFORMATION_ATTRIBUTE
    CreationTime           As LARGE_INTEGER
    ChangeTime            As LARGE_INTEGER
    LastWriteTime         As LARGE_INTEGER
    LastAccessTime        As LARGE_INTEGER
    FileAttributes         As NtfsFileAttributes
    MaxVersionNumber      As Long
    VersionNumber         As Long
    ClassID               As Long
    OwnerID               As Long
    SecurityID            As Long
    QuotaCharge           As LARGE_INTEGER
    Usn                   As LARGE_INTEGER
End
```

Structure 4.3.2.1 - \$STANDARD_INFORMATION Attribute

```
Public Enum NtfsFileAttributes
    ReadOnly = &H1&
    Hidden = &H2&
    System = &H4&
    Archive = &H8&
    Device = &H40&
    Normal = &H80&
    Temporary = &H100&
    SparseFile = &H200&
    ReparsePoint = &H400&
    Compressed = &H800&
    Offline = &H1000&
    NotContentIndexed = &H2000&
    Encrypted = &H4000&
End Enum
```

Enumeration 4.3.2.1 - NTFS File Attributes

The OwnerID is the key in the \$O and \$Q Indexes of the \$Quota Metafile, and SecurityID is the Key in the \$SII and \$SDS Index and Stream of \$Secure. Finally, Usn is an index into \$UsnJrnl. Basically this attribute describes all the Filesystem level information of a file.

4.4 \$ATTRIBUTE_LIST

4.4.1 Description

Type:0x20. Nonresident. Min/Max Size: Variable.

NTFS will always try its best to fit all of a File Record's Attributes into the MFT. If it cannot, it usually moves data out of it and makes it non-resident. This however still keeps the header information inside the MFT. Although it is very rare, a file could have an unusually large amount of attributes needed to be defined, and whose headers cannot fit inside the MFT. To solve this, a special attribute, \$ATTRIBUTE_LIST, is used to point into a completely separate MFT File Record that describes the rest of the Attributes.

4.4.2 Structure

The Attribute List is a series of records organized as shown below:

```
Public Type ATTRIBUTE_LIST_ATTRIBUTE
    AttributeType           As NtfsAttributes
    Length                  As Integer
    NameLength              As Byte
    NameOffset              As Byte
    LowVcn                  As LARGE_INTEGER
    FileReferenceNumber     As LARGE_INTEGER
    AttributeNumber         As Integer
End Type
```

Structure 4.4.2.1 - \$ATTRIBUTE_LIST Attribute

It is very important to note that you must align the structure to 8-bytes to read the next attribute.

4.5 \$FILE_NAME

4.5.1 Description

Type:0x30. Resident. Min Size: 68 bytes. Max Size: 578 bytes.

The job of the \$FILE_NAME Attribute is quite simple. It is where the name of file contained in the File Record is held (or the name of the Directory). This Attribute is also present for each Hard Link (folders actually pointing to other folders or volumes), one for each copy, with the appropriate other information contained. The \$FILE_NAME Attribute has a lot of similarity to the \$STANDARD_INFORMATION Attribute, however there is one subtle change to be aware of: the changes to the different times are only updated if the filename is changed. This means that it can be out-of-date and shouldn't be used to get File Access Times.

4.5.2 Structure

The \$FILE_NAME Attribute has a similar structure to \$STANDARD_INFORMATION, but also includes certain information that might be needed by EA (*Extended Attributes*) (the buffer needed). If this is a Reparse Point however, the Reparse Type will be located in the respective field. Once again, all times can be decoded using the appropriate API. Finally, NTFS supports many "Namespaces" for a certain file, but the most common ones you will meet are Win32 or DOS. It is important to first read which one the file uses, so one can be aware of the proper buffer size (510 bytes or 14 bytes respectively, the first one being Unicode and the second one ANSI). The \$FILE_NAME Attribute can be read using the structure definition below:

```

Public Type FILENAME_ATTRIBUTE
    ParentDirFileRefNumber As LARGE_INTEGER
    CreationTime           As LARGE_INTEGER
    ChangeTime            As LARGE_INTEGER
    LastWriteTime         As LARGE_INTEGER
    LastAccessTime        As LARGE_INTEGER
    AllocatedSize         As LARGE_INTEGER
    RealSize              As LARGE_INTEGER
    FileAttributes        As NtfsFileAttributes
    EABufferOrReparseType As Long
    NameLength            As Byte
    NameSpace             As NtfsFileNameSpaces
    FileName()            As Byte
End Type

```

Structure 4.5.2.1 - \$FILE_NAME Attribute

```

Public Enum NtfsNameSpaces
    POSIX = &H1      ' // Unix-style Unicode
    Win32 = &H2     ' // 255 bytes Unicode
    DOS = &H3       ' // 8.3 Notation
    Win32DOS = &H4  ' // Win32 and DOS equivalent
End Enum

```

Enumeration 4.5.2.1 - NTFS NameSpaces

4.6 \$OBJECT_ID

4.6.1 Description

Type:0x40. Resident. Max Size: 256 bytes.

The \$OBJECT_ID Attribute, new to NTFS 3.0+, contains up to four different GUIDs that reference the file as an object (much like the CLSID of a COM Object). A file can be opened using this GUID by using Native API, and it can also be useful in certain situations when using Active Directory.

4.6.2 Structure

The Attribute has a straight-forward structure of 4 GUIDs:

```

Public Type OBJECT_ID_ATTRIBUTE
    ObjectId                As GUID
    BirthVolumeId          As GUID
    BirthObjectId           As GUID
    DomainId                As GUID
End Type

```

Structure 4.6.2.1 - \$OBJECT_ID Attribute

Because the ObjectID can change in rare situations, the BirthObjectID holds the original assigned value. DomainID is only used on Active Directory. Please be aware that it is possible for only one or two GUIDs to be present.

4.7 \$SECURITY_DESCRIPTOR

4.7.1 Description

Type:0x50. Nonresident. Min/Max Size: Variable.

The \$SECURITY_DESCRIPTOR Attribute, which in NTFS 3.0+ is only present in the \$SDS Stream of the \$Secure Metafile is the pillar of NTFS Security Architecture. Because a lot of background information on terms like *ACL*, *ACE*, *SID* is needed before being able to completely digest the information, a look at Chapter 5 is strongly recommended. The Attribute basically contains the Owner User or Group, as well as the permissions applicable to the file in respective *Access Control Lists/Entries*. It can also optionally include *Auditing Information*, which will describe which kind of accesses to monitor and log. As repeated before, this information is not really an Attribute anymore, but has become a repeating structure in the \$SDS Index of \$Secure.

4.7.2 Structure

The Security Descriptor uses the following structure:

```

Public Type SECURITY_DESCRIPTOR
    Revision                As Byte
    Padding                 As Byte
    ControlFlags            As SecurityControlFlags
    UserSIDOffset           As Long
    GroupSIDOffset         As Long
    SACLOffset              As Long
    DACLOffset              As Long
End Type

```

Structure 4.7.2.1 - \$SECURITY_DESCRIPTOR Attribute

```

Public Enum SecurityControlFlags
    OwnerDefault = &H1
    GroupDefault = &H2
    DACLPresent = &H4
    DACLDefault = &H8
    SACLPresent = &H10
    SACLDefault = &H20
    DACLAutoInheritRequest = &H100
    SACLAutoInheritRequest = &H200
    DACLAutoInherited = &H400
    SACLAutoInherited = &H800
    DACLProtected = &H1000
    SACLProtected = &H2000
    RMControlValid = &H4000
    SelfRelative = &H8000
End Enum

```

Enumeration 4.6.2.1 - Security Descriptor Flags

The different ControlFlags will allow you to check which bits are set, and in this scenario, quickly know if a *Security Access Control List (SACL)* is being used for Auditing. The structure applies to Revision 1 only, so take notice if you are reading a different number (no newer versions are known to exist). The structures to which the different offsets point to will not be presented here, but instead in Chapter 5 under the Security Concepts. Theoretically, accessing this security information with the proper structure definitions could allow an attacker to modify Quota and File Permissions with a relative ease, since nothing is stored encrypted, but simply in normal ACLs/ACEs.

4.8 \$VOLUME_NAME

4.8.1 Description

Type:0x60. Resident. Min Size: 2 bytes. Max Size: 256 bytes.

The Volume Name Attribute is unique in the sense that it is only used in the \$Volume Metafile. Although NTFS could've placed the Volume Name in a Data Stream, perhaps by using a special Attribute it is easier to read by the bootloader. You should use this Attribute whenever you want to read the partition's volume name.

4.8.2 Structure

The attribute is a simple 255-byte null terminated string containing the name of the Volume.

4.9 \$VOLUME_INFORMATION

4.9.1 Description

Type:0x70. Resident. Min Size: 12 bytes. Max Size: 12 bytes.

The Volume Information Attribute is also only used in the \$Volume Metafile. It contains however more information about the Volume than just the name, and it can be useful to determine which NTFS Version you are dealing with (and make your program return an error if it's unsupported). There are other flags which are useful for the NTFS Driver to determine the current state of the Volume.

4.9.2 Structure

The Attribute is a simple structure that is very easy to read according to the following format. Version 3.1 is Windows XP and is compatible with version 3.0 (Windows 2000):

```

Public Type VOLUME_INFORMATION
    Paddin(1)                As Long
    MajorVersion              As Byte
    MinorVersion              As Byte
    Flags                     As VolumeFlags
    Padding                   As Long
End Type

```

Structure 4.9.2.1 - \$VOLUME_INFORMATION Attribute

The Volume Flags have the following meanings below. The Dirty flag means that this Volume will be scanned by chkdsk on next boot:

```

Public Enum VolumeFlags
    Dirty = &H1&
    ResizeLogFile = &H2&
    UpgradeInMount = &H4&
    MountedInNT4 = &H8&
    DeleteUSNUnderway = &H10&
    RepairObjectIDs = &H20&
    ModifiedByChkdsk = &H8000&
End Enum

```

Enumeration 4.9.2.1 - Volume Information Flags

4.10 \$DATA

4.10.1 Description

Type:0x80. Nonresident. Min/Max Size: Variable.

The \$Data Attribute is where all the main information of a file is contained that is not in a special Attribute. For a normal user file, this would correspond to the binary text inside that file. It is important to note however that a file can have multiple \$Data Attributes, which are called Data Streams. These Streams can be named or unnamed, and will not be accounted for (in size) or visible by normal methods (not even by Windows itself). [See ADS Article]

4.10.2 Structure

The \$Data attribute or other streams don't have any structure for they are simply binary data in no organized structure. Although the ADS Article presented a nice and clean way of reading multiple streams, it is also possible to read the \$Data Attributes in the file for the same function.

4.11 \$INDEX_ROOT

4.11.1 Description

Type:0x90. Resident. Min/Max Size: Variable.

Indexes are without any doubt the most complicated part of the NTFS design, and are also responsible for most of its speed and extensibility. In fact, Indexes are responsible for all Directories, Security Descriptors, Quotas, GUIDs and Reparse Points on NTFS. Without them, NTFS would be a flat and slow filesystem, in which data would have to be recursively read from the MFT, slowing down the system to a crawl.

There are two kinds of Index Attributes: \$INDEX_ROOT and \$INDEX_ALLOCATION. The first one is used for very small Indexes, and is always resident, while the second one is a Nonresident Attribute whose *Data Runs* point to *Index Blocks*. This section will deal with \$INDEX_ROOT, which is the *Root Node* of the current Index.

\$INDEX_ROOT usually indicates what kind of Index that is being held, and includes an *Index Header* which points to the *Index Entries* contained within it. The *Index Entry* itself depends on the kind of information being Indexed. The most commonly used example are Filenames, used in Directories.

4.11.2 Structure

The attribute is defined by the following “Master Header”:

```
Public Type INDEX_ROOT_ATTRIBUTE
    Type                As NtfsAttributes
    CollationRule       As Long
    BytesPerIndexBlock  As Long
    ClustersPerIndexBlock As Long
    IndexHeader         As INDEX_HEADER
End Type
```

Structure 4.11.2.1 - \$INDEX_ROOT Attribute

Simply by looking at the Type, which corresponds to NTFS Attribute Types (such as 0x30 for Filename), it is possible to determine what kind of Index Entries will be present. Before the *Index Entries* start however, the *Index Header* will follow:

```
Public Type INDEX_HEADER
    EntriesOffset       As Long
    IndexBlockLength    As Long
    AllocatedSize       As Long
    LargeIndexFlag      As Long
End Type
```

Structure 4.11.2.2 - Index Header

If the flag is set to 1, it means that an *Index Allocation* Attribute will be present. Once the offset to the *Index Entries* has been obtained, they can be read. *Index Entries* however, also have their own header, the *Index Entry Header*, shown below:

```
Public Type INDEX_ENTRY_HEADER
    FileReferenceNumber As LARGE_INTEGER
    Length              As Integer
    KeyLength           As Integer
    Flags               As IndexEntryFlags
End Type
```

Structure 4.11.2.3 - Index Entry Header

```

Public Enum IndexEntryFlags
    PointsToSubnode = &H1&
    LastEntryInNode = &H2&
End Enum

```

Enumeration 4.11.2.1 - Index Entry Header Flags

If the Flag is, or has the Last Entry value (flag is 2 or 3), then the data is either invalid or empty. The KeyLength value is not aligned to 8 bytes, so alignment must be done by the program. Furthermore, if the Flag is, or has the Subnode value (flag is 1 or 3), then 8 bytes pointing to the LCN of the Subnode will follow after the *Index Entry* (or directly after the *Entry Header* if the KeyLength is 0). Finally, note that there is no *Index Entry* structure per-se. As said before, *Index Entries* can be anything, depending on what is being indexed. In the case of directories, this would be the \$FILENAME Attribute for example. For some examples on how to read a Directory, or other indexes, please refer to the Chapter 5, Indexes.

4.12 \$INDEX_ALLOCATION

4.12.1 Description

Type:0xA0. Nonresident. Min/Max Size: Variable.

Many (most) times, \$INDEX_ROOT won't be enough to contain all the indexes in the file. For this reason, \$INDEX_ALLOCATION is present to contain a much larger array of Indexes, because it is Nonresident and can be tracked by its *Data Runs*.

This type of Attribute isn't as easy as it seems though, because the *Data Runs* point not to the *Index Entries* themselves, nor even their headers, but to a completely different structure called an *Index Record* or Block. Much like *File Records*, *Index Records* map out the different indexes that are present.

These *Index Records* are actually the true Nonresident Attribute Data, and as such they are a mix between an NTFS Record and an NTFS Attribute, which does sometimes lead to confusion. It is best to think of the \$INDEX_ALLOCATION Attribute as simply being the *Data Runs* which point to the respective Data: the *Index Records*.

4.12.2 Structure

The \$INDEX_ALLOCATION Attribute should actually be referred to as an *Index Record*, because it isn't an attribute per-se. It's a bit similar to the *File Record*, and shares the same NTFS Record Header and USN Structures at the end (see Chapter 5, USN). The data in between however, is largely different:

```
Public Type INDEX_ALLOCATION_ATTRIBUTE
    NtfsRecordHeader      As NTFS_RECORD_HEADER
    IndexBlockVcn         As LARGE_INTEGER
    IndexHeader           As INDEX_HEADER
    UpdateSequence        As Integer
    UpdateSequenceArray() As Byte
End Type
```

The Index Block VCN simply is the VCN of the current *Index Record* in the current *Data Run* (see Chapter 5, Data Runs). The *Index Header* is exactly the same one as has been defined before, however it is not followed by the respective *Index Entries*, which actually come after the USN data.

The whole INDEX_ALLOCATION_ATTRIBUTE structure is actually what precedes all Indexes such as \$I30, \$SII and \$Q. Once again, a look at Chapter 5, Indexes, will present a much clearer view on how to properly read these Indexes.

4.13 \$BITMAP

4.13.1 Description

Type:0xB0. Nonresident. Min/Max Size: Variable.

The \$Bitmap Attribute is similar to the \$Bitmap Metafile, but in this case it only refers to the Indexes Entries present in the file. This allows NTFS to know which Indexes are free and which are not. The Attribute is also used in \$MFT, for the purpose of describing which File Records are in use, and which are free (either empty of that can be overwritten).

4.13.2 Structure

The structure of the \$Bitmap Attribute is exactly the same as the one for the \$Bitmap Metafile, except that in this case, each bit represents one VCN of Index Allocation.

4.14 \$REPARSE_POINT

4.14.1 Description

Type:0xC0. Nonresident. Max Size: 16384 bytes.

A Reparse Point is a collection of user meta-data, which is structured according to the specific application filter that stores it. It can be thought of as a special area of NTFS where a filesystem extension can store its data in a format that it can read. It is also fair to think of a Reparse Point as a way to store extended data which should not be placed in a Data Stream. For example, Microsoft's Backup Server tools use Reparse Points to store the data about each file on the filesystem. When the filesystem driver extension reads the Reparse Point Data, it finds the original file and opens it. NTFS Hard Links also use this same mechanism. This Attribute is very useful to

implement Symbolic links, to link a directory to another directory, volume, or link a file to another file.

4.14.2 Structure

The \$REPARSE_POINT Attribute has a generic structure for the identification of the type of data, plus two defined structures that Microsoft uses (which are actually the same). They are used by Volume (Hard) and Symbolic Links. The structure is generally like this:

```
Public Type REPARSE_POINT
    ReparseTag           As ReparseTags
    ReparseDataLength   As Integer
    Reserved             As Integer
    ReparseData ()      As Byte
End Type
```

Structure 4.14.2.1 - \$REPARSE_POINT Attribute

The bad news however is that the structure is only valid for Microsoft Tags. As a matter of fact, the Reparse Tag actually indicates if the Reparse Point was made by Microsoft or not. The Microsoft bit is the first bit in the Reparse Tag, so you can check for it by a simple “And &H80000000”, and check if the result is 0 or not. If it’s 0 (Not a Microsoft Tag), then an additional member must be added:

```
Public Type REPARSE_POINT_NONMS
    ReparseTag           As ReparseTags
    ReparseDataLength   As Integer
    Reserved             As Integer
    GUID                 As GUID
    ReparseData ()      As Byte
End Type
```

Structure 4.14.2.2 - \$REPARSE_POINT Attribute

The GUID represents the Filesystem Extension Manufacturer’s GUID that created this Reparse Point.

Additionally, another way to check if the Reparse Point is from Microsoft or not, is to verify against the following known Microsoft Tags, or perform an AND operation against them:

```
Public Enum ReparseTags
    NativeStorage = &H68000005
    NativeStorageRecovery = &H68000006
    SIS = &H68000007
    DistributedFileSystem = &H68000008
    VolumeMountPoint = &H88000003
    HSM = &HA8000004
    SymbolicLink = &HE8000000
End Enum
```

Enumeration 4.14.2.1 - Microsoft Reparse Tags

Finally, if the Reparse Point is a Symbolic Link or Volume Mount Point, the data can be read as the following structure:

```
Public Type NTFS_LINK
    SubstituteNameOffset As Long
    SubstituteNameLength As Long
    PrintNameOffset As Long
    PrintNameLength As Long
    PathBuffer() As Byte
End Type
```

Structure 4.14.2.3 - NTFS Link Reparse Point

4.15 \$EA_INFORMATION

4.15.1 Description

Type:0xD0. Nonresident. Min Size: 8 bytes. Max Size: 8 bytes.

In order to have compatibility with OS/2 (which Microsoft was co-developing with IBM before the two separated and NT was born), Microsoft allowed NTFS Files to contain “Extended Attributes”, or EAs. Just like Alternate Data Streams were used to support Macintosh Resource Forks, EAs are used for compatibility with HPFS

(the OS/2 filesystem) so that HPFS can properly read the information it needs from an NTFS volume.

There should be no need to programmatically access or read EAs, but in order to strive to be complete, this document includes it. \$EA_INFORMATION is the first Attribute dealing with EAs, and contains information about the size and number of EAs present.

4.15.2 Structure

The Attribute is an extremely simply 8-byte structure with the following typecast:

```
Public Type EA_INFORMATION
    EaPackedLength      As Integer
    EaCount             As Integer
    EaUnpackedLength    As Long
End Type
```

Structure 4.15.2.1 - \$EA_INFORMATION Structure

4.16 \$EA

4.16.1 Description

Type:0xE0. Nonresident. Max Size: 65535 bytes.

The \$EA Attribute is an Extended Attribute needed for HPFS Compatibility (see above). An EA is basically a collection of name and value pairs, and there can be many EAs in a File.

4.16.2 Structure

All EAs have the same structure, but the Name/Value pair can have a variable size which is defined in the structure below:

```

Public Type EA_ATTRIBUTE
    NextEntryOffset    As Long
    Flags              As Byte
    EaNameLength       As Byte
    EaValueLength      As Integer
    EaName()           As Byte
End Type

```

Structure 4.16.2.1 - \$EA Structure

The Value itself follows after the EaName, and has the size of EaValueLength. The NextEntry Offset points to the next EA, and indirectly the size of the current EA itself.

4.17 \$LOGGED_UTILITY_STREAM

4.17.1 Description

Type:0x100. Resident. Max Size: 65535 bytes.

This Attribute is used by EFS (Encrypting File System) when encrypting a file on an NTFS Volume. It is where EFS stores two important data that it needs to decrypt the file. One is the DDF (Data Decryption Field) and the other is the DRF (Data Recovery Field). In the DDF, the FEK (File Encryption Key) is encrypted with the users's public key, and in the DRF, the FEK is encrypted with the DRA (Data Recovery Agent) public key. Because it's encrypted with both keys, EFS can use the DRA Private Key of the system to decrypt an encrypted file if the user key is not available. This Attribute should not normally be accessed by any application, unless the Private Key is somehow known and the data needs to be recovered. More information about EFS is available in Chapter 5, Encryption.

4.17.2 Structure

The Structure of the \$LOGGED_UTILITY_STREAM will be available in a future revision of this document.

5.0 Advanced Concepts

5.1 VCNs and LCNs

Along this document, it has often been talked about VCNs (*Virtual Cluster Numbers*) and LCNs (*Logical Cluster Numbers*). NTFS uses both terminologies in certain structures, and it is important to be able to differentiate between them. LCNs are simply the normal cluster number found on a volume, which is sequentially arranged. The first bytes on the volume (the Boot Sector) is LCN 0, and every other cluster is sequential. The number of bytes in a cluster is of course defined in the Boot Sector, as it has been discussed before. As such, to get the offset of an LCN, simply multiply the LCN by the number of bytes/cluster.

The VCN is a completely different matter. It is called Virtual because its starting offset (VCN 0) depends on the file that contains it. VCNs are used in Nonresident Attributes to indicate where the *Data Runs* are located, in regards to VCN 0 (the first *Data Run*). Consequently, VCN 20 means 20 clusters after the first *Data Run*. To further understand this concept, please refer to section 5.2 below.

5.2 Data Runs

5.2.1 Definition

As has been previously said, NTFS holds Nonresident Attributes, usually containing data or indexes (such as \$DATA) according to streams called *Data Runs*. A *Data Run* simply indicates, using a compressed structure, which VCNs contain the next “Run”, or stream of bytes. By reading all the *Data Runs*, it is possible to have an image of where the file is contained on the disk, get the LCNs of the various runs, and use their offsets to read the data.

5.2.2 Structure

Because of the tightly packed structure of a *Data Run*, it is impractical to describe it as a regular VB Type. It consists of four main pieces of data:

<i>Nibble</i>	<i>Size</i>	<i>ID</i>	<i>Description</i>
0	1 nibble	F	Size of the Offset Field (O)
1	1 nibble	L	Size of the Length Field (RL)
2	2 nibbles * L	RL	Length Field
2 + (2*L)	2 nibbles * F	O	Offset Field

Table 5.2.2.8 - Nibble Table

A *nibble* is simply the position of a number. For example, in 345, the first *nibble* (0) is 3, and the last *nibble* (2) is 5. Because the table is abstract, and the structure can get complicated, some examples are needed to better understand the concept of *Data Runs*.

Simple Example

Assume the following *Data Run*

```
>628308017 18100788 822219237 50375234 0
```

First, convert the four long values to hexadecimal:

```
>25733831 01143234 310211E5 0300AA42 0
```

Next, convert this to little-endian

```
>31387325 34321401 E5110231 42AA0003 0
```

Now, split the numbers into bytes

```
>31 38 73 25 34 32 14 01 E5 11 02 31 42 AA 00 03 00
```

The *Data Run* is now arranged and ready to read. Starting by the first *nibble*:

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x3	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x38	Length Field
2 + (2*1) =4	2 nibbles * 3	0x342573	Offset Field

Table 5.2.2.2 - Nibble Table

It might seem tempting at first to write 0x732534 as the Offset Data, however it is important to remember that these values are in little-endian, as such they must be read backwards (12 34 becomes 0x3412; 72 25 32 becomes 0x342573).

Finally, by using the table above, write the information in a logical manner:

0x38 clusters starting at LCN 0x342573, or

56 clusters starting at LCN 3417459.

The next bytes must now be interpreted as a new run.

>32 14 01 E5 11 02 31 42 AA 00 03 00

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x3	Size of the Offset Field
1	1 nibble	0x2	Size of the Length Field
2	2 nibbles * 2	0x114	Length Field
2 + (2*2) =6	2 nibbles * 3	0x211E5	Offset Field

Table 5.2.2.3 - Nibble Table

The only thing to be aware here is that the size is now 2 bytes as well, so it must also be read in little-endian (0x114, not 0x1401).

Therefore, the next run has:

0x114 clusters starting at LCN 0x342573 + 0x211E5.

The crucial fact to notice above was that offsets are actually relative, so they must be added to the previous value. The next bytes are now:

>31 42 AA 00 03 00

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x3	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x42	Length Field
2 + (2*1) =4	2 nibbles * 3	0x300AA	Offset Field

Table 5.2.2.4 - Nibble Table

It is now easy to understand that the last run has:

0x42 clusters at LCN **0x342573 + 0x211E5 + 0x300AA.**

By adding all the offsets and tabulating the data, the *Data Runs* can be expressed as follows:

38:342573;114:363758;42:393802 or,

Data Run 1: 0x38 clusters starting at LCN 0x342573

Data Run 2: 0x114 clusters starting at LCN 0x363758

Data Run 3: 0x42 clusters starting at LCN 0x393802

Complex Example

Assume the following *Data Run*

>18886673 809505120 12911121

Convert to Hexadecimal, change to little endian, split bytes

>11 30 20 01 60 11 40 30 11 02 C5 00

The *Data Run* is now arranged and ready to read. Starting by the first *nibble*:

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x1	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x30	Length Field
2 + (2*1) =4	2 nibbles * 1	0x20	Offset Field

Table 5.2.2.5 - Nibble Table

Therefore, this run has:

0x30 clusters starting at LCN 0x20.

The next bytes are:

>01 60 11 40 30 11 20 D5 00

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x0	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x60	Length Field
2 + (2*0) =2	2 nibbles * 0	NULL	Offset Field

Table 5.2.2.6 - Nibble Table

Something immediately can be seen from this table. It would seem that there are 60 clusters which have no physical location on disk. How can this be? It turns out that the file we are actually reading is therefore a *Sparse File*, or could even be a *Compressed File*. For the sake of simplicity, let's assume it's Sparse. This run tells us that there are 60 "empty" clusters. The file size will report the space occupied by these non-existent clusters as part of the whole file size, hence the whole usefulness of a *Sparse File*. Assuming a cluster size of 4096 bytes, this means that 393216 bytes (96 * 4096) are logically (virtually) being taken up on the disk, but not physically. When the file needs to grow, the physical clusters will be filled up with the new data.

The next bytes are:

>11 40 30 11 02 C5 00

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x1	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x40	Length Field
2 + (2*1) =4	2 nibbles * 1	0x30	Offset Field

Table 5.2.2.8 - Nibble Table

Therefore there are:

0x40 clusters starting at LCN 0x20+0x30.

The next bytes are:

>11 02 C5 00

<i>Nibble</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
0	1 nibble	0x1	Size of the Offset Field
1	1 nibble	0x1	Size of the Length Field
2	2 nibbles * 1	0x2	Length Field
2 + (2*0) =4	2 nibbles * 1	0xC5	Offset Field

Table 5.2.2.8 - Nibble Table

Therefore it would be tempting to say that we have:

0x20 clusters starting at LCN 0x20+0x30+0xC5.

However, it must be mention that the values for each byte are actually signed integers. This means that any byte value over 0x80 is actually a negative number. As such, D5 is not 0xD5, but -0x3B.

Therefore, we have:

0x20 clusters starting at LCN 0x20+0x30-0x3B.

By adding all the offsets and tabulating the data, the *Data Runs* can be expressed as follows:

30:20;60:0;40:50,20:2 or,

Data Run 1: 0x30 clusters starting at LCN 0x20

Data Run 2: 0x60 clusters to be filled later.

Data Run 3: 0x40 clusters starting at LCN 0x50

Data Run 4: 0x2 clusters starting at LCN 0x15

Why are those 2 clusters marked at the end, instead of being put as the first *Data Run*? Remember that *Data Runs* are actually in order of the data they contain. As such, those last two clusters are actually the last 8192 bytes of the file (assuming 4096 bytes per cluster), located before the beginning. This indicates a severely fragmented file, and most probably volume.

5.3 Security Concepts

5.3.1 SIDs

5.3.1.1 Definition

A *SID* is exactly what its acronym stands for: a *Security Identifier*. It is the basic mechanism by which NT assigns a special ID to each User or Group on the system, and maps it to the corresponding Authority that owns it. S-1-5-21-646518322-1873620750-619646970-1110 is a typical SID. Any piece of information that requires to be linked to a user or group of the NT Security Subsystem needs to be linked to its SID. Furthermore, some SIDs belong specifically to built-in accounts or groups of NT, while others correspond to specific users on a specific computer on a specific domain. SIDs should never be allowed to duplicate (a big problem with cloning utilities).

5.3.1.2 Structure

SIDs don't have a fixed length, consequently their structure below is more of a architectural design than a structure that can be used in VB.

<i>ID</i>	<i>Size</i>	<i>Data</i>	<i>Description</i>
S	1 byte	"S"	Security String
R	1 byte	0x1	Revision (1.0)
A	6 bytes	See Below	NT Authority
SAs	Variable	See Below	NT Sub-Authorities

Table 5.3.1.2.1 SID Structure

A SID can only have one NT Authority, but an unlimited number of Sub-Authorities, therefore its size is always variable. Worse, the NT Authority, although it is a 6-byte number, is not completely written if not all 6 bytes are taken., making it also a variable-length field. NT has some predefined NT Authorities and Sub-Authorities, shown below:

<i>Identifier</i>	<i>Code</i>	<i>SID</i>
Null	0	S-1-0-0
World	0	S-1-1-0
Local	0	S-1-2-0
Creator Owner	0	S-1-3-0
Creator Group	1	S-1-3-1
Creator Owner Server	2	S-1-3-2
Creator Group Server	3	S-1-3-3
Non-unique	0	S-1-4-0
Dialup	1	S-1-5-1
Network	2	S-1-5-2
Batch	3	S-1-5-3
Interactive	4	S-1-5-4
Logon IDs	5	S-1-5-5
Service	6	S-1-5-6
Anonymous Logon	7	S-1-5-7
Proxy	8	S-1-5-8
Enterprise Controllers	9	S-1-5-9
Server Logon	9	S-1-5-9
Principal Self	10	S-1-5-10
Authenticated User	11	S-1-5-11
Restricted Code	12	S-1-5-12
Terminal Server	13	S-1-5-13
Local System	18	S-1-5-18
NT Non-Unique	21	S-1-5-21
Built-in Domain	32	S-1-5-32

Table 5.3.1.2.2 - Well Known SIDs

The following table also lists some well-known Sub-Authority RIDs (Relative Identifiers):

<i>Identifier</i>	<i>Code</i>
Admin	500
Guest	501
Kerberos Target	502
Admins (group)	512
Users (group)	513
Guests (group)	514
Computers (group)	516
Controllers (group)	516
Certificate Admins (group)	517
Schema Admins (group)	518
Enterprise Admins (group)	519
Policy Admins (group)	520
Admins (alias) (group)	544
Users (alias)	545
Guests (alias)	546
Power Users (alias)	547
Account Ops (alias)	548
System Ops (alias)	549
Print Ops (alias)	550
Backup Ops (alias)	551
Replicators (alias)	552
RAS Servers (alias)	553
Pre-WIN2K Compatibility (alias)	554
Remote Desktop Users (alias)	555
Network Config Ops (alias)	556

Table 5.3.1.2.3 - Well Known RIDs

An example should help in clearing things out:

> S-1-5-32-544

<i>ID</i>	<i>Size</i>	<i>Data</i>	<i>Meaning</i>
S	1 byte	"S"	Security String
R	1 byte	1	Revision (1.0)
A	1 byte	5	NT Logon IDs
SA1	2 bytes	32	NT Built-in Domain
SA2	3 bytes	544	NT Admins

Therefore this is the SID of the Local Administrators on any machine. As a final note, there are many security APIs that can use SID strings to get more information about the user. It would be useful to take a look at the MSDN Security Documentation for more information on the subject.

5.3.2 ACLs

5.3.2.1 Definition

5.3.2.2 Structure

5.3.3 ACEs

5.3.2.1 Definition

5.3.2.2 Structure

5.4 Indexes

5.5 Sparse Files

5.6 Encryption

5.7 Compression

5.8 USNs