File System Data Recovery

- Overview
- Unallocated Cluster
- Cluster Slack
- RAM Slack
 Data Fragm
- Data FragmentationMaster File Table
- References

Overview

There are numerous areas of a file system where vital evidence may be located and possibly recovered. The following article looks at some common file system areas.

Unallocated Cluster

One of the most neglected areas in forensic data recovery is that of unallocated clusters. When a file on a hard drive containing data is deleted, the data belonging to the file remains on the disk. This area is known as unallocated clusters or unallocated space. The data will remain there until it is overwritten by another file. Any data found in this area at one time belonged to a file which has since been deleted.

Cluster Slack

Another area which is frequently neglected is that of file or cluster slack. Files are created in varying lengths depending on their content. NTFS (New Technologies File System) and FAT (File Allocation Table) file systems store files in fixed length blocks of data commonly referred to as clusters. It is rare for file sizes to exactly match the size of one or multiple clusters perfectly. The data storage space that exists from the end of the last sector of the file to the end of the last cluster assigned to the file is called file or cluster slack. As files are deleted and created, it is possible that fragments from the file previously located in that cluster can be recovered. For example, a cluster which is 4096 bytes in size (8 sectors of 512 bytes each) may hold a file as small as a single byte. This would result in 3584 bytes of cluster slack. Across a single hard disk, this could amount to a significant amount of data which is potentially full of evidence.

RAM Slack

From the previous example, there is also an additional 511 bytes which are not referenced. This is located outside the logical file and is known as RAM slack. This data is located in the last sector allocated to a file from the end of the logical file to the end of that sector. As the operating system must write data in sector aligned blocks, the write buffer may need to be filled with data to align it to the correct size. In earlier version of the Microsoft Windows operating system (pre Windows 95B), this padding data was randomly copied from memory; later versions fill this area with 0x00. From a forensic computing point of view, in older versions of Microsoft Windows, RAM slack could be filled with sensitive information associated with the use of the computer. Figure 1 shows an example file which is 2,248 bytes in size which would result in 1,536 bytes of cluster slack and 312 bytes of RAM slack:



Figure 1

/!\

Some record based files (usually binary files where the records contain length markers) can also have what is known as Record Slack. Record Slack is the data area immediately after the end of a live record (or block of data) to the end of the allocated block or the start of next record.

Data Fragmentation

When recovering data from hard disks (or disk images) it is important to understand how fragmentation impacts upon the process. The NTFS file system is very bad at avoiding fragmentation on some files, partly due to its allocation strategy of intentionally placing gaps between files; which is good if those files expand, but bad if they don't. Under ideal conditions, file system read and write transfer performance is maximised when files are contiguous on the disk. This means that all of the data in each file would be located in consecutive clusters or blocks within the volume. Contiguous storage improves performance by reducing unnecessary seek motions that are required when data is located in many different places. When files are broken into many pieces they are said to be fragmented. The NTFS file system handles the storage of files and directories in a very different way than the FAT file system does. FAT is a very simple and relatively "unintelligent" file system that pays little attention to how much fragmentation will result from how it operates. In contrast, NTFS is smarter about how it manages the storage of data. For example, NTFS reserves space for the expansion of the Master File Table, reducing fragmentation of its structures. In fact, due to their complexity, NTFS volumes suffer from a variety of different types of fragmentation. Unlike FAT, where a simple cluster allocation system is used, NTFS uses the Master File Table and a combination of resident and non-resident attributes to store files. Due to the flexible way that data is stored, and that additional data storage areas are added as needed, the result can be pieces of data spread out over the volume, particularly when small files grow into large ones.

Remember that while NTFS has a much better design than FAT, at its core it does still store data in clusters. The addition and removal of data storage extents causes much of the fragmentation of files and directories. As the MFT grows, it itself can become fragmented, reducing performance further. From a recovery point of view, any data from a deleted file which crosses a cluster boundary where the clusters are not stored contiguously on the disk is very difficult to recover. Figure 2 shows a fragmented INDEX.DAT file from Internet Explorer:



Figure 2

As you can see towards the end of this file representation, the data stored in cluster 327 and 330 is not contiguous. If we identified the start of a record at the end of cluster 327 and the data crossed the cluster boundary into 330, sector based recovery would expect the record to be in cluster 328. In this case, the record would likely be recovered in a corrupted state as it would contain JPEG data as well as data from a URL record.

Master File Table

The NTFS file system contains a file called the master file table, or MFT. There is at least one entry in the MFT for every file on an NTFS file system volume, including the MFT itself. All information about a file, including its size, time and date stamps, permissions, and data content, is stored either in MFT entries, or in space outside the MFT that is described by MFT entries.

As files are added to an NTFS file system volume, more entries are added to the MFT and the MFT increases in size. When files are deleted from an NTFS file system volume, their MFT entries are marked as free and may be reused. However, disk space that has been allocated for these entries is not reallocated, and the size of the MFT does not decrease.

The NTFS file system reserves space for the MFT to keep the MFT as contiguous as possible as it grows. The space reserved by the NTFS file system for the MFT in each volume is called the MFT zone. Space for file and directories are also allocated from this space, but only after all of the volume space outside of the MFT zone has been allocated.

Depending on the average file size and other variables, either the reserved MFT zone or the unreserved space on the disk may be allocated first as the disk fills to capacity. Volumes with a small number of relatively large files will allocate the unreserved space first, while volumes with a large number of relatively small files allocate the MFT zone first. In either case, fragmentation of the MFT starts to take place when one region or the other becomes fully allocated. If the unreserved space is completely allocated, space for user files and directories will be allocated from the MFT zone. If the MFT zone is completely allocated, space for new MFT entries will be allocated from the unreserved space.

The default MFT zone is calculated and reserved by the system when it mounts the volume, and is based on volume size. You can increase the MFT zone by means of the registry entry detailed in Microsoft Knowledge Base Article 174619, but you cannot make the default MFT zone smaller than what is calculated. Increasing the MFT zone does not decrease the disk space that users can use for data files.

References

• MSDN: Master File Table