

www.FreeRaidRecovery.com

Practical issues in DIY RAID Recovery

Based on years of technical support experience



© 2012

www.FreeRaidRecovery.com

This guide is provided to supplement our ReclaiMe Free RAID Recovery software, available at www.FreeRaidRecovery.com. In this guide, we covered the most typical issues our customers have to deal with, based on several years of technical support experience. Reading this guide thoroughly before recovery will help you avoid all the common pitfalls in RAID recovery. The guide is short enough and is best read in its entirety, start to end.

Contents

Equipment and preparation.....	3
System requirements	3
Requirements from software	3
Hardware requirements	4
How the member disks are connected to the PC	4
USB and SATA (direct connection to a motherboard).....	4
Disk image files.....	5
Mixed cases	6
RAID recovery output.....	6
Tips on the array type detection.....	6
Using Disk Management	6
Using md-raid data in case of a NAS device	7
Disk set for the analysis	8
Hot spare in RAID 5	8
RAID 5 with all the disks available cannot be recovered.....	8
Exclude failed disk off fault tolerant array	8
ReclaiMe Free RAID Recovery messages.....	9
Expected analysis time.....	9
Copying.....	10
Array image files and OSFMount.....	10
Disks and reboots after copying	10
Partition table and TestDisk	10

Equipment and preparation

Before you start to recover RAID at home, you need to understand clearly what requirements for the equipment are. If the system on which you plan to recover RAID doesn't meet the basic requirements, you need either to upgrade the existing system or to recover the array using another computer.

Also you should decide in advance what the recovery result you want. Maybe it is enough for you just to recover the array parameters (and then try to assemble the original array) or else you wish to copy data from the array on a file-by-file basis. Depending on what you need additional configuration of the system may be required, for example, you have to add disks to copy data from the array.

System requirements

System requirements can be divided into two groups:

- RAID recovery program imposes memory and processor power (CPU) requirements.
- The disks of your array impose power (PSU) and number of SATA ports requirements.

Software requirements

Memory and CPU

RAID recovery program processes large amount of data doing various calculations which require significant CPU resources. So it is better to choose more powerful computer to recover array configurations and save on the time.

RAID recovery program uses statistical approach for data analysis. It means that RAID recovery program reads data from the array member disks and then processes this data using special algorithms, requiring CPU power. Until RAID recovery program accumulates analysis results enough for the final decision, it stores the accumulated data in a computer memory. Therefore, the more memory available, the better for RAID recovery software. Despite the certain resource requirements depend on the particular array, it may be said that:

- laptop Intel Atom with 512 MB memory - too weak a system,
- a dual core processor and 4 to 8 GB memory - a fair choice,
- a six core server processor and 128 GB memory - quite an overkill.

Reboots

From RAID recovery point of view, reboots are bad since they may change the disk paths. If you redo RAID recovery on a PC after reboot, be careful when selecting disks for the analysis because their names (disk paths) may not match the previous ones. It

should be noted that changes of disk paths due to a system reboot are typical for disks connected via USB.

One more common case when reboot spoils the recovery is when it occurs between the recovery of the array parameters and the extraction of the data from the array using these parameters. In this case file describing the array configuration doesn't match the true configuration any longer for disk order has changed due to a reboot. If this happens, you should rerun RAID recovery to get the up-to-date array configuration.

A system reboot can be both planned and not. Planned reboots include those cases when you change computer hardware configuration in some way, for example bringing it into line with RAID recovery hardware requirements. Non-planned reboots take place for example when a system reboots after regular Windows update. Thus, it is recommended to turn off the automatic Windows Updates.

Hardware requirements

Power Supply Unit (PSU)

In case of a hardware RAID and sometimes for a software RAID to recover array configuration you need to get the disks out of the RAID device. Once you pull the disks out of the NAS or disconnect them from the RAID controller, you should connect them to the PC on which you plan to recover array configuration. At this point you should be aware that if you bring new disks to the old system, the power of the existing PSU may be insufficient. Sometimes you know about this immediately since the system just would not work. However, there is an intermediate state when there are no evident error messages, but the disks do not work properly resulting in significant analysis speed decrease.

How the member disks are connected to the PC

When selecting how the disks will be connected to the PC on which you are going to carry out the recovery, you need to realize what performance you can expect in each case.

USB and SATA (direct connection to a motherboard)

To get the maximum performance in RAID recovery, it is recommended to connect the disks to the motherboard ports directly. However, it is often not possible to set up the system as recommended (e.g. not enough ports), so the disks are connected via USB. Let's take a look at what the average read speed you get using different ways of connection.

Speed estimation

The disk read speed involves the following:

- the speed at which disk itself returns the data,
- the transfer speed of the data.

On the average, the read speed of a regular hard drive is about 60 MB/sec. The speed at which the data is transferred to the RAID recovery program depends on the connection type. The throughput of a SATA connection is greater than the average disk read speed. This is not always true for the USB protocols. There are several USB protocols, each of which has its own transfer speed.

USB 1.1 - the slowest protocol with an average speed of about 1.5 MB/sec, which is absolutely unacceptable for RAID recovery.

USB 2.0 - the most common USB protocol which provides transfer speed of about 10-15 MB/sec. This is a minimum acceptable speed for the recovery.

USB 3.0. Although the throughput of the protocol is up to 600 MB/sec, the overall read speed will be the same as in case of a SATA connection, specifically about 60 MB/sec, since it will be limited by the disk read speed. Having USB 3.0 is great if you managed to make it work. In real life configurations with many disks and especially hubs often fall back to USB 2.0.

Thus, if you have an array of four 2 TB disks which are connected via USB 2.0, in the worst case it will take eight days to complete the scan.

$$\frac{4 * 2\,000\,000}{10} / 3\,600 * 24 \sim 8 \text{ days}$$

If you use a SATA connection, you recover the array configuration five times faster, in a couple of days.

Bad sector handling

You should know that the implementations of SATA-USB convertors handle bad sectors worse than a plain SATA port. This sometimes leads to the data recovery software hanging when attempting to read data from a bad sector. You do not have such a problem with the disks connected to the motherboard ports, because when encountering a bad sector, the disks are still capable to return control to the program without causing it to hang.

Disk image files

When placing disk image files onto physical disks, try to put one disk image file per physical disk. Doing so gets you the maximum possible read speed. If you place several

disk image files onto one physical disk, you lose in read speed significantly due to seek time. Read speed in case of the "one disk image file per physical disk" placement is ten times greater than if there are two disk image files on one physical disk.

Mixed cases

If you connect the disks using several ways, for example you connect some disks via USB and others via SATA, the read speed would be determined by the slowest disk.

- SATA only - something like 40-50 MB/sec,
- mixed SATA and USB gives something like 15-20 MB/sec,
- USB only - expect something like 7-10 MB/sec.

Note: All of the above implies USB 2.0.

RAID recovery output

Before you start to recover the array configuration you should decide in advance what you want to do with the recovered parameters. If you are going to write an array copy to a disk and then try to access the data using Linux, you need storage with the same size as useful capacity of your array.

However, if you plan to recover files, you need storage with the same size as estimated size of the files you want to recover. If you just want to get a description of your array configuration and continue the recovery process later, keep in mind that if by the time you decide to recover data from the array the system was rebooted, the description of the array may have become incorrect.

Tips on the array type detection

In our practice one of the most common errors is when people specify wrong array type. There are several easy ways to check what the array type might be and what the type definitely cannot be on the disks.

Using Disk Management

Disk Management displays the disks either with partitions or not, when all the disk space is marked as unallocated. Partition sizes on the array member disks can be greater than the physical disk sizes (except for a RAID 1) since sizes of the partitions correspond to the original array capacity.

Typically, only one disk from the set will have partitions. However, you may see that two disks contain partitions in case of mirrored RAIDs (RAID 1 or RAID 10), or when a RAID 5 parity computes as an exact copy of the partition table.. Sometimes you can see all the drives without partitions and then you are not able to determine the array type. This means that the partition table is not in the right place. Such a situation can occur

when a controller places its own metadata before the partition table resulting in displacement of the partition table.

On the other hand, it also happens that all the disks are displayed with partitions but Windows doesn't show the filesystem type for them. This means you deal with a Linux software RAID (*md-raid*). Most likely, you do not need ReclaiMe Free RAID Recovery. Instead of this try ReclaiMe File Recovery (www.ReclaiMe.com) to get access to the array data.

RAID 0

In Disk Management you should see only one disk with partitions while all the remaining disks should be marked as unallocated.

RAID 5 and RAID 10

There will be one or two disks with partitions in Disk Management depending on what the data was on the disks and whether there is a disk missing.

RAID 1

In case of a RAID 1 you should see only two drives and both with identical sets of partitions in Disk Management. Most likely you do not need any recovery at all.

Checking the array type in Disk Management

So, what to do to check that your idea about the array type is consistent with the data Disk Management provides.

1. If you have all the member disks (there are no missing disks) and Disk Management shows
 - one disk with partition(s) and the remaining as unallocated, you deal with RAID 5 or RAID 0.
 - two disks with partitions and the remaining as unallocated, you have RAID 5 or RAID 10.
2. If one member disk is missing,
 - RAID 0 is not considered because it is not a fault tolerant array and therefore cannot survive a failure of one member disk.
 - it is impossible to tell RAID 5 from RAID 10.
3. If all the disks belonging to the array are shown with partitions, most likely you deal with Linux *md-raid* and therefore you do not need RAID recovery. Try ReclaiMe File Recovery (www.ReclaiMe.com) instead.

Using *md-raid* data in case of a NAS device

If the disk set came from a NAS device, you can try to access the data without RAID recovery. Just get ReclaiMe File Recovery and start it having all the member disks connected separately to a PC. If you are lucky and *md-raid* data is not damaged too

badly, ReclaiMe File Recovery assembles the array and provides access to the array data.

Disk set for the analysis

It is often believed that the more disks from the array will be used in RAID recovery, the better. In practice, it is not always true. For example, if a hot spare disk is included in the disk set of a RAID 5 array, RAID recovery software is not able to solve the case.

Surely, to recover a RAID 0 array, you should have all the drives; otherwise, you cannot recover data from a RAID 0. In case of fault tolerant arrays, you have certain freedom of choice not to include one of the disks in the analysis.

Hot spare in a RAID 5

When analyzing a RAID 5, always exclude a hot spare disk, if any. If you do not know definitely whether there is a hot spare disk, firstly try including all the disks. If you cannot get a solution, try to identify a hot spare disk, for example looking for a disk somewhat different from the others or just excluding disks one by one.

If all the member disks are present, but ReclaiMe Free RAID Recovery displays a warning message "*It was determined that the RAID 5 set is missing a drive*", exclude the disks one by one until you find a disk set for which this message is not shown.

If one member disk is missing, the trick does not work and you need to do RAID recovery as many times as there are member disks.

RAID 5 with all the disks available cannot be recovered

If you have a RAID 5 with all the disks available but you cannot recover its configuration, try excluding one of the disks. If one of the disks contains data that is not consistent with data on the other disks, it will be difficult to obtain the correct solution. You can face such a situation, for example if you have replaced a failed disk with a hot spare, but rebuild was not completed. Since ReclaiMe Free RAID Recovery can handle a RAID 5 with one disk missing, it is better to just exclude such a suspicious disk from the analysis. First pay attention to a disk which:

- has been replaced recently,
- has been replaced automatically with a hot spare,
- differs in some way from the others, for example by size, vendor, or model.

Exclude a failed disk off a fault tolerant array

Sometimes ReclaiMe Free RAID Recovery works too slowly so you just cannot wait for the solution. Often it means that one of the disks is physically damaged. If you have a fault tolerant array with all the disks available you can try to exclude this

malfunctioning disk and redo the recovery. If you cannot detect which disk has a problem, exclude the disks one by one and see what it brings.

ReclaiMe Free RAID Recovery messages

Sometimes ReclaiMe Free RAID Recovery gives messages which should be considered thoroughly.

"Looks like wrong disks are selected". This message is displayed when there is a big chance that not all the selected disks belong to the same array, for example if you specified the disks with very different sizes.

"There were no mirrored pairs of drives detected". You may encounter such a message only when you recover a RAID 10. Typically, this message means that the disks do not belong to a RAID 10 and you should specify another RAID type (RAID 5 or RAID 0) or review the disk set.

"It was determined that the RAID 5 set is missing a drive". Such a message is shown when ReclaiMe Free RAID Recovery is not able to find enough samples in which parity blocks are consistent with data blocks. You see this message if:

- You have a RAID 5 with one disk missing, and the message is just confirming this fact. Obviously, no action is required in this case.
- There is a hot spare or other non-array member disk among the disks specified for the RAID 5 analysis.
- Specified disks do not belong to a RAID 5.

If you get one of these two messages above, review the specified array configuration once again.

Expected analysis time

Expected analysis time depends on what and how much data was stored on the array.

- Less filled arrays are recovered slower than more filled.
- If there are a lot of small files on the array, expect quicker recovery. If just only a few large files were stored on the RAID, the recovery will take more time.
- If there is a lot of high entropy data on the array such as movies, music or compressed files, count on that the recovery would be slower. Otherwise, with low entropy data like text or database files, you get the result faster.

Depending on the data stored on the array, analysis time may decrease. Sometimes it may happen that a solution is achieved earlier than all the array capacity is read. In the

worst case it is needed to read all the data from all the disks and here analysis time depends on the disk read speed, estimation of which was given in the "speed estimation" paragraph.

It should be noted that analysis time may differ on the same set of the disks due to ReclaiMe Free RAID Recovery algorithm specifics. The program reads the data randomly, producing a new sequence each time it is started.

Copying

Array image files and OSFMount

ReclaiMe Free RAID Recovery can create array image files. We are often asked what to do with these images next.

To read data off such an image file, you can try, for example, OSFMount (<http://www.osforensics.com/tools/mount-disk-images.html>). This freeware is capable of mounting a disk image file as if it were a hard drive, thus providing access to the data.

Disks and reboots after copying

If you select the *Write array to disk* option in ReclaiMe Free RAID Recovery, then once the process completes, you get a message telling that you need to reboot the system. In 9 cases out of 10 you do not need to reboot as you ascertain the results immediately - either you can access the data or see that the data is not readily available. In the remaining one case you can't see the data because the changes of the partition table have not taken effect and therefore you need to reboot the system.

Partition table and TestDisk

If you used the *Write array to disk* option of ReclaiMe Free RAID Recovery and after the reboot you still cannot access the data, try correcting the partition table using free TestDisk utility (www.CgSecurity.com). This situation often occurs with a software RAID on which the array metadata is placed before the user data. When scanning disk for the partitions, an operating system is not able to find the partition table due to such a displacement and therefore it can't provide access to the data.

Sometimes, hardware controllers place the metadata before the array data causing the same problems. As in case of a software RAID, try TestDisk to correct the partition table.