CPR TOOLS
Media Research & Recovery

Data Recovery and RAID Drives

This paper addresses the recovery of information (user files) from hard drives which have been used in RAID systems. There is a history of successful data recovery efforts when working with drives from RAID systems. RAID systems are unique in that multiple drives can be used to store the data and as the 'R' in the acronym states they are 'redundant'. This redundancy not only helps the system administrator or IT staff keep the systems running when a drive fails but also aids the data recovery engineer in recovery efforts when multiple drives fail.

Simple mirror RAID (RAID 1) recoveries are far less complex than any other as the file structure(s) are intact on a complete physical volume. This RAID mechanism simply makes a runtime copy of the primary drives' data onto the secondary drive. This keeps a complete 'backup' of all information; losing one drive in this scenario is not an issue because the other drive contains all of the same information.

We are more interested in RAID topologies which use spanning across multiple physical volumes. The most popular of these is RAID 5 as shown in Figure 1.

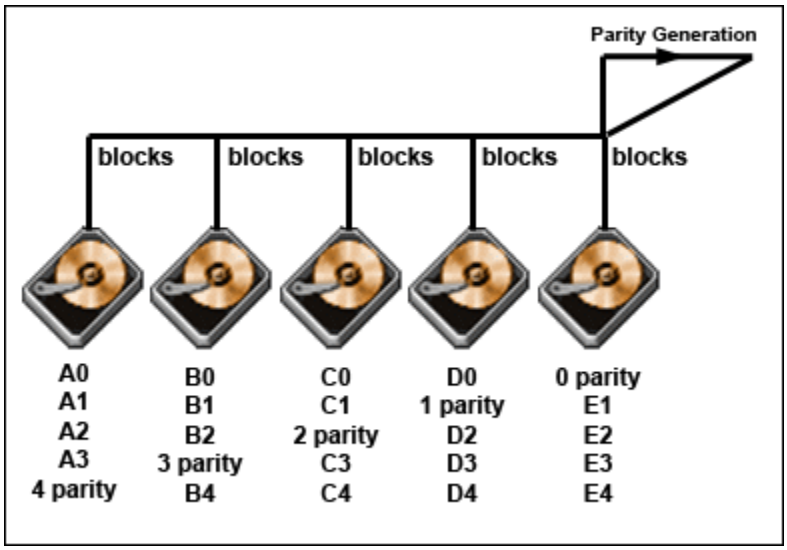

Figure 1 - RAID 5 is a very popular RAID topology

Systems which use multiple drives in RAID topology other than RAID 1 employ a data storage method called 'striping'. If one drive in a striped RAID fails, the administrator can replace the failed drive with a new one; using stored parity information, the missing data will be rebuilt automatically.

If more than one drive fails a data recovery specialist is typically called in to retrieve the lost data. Retrieving data from multiple drive arrays does raise the bar in terms of being able to recover data from a single participant drive; in these cases, the definition of 'success' is worthy of discussion.

In arrays which span multiple physical volumes the key factor in determining successful recovery of user data is the stripe size. By default, most RAID systems utilize a stripe size of 64K. A common variant to this default is a 128K stripe size which, in certain hardware configurations, may produce better throughput and therefore is chosen by system administrators. Given that the stripe size defines

contiguous storage block on a physical volume, files which are smaller in size than the stripe size are easily recovered on any single physical volume which had participated in a multi-drive RAID.

For the sake of clarity, let's quickly break down a typical hard drive. The smallest informational component on a hard drive (for the sake of this paper) is a sector. A sector holds 512 bytes of information. While this is a small amount of data given ready comparisons to hard drives which store data measured in Gigabytes, Figure 2 shows just how much information can be stored in a single sector.

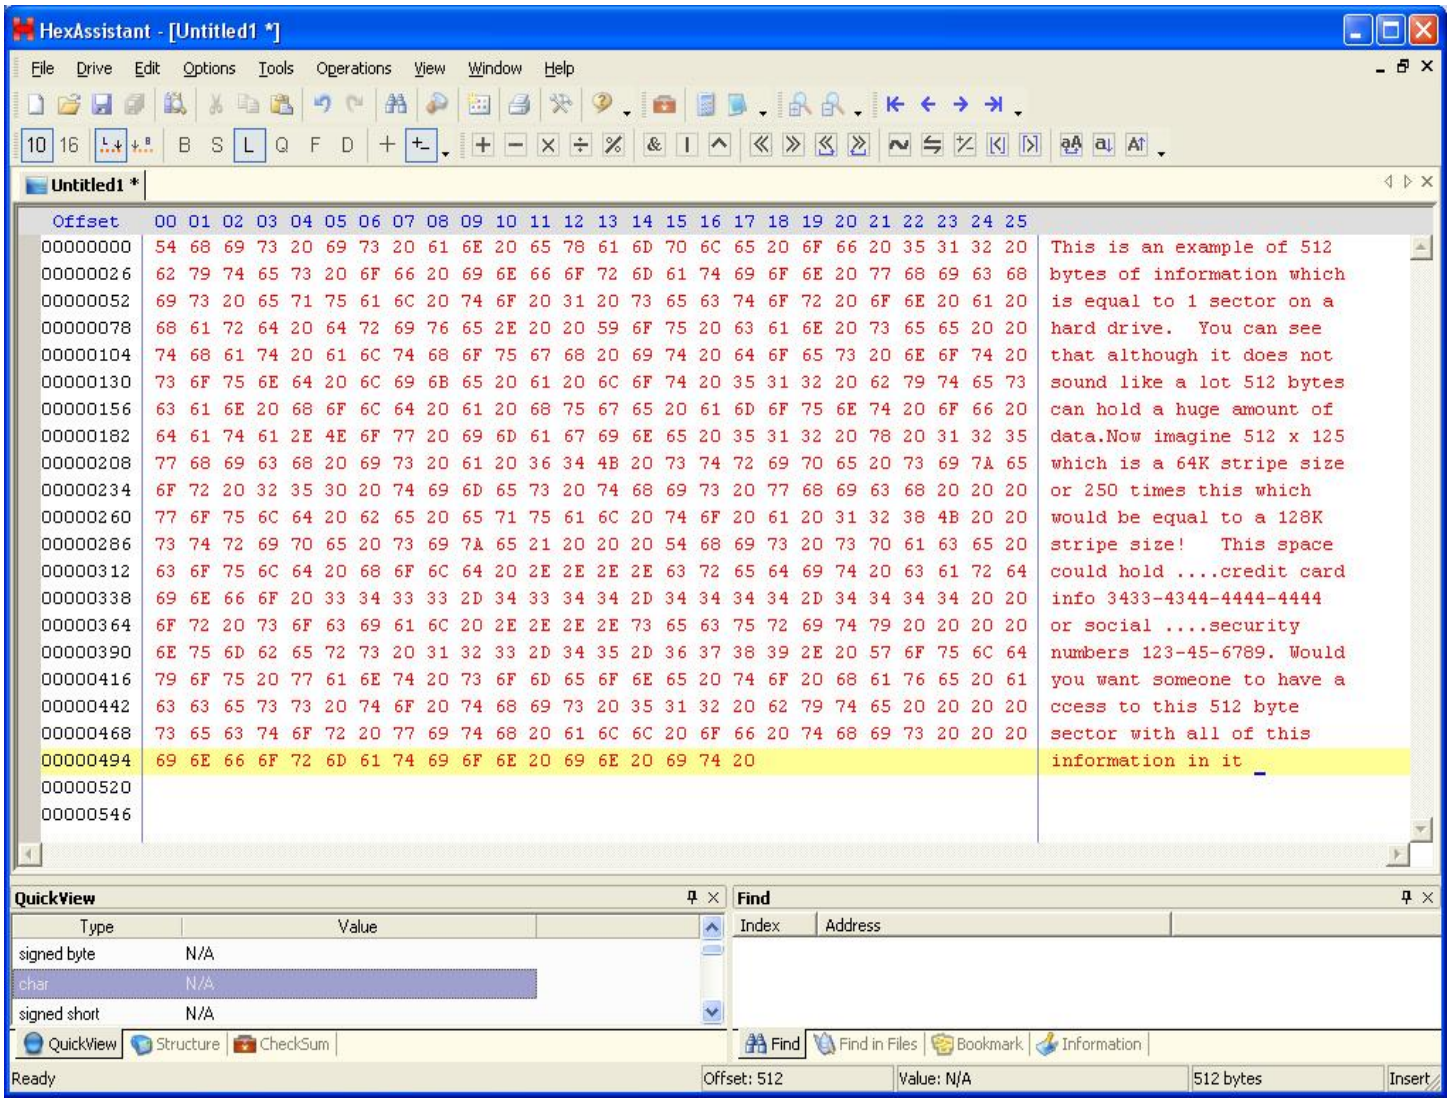

Figure 2 - 512 bytes of data - one sector holds all of this information

CPR Tools recently worked with CBS news affiliate WINK-TV in Ft Myers Florida[1]. The news agency purchased ten (10) used SCSI drives from a vendor on eBay and brought them to us to determine if any data was present. All of these drives had participated in multi-drive RAID systems. All of these drives

[1]The WINK-TV news segment can be viewed here http://www.cprtools.net/securityinfo.php

contained recoverable information from companies like GIANT Foods, Sears and a financial company purchased by Bank of America.

None of the purchased drives were siblings in any one RAID system; from these single drives we were able to recover over four thousand (4,000) credit card numbers, hundreds of prescriptions including physician and patient names, medicine prescribed, underlying medical reason for the prescription, phone numbers and addresses of both patients and doctors and a wealth of other information.

When considering data security and adequately protecting information, there are factors that must be kept in mind regarding drives which have participated in multi-drive RAID systems. These include:

1. Sector size
    a. As shown in our recent work with WINK-TV news, individual sectors were found to contain between 32 and 38 complete credit card numbers.
2. Stripe size
    a. Documents whose size are less than or equal to the RAID stripe size will be easily recovered in their entirety.

Average file size is also worth mentioning. The University of Toronto, in a study[2], found that while Microsoft Word and Excel documents may range from very small to very large file sizes, the bulk of the larger files were taken up by graphs, embedded images and formatting.

> *Interestingly, these results are similar to the findings of Cunha* et. al. *where the size of HTML-based Web documents was found to follow the power-law distribution. However, while Cunha* et. al. *found that most HTML documents are quite small (usually between 256 and 512 bytes), Office documents tend to be much larger. Common sizes of Word and Excel documents size range from 12 KB to 24 KB, and common PowerPoint documents range from 48 KB to 80 KB.*

Comparing the results from the study to our 64K and 128K stripe size we see that all of the Microsoft™ Word, Excel, Powerpoint and HTML files fall within the stripe size and therefore could be recovered from a single drive removed from a RAID system. Raw data such as text almost always falls within the < 64K range.

Summary

Systems administrators, corporate or organizational management and IT policy makers should be cognizant of the recoverability of data from individual drives which have participated in RAID systems. Protecting data is often legislatively mandated and can carry severe penalties for security breaches which expose private information. Even when legislation and the threat of penalties is not a driving force, organizations should be aware of potential exposures and move to ensure that they are acting responsibly where data security is concerned.

Of the 300 million data breaches that have occurred since 2005, it has been said that at least half were preventable using the proper tools, verification methods and of course through educating IT managers and employees as well as users as a whole.

Questions regarding this document can be sent to John Benkert by email to john.benkert@cprtools.net.

[2] http://www.cs.toronto.edu/~delara/papers/usenix_win2000/html/node9.html