

RAID Recovery course

Lesson 2 – RAID levels

Practice part

Task 1 - Determining storage capacity needed to copy the recovered data

- a) A client brought four 2 TB disks from a RAID10. How much disk space you should have to copy the data off the RAID?

Let's draw a 4-disk RAID10:

Disk 1	Disk 2	Disk 3	Disk 4
1	2	1	2
3	4	3	4
5	6	5	6

We see that every data block is stored in two copies which are located on the array member disks. Therefore, exactly the half of the array capacity is used to store "useful" data while another half is used for redundancy. Thus, to copy data from a RAID10 of 4 disks, 2TB each, we need $4/2 * 2 \text{ TB} = 4 \text{ TB}$ disk space.

- b) A client brought a 5-disk RAID6, the capacity of each disk is 500 GB. How much disk space you should have to copy the data off the RAID?

A RAID 6 of 5 disks looks like:

Disk 1	Disk 2	Disk 3	Disk 4	Disk 5
Q	1	2	3	P
4	5	6	P	Q
7	8	P	Q	9
10	P	Q	11	12
P	Q	13	14	15

We see that in each row exactly **two** blocks (P and Q) are used to store parity meaning that in the whole RAID6 the capacity of **two** member disks is used to store parity. Therefore, in case of a 5-disk RAID6 you need to prepare $(5-2)*500 = 1.5 \text{ TB}$. Notice that you do not need to know what kind of RAID6 (synchronous or asynchronous, left or right, and so on) is used when estimating the capacity requirements.

- c) A client brought a RAID5 with delayed parity, consisting of three disks, 2 TB each. How much disk space you should have to copy the data off the RAID?

A RAID 5 of 3 disks with delayed parity looks like:

Disk 1	Disk 2	Disk 3
1	2	P
3	4	P
5	P	6
7	P	8
P	9	10
P	11	12

Here as well as in a regular RAID5, only one block in a row is used for parity which, on the scale of the whole array, gives the capacity of one member disk. Therefore, to copy data from this array you need $(3-1)*2\text{ TB} = 4\text{ TB}$.

Notice that you do not need to know what delay is used when estimating the capacity requirements for such an array. Like the previous example with 2-block delay, for the array with another delay, say 3-block as in the picture below, only the fact that exactly one block in a row is used to store parity is important.

Disk 1	Disk 2	Disk 3
1	2	P
3	4	P
5	6	P
7	P	8
9	P	10
11	P	12
P	13	14
P	15	16
P	17	18

Task 2 - Determining RAID type based on the Disk Management screenshot

Screenshot 1

On the Disk Management screenshot below we see that two disks are "Unallocated" while two other disks contain partitions. This means that Disk 4 and Disk 5 contain partition tables which Disk Management analyzes when you connect the disk set to a PC.

Disk 4 Basic 55.90 GB Read Only	(G:) 55.90 GB RAW Healthy (Primary Partition)
Disk 5 Basic 55.90 GB Read Only	(F:) 55.90 GB RAW Healthy (Primary Partition)
Disk 6 Unknown 55.90 GB Not Initialized	55.90 GB Unallocated
Disk 7 Unknown 55.90 GB Not Initialized	55.90 GB Unallocated

From possible array types - RAID0, RAID5, and RAID10 - we can exclude only RAID0, since only one disk in RAID0 can contain a partition table. So, we need to make a choice between RAID10 and RAID5. RAID10 is ideally suited for this picture since in a RAID10 the copy of each data block is stored resulting in two blocks with a partition table stored on two different disks. In Disk Management, as it should be, we see two disks with partitions.

Now let's consider a RAID5. For RAID5, it is typical that a partition table is written on the first disk at the beginning while all other disks contain zero blocks in the first row:

Disk 1	Disk 2	Disk 3	Disk 4
Partition table	0	0	P₁
0	0	P₂	0
0	P₃	0	0
P₄	0	0	0

Now let's see what data P₁ block stores. Due to the fact that $A \oplus 0 = A$ and the fact that Disk 2 and Disk 3 contain zero in the first row, P₁ will be a copy of block with a partition table. Therefore, in a RAID5, two disks store a partition table that should be reflected in Disk Management.

Conclusion: for this screenshot, RAID10 and RAID5 are possible array levels. To determine which one, you need to use some other methods.

Screenshot 2

On the screenshot below we see that Disk 4 differs from other disks. Disk 5, Disk 6, and Disk 7 are "Unallocated" while Disk 4 contains a partition.

Disk 4 Basic 55.90 GB Read Only	(G:) 55.90 GB RAW Healthy (Primary Partition)
Disk 5 Basic 55.90 GB Read Only	55.90 GB Unallocated
Disk 6 Basic 55.90 GB Read Only	55.90 GB Unallocated
Disk 7 Basic 55.90 GB Read Only	55.90 GB Unallocated

From possible array types - RAID0, RAID5, and RAID10 - we can exclude only RAID10, since for a RAID10 we should expect to see two disks having partitions.

Further the choice between RAID0 and RAID5. As in the previous case, in a RAID5 due to the specifics of parity calculation and the specifics of data typically stored at the beginning of the array, two disks from the disk set usually contain partitions. Therefore, RAID 5 is less likely.

Conclusion: only a RAID0 matches the picture. Additionally, you can make a conclusion that Disk 4 is the first disk in the array.

Screenshot 3

On the screenshot below, we see that two disks are "Unallocated" while two other disks have partitions.

Disk 4 Unknown 55.90 GB Not Initialized	55.90 GB Unallocated
Disk 5 Basic 111.67 GB Read Only	(F:) 111.67 GB RAW Healthy (Primary Partition)
Disk 6 Unknown 55.90 GB Not Initialized	55.90 GB Unallocated
Disk 7 Basic 111.67 GB Read Only	(H:) 111.67 GB RAW Healthy (Primary Partition)

Unlike the previous screenshots where all disks have the same sizes (55.9 GB), in this case the disks with partitions have different sizes - 111.67 GB. This is explained by the fact that in the previous examples MBR partitioning scheme is used while in this case we deal with GPT. MBR partitioning scheme doesn't show the full partition size if it is larger than the physical disk size; that's why we see that all disks and

partitions in the previous cases have the same sizes truncated to the physical disk size. When it comes to GPT, as in this case, Disk Management shows the actual partition size.

Now let's see what RAID type it can be. Similar to the previous cases, RAID0 is excluded immediately since we have two disks with partitions rather than one. We can choose between RAID5 and RAID10 based on GPT information about the partition size. Since the partition size is 111.67 GB, which approximately equals $55.9\text{GB} \times 2$, we can conclude that we deal with a RAID10 where exactly the half of disk space (two disks in this case) is used for redundancy. If it would be a RAID5, the partition size would be about 167 GB ($55.9\text{GB} \times 3$) since the capacity of one disk is used to provide redundancy.

Conclusion: only a RAID10 matches the screenshot. Further, RAID recovery boils down to RAID0 recovery of any of the two disk sets: {Disk5, Disk4} or {Disk 7, Disk6}. In each set, the first disk is the disk with the partition.