

Following is a slightly edited copy of advice to a customer regarding data recovery. It was written with regard to a corrupted hard drive but the same advice and methods are valid for a CompactFlash card. Note that corruption on the CompactFlash card is usually due to power loss or application crash while files are opened, resulting in directory errors, because the files were not closed and the directory file not updated. If PicoZOOM is enabled then the corruption can be worse because the directory is updated less often. In this scenario (that files are not closed properly) it is likely that some of the data recorded before the failures will probably still be there. Sometimes the corrupted directory will allow the OS to overwrite some portions of some files. It is less likely that the media itself is faulty.

1. First make a copy of the media.

My experience with data recovery from a system that experienced failures has taught me to be very cautious to not do any kind of writing or repairs on the failed card or disk. Some well-meaning automated recovery programs may overwrite critical sections of the file system and obscure key clues that might have allowed for a manual recovery. The more important the data, the more important it is to first make a true bit-copy of as much of the damaged media as possible.

I use a PC and Runtime Software's "DiskExplorer for FAT" to create a bit-image of the drive as a file on the PC, and this file can then later be used by DiskExplorer or Runtime's "GetDataBack for FAT" to attempt recovery. You can also make multiple copies of this image file, which is important because the early recovery attempts are sometimes just a learning experience, from which you may well want to start over. Both of these programs are inexpensive and they offer free trials which I believe are for 30-days with no limitations.

Runtime (www.runtime.org) also publishes technical information on the mechanics of file recovery.

Notes:

Both DE and GDB allow you to choose either the physical (entirety) or logical drive (volume/partition). The physical drive is what you want to make a backup of.

Remove the damaged drive from the PC after backing up. Destinations and sources are likely to get confusing during a prolonged recovery session and you don't want to inadvertently write to the damaged drive.

2. Zero length files.

A file's length and starting-cluster-number are recorded as zero in the file's directory entry when the file is created and these are only updated when the file is closed. A file's starting-cluster-number (in the FAT table) is recorded in the file's control block after the first write request and does not change. The cluster number both identifies a block of sectors in the data area and forms an index into the FAT that contains the next cluster number for data storage; a singly linked list. Each time a cluster fills (typ 32KB for big drives) and new cluster is allocated (sequentially after a format, by search after deletions), the FAT is updated, and the sector containing that FAT entry is written. If PicoZOOM is not enabled, that write occurs immediately, otherwise its buffered and deferred until that cache sector is needed.

If an ongoing recording ends without an orderly file close, the file's directory entry length and starting cluster number will remain what they were when the file was opened. If PicoZOOM was enabled but flushed or released, up to 2 sectors (32 cluster numbers representing 1MB of data for a large drive) may

not be recorded and these clusters will be reallocated by whatever file requests write storage and their original data may be overwritten. If PicoZOOM was not enabled, all allocated clusters and associated data will have been written. In either case, a directory entry with zero file length and zero starting cluster has no data in the data region.

The DIR command and directory functions return length based solely on the files directory entry length so a never closed file will show up a zero length regardless of the actual data recorded. Similarly, a never closed file will have a zero starting cluster and will not be associated with the recorded data. If you request the verbose version of DIR from PicoDOS with DIR /V, you will get summary drive space statistics that use first directory file length and then used cluster counts. If you subtract "KB free" from "KB total disk space" you will know how much data was recorded. If you subtract "KB allocated" from that data recorded number, you will have some idea of how much other data was recorded that is not associated with directory listings.

3. CHKDSK and lost clusters.

This fragility is one of the weaknesses of the FAT file system and was recognized very early on prompting the creation of CHKDSK for DOS and SCANDISK for Windows. These programs analyze the directory and FAT tables looking for inconsistencies in the linked list, and can generally find the lost chains and create new directory entries that simply point to the orphaned starting clusters and declaring a size equal to the number of clusters in the linked list times the cluster size.

Once you have a safe bit-copy backup, you can move this to another 60GB or larger drive using the same Runtime tools, and then perform a CHKDSK /F as a first pass at data recovery. If successful, you may be able to manually associate the file with zero length to the batch of FileXXXX.chk that are created by sequentially matching the names. This will be most successful if your program generated files in simple and deterministic fashions. It will also be helpful for assigning confidence levels if your data has embedded interpretable time or sequence information.

4. GetDataBack or other third party recovery tool.

I have not had the best results from using GetDataBack with the relatively odd (to a PC) types of binary data associated with data acquisition systems.

Many tools have heuristics and tables to help identify standard files like Word docs, PDFs, JPGs, and such, and get hopelessly confused with our data.

I have sometimes found it useful for confirming that a CHKDSK or manual extraction session has done a reasonable job. If you use this, you may want to attempt runs on both a copy of the original drive image and a copy of the post-CHKDSK drive image.

5. Data recovery companies.

They seem to be everywhere, but I've never used one or known anyone who has, so I can't offer any recommendations. If you follow the bit-copy path and just send them a drive with a copy, you have nothing to lose but the fee. I would guess that a good outfit would be able to tell you if they could handle drives from a non-PC system and especially if you could provide them with what you would expect to find for files and lengths on system that did not experience failures.