RAID Recovery course

Lesson 5 - RAID entropy analysis

Practice part

Task 1 – Determine configuration of a parity-based RAID

Frist we need to load disk image files. Notice that disk image files you are offered to work with in this task are in VHD format rather than sector-by-sector disk copies, that's why they cannot be loaded in ReclaiMe Pro directly. So before loading you need to mount them using Disk Management. Open Disk Management, click *Action-Attach VHD,* and specify the location of disk image files one by one. You should get the following picture:

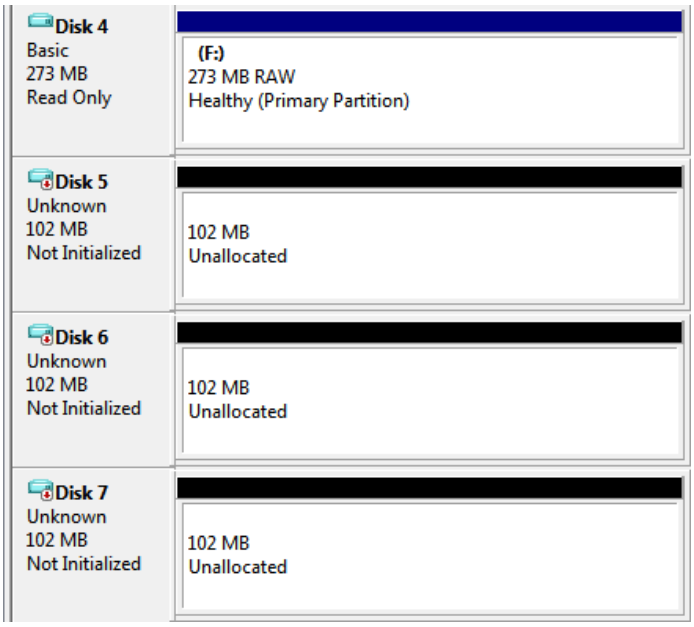

Let's analyze what we see in Disk Management. We see that one of the disks has a 273 MB partition that, given the size of each disk (102 MB), is very close to the **triple** capacity of array member disk: $102\ MB * 3 = 306\ MB$. The first conclusion is that it can be a 4-disk RAID5 in which we should expect a partition three times larger in size than a RAID member disk.

Then launch ReclaiMe Pro, select the just mounted disks for the content analysis, and get the following picture:

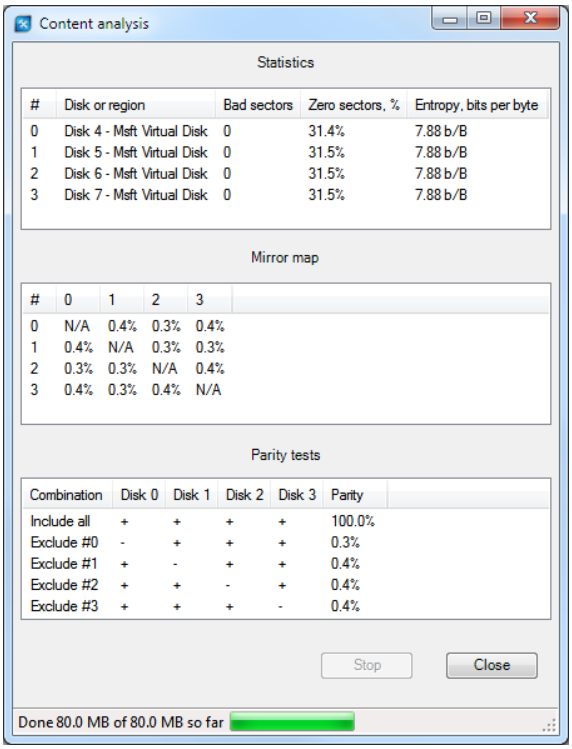

First, look closely at the statistics as to ratio of zeros and average entropy. We see that all disks have the same ratio of zeros (~31%) and average entropy (7.88 b/B) meaning that all disks belong to one array.

Based on Disk Management, we guess that we deal with a 4-disk RAID5 or RAID6 and mirror map does not contradict this layout.

Now it's time for parity analysis. Parity tests confirm our assumption about a 4-disk RAID5 or RAID6:

- There is 100% parity for full disk set meaning the correct disk number (4 disks) for the RAID5.
- There is no parity in all exclusion combinations.

Let's move on to the entropy analysis to get more information about RAID5 configuration. Close the window with content analysis and launch entropy analysis. With the default block size (1024 KB), ReclaiMe Pro gives you the following picture:

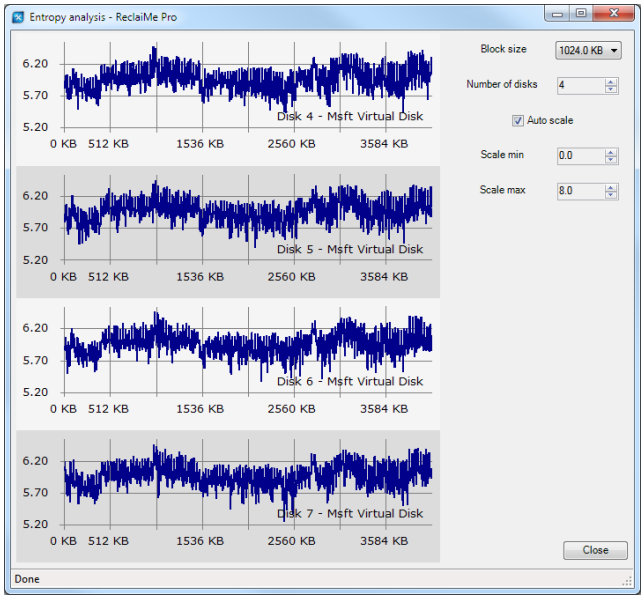

Too many peaks tell us that the block size is a way too large. We need to decrease the block size, say down to 128 KB and get the following:

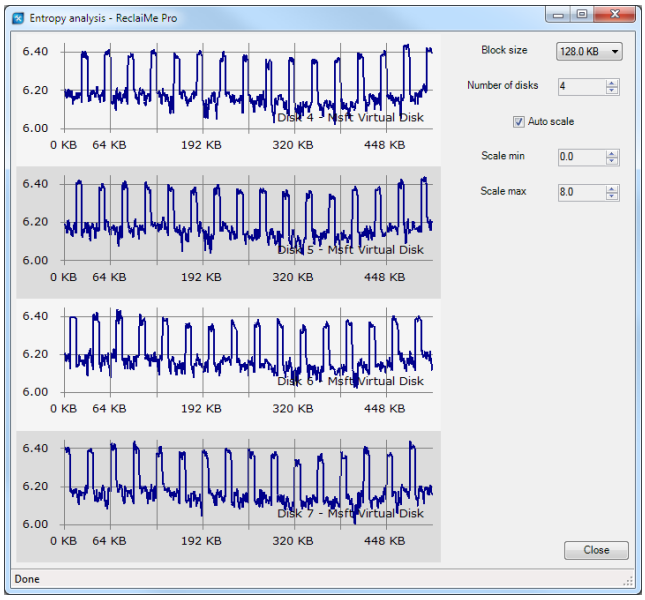

Looks better. Now let's calculate the number of peaks – 16, divide the current block size (128 KB) by 16 and get the true block size (8 KB). Set the block size to 8 KB and get the beautiful picture:

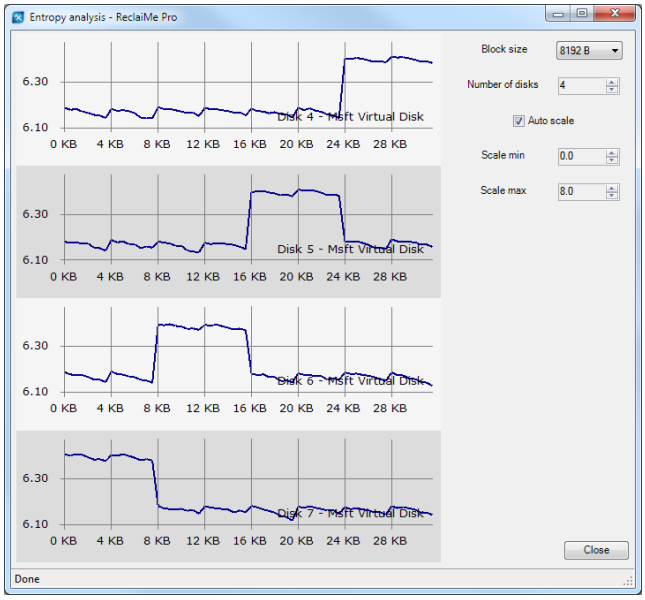

We see that there is only one peak per disk and peaks do not overlap; therefore, all the criteria of a typical picture are met. So we deal with 4-disk RAID 5 with 8 KB block size.

Additionally, we can get an idea about disk order by placing disks so that peaks follow each other (in our screenshot they are already in correct order). Note that you can not identify the first disk using this analysis; however, based on Disk Management we know that the first disk is Disk 1, because it stores the partition table. So the disk order is either Disks 4 - 5 - 6 - 7 or Disks 4 - 7 - 6 - 5. As for start offset, most likely it is zero since the peak on the Disk 7 starts right from the beginning.

Conclusion

The analyzed array is a full 4-disk RAID 5, block size is 8 KB, start offset is 0, the disk order is either Disks 4 - 5 - 6 - 7 or Disks 4 - 7 - 6 - 5. Then we need to recover RAID configuration in automatic mode specifying RAID5 level with all disks present and 8 KB block size.

Task 2 – Determine configuration of a parity-based RAID

First, you need to load disk images, similarly to Task 1. You should get the following picture:

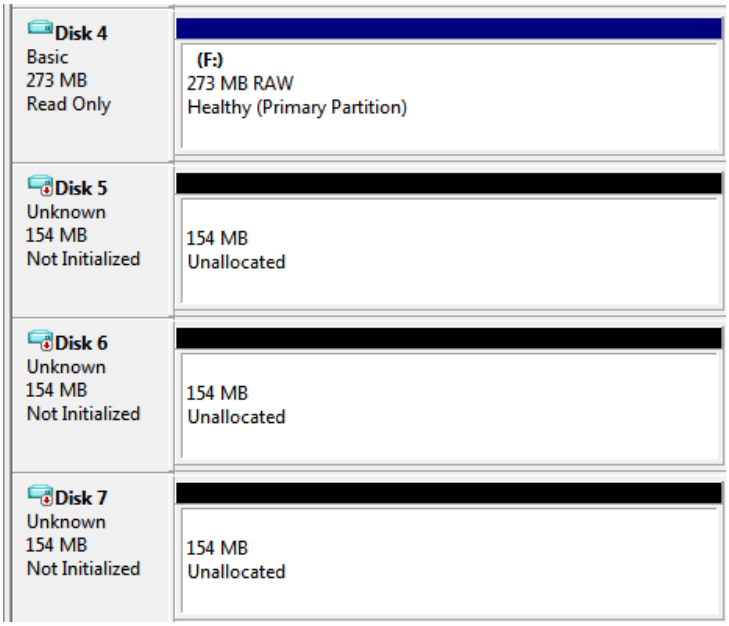

Let's analyze what we see in Disk Management. We see that one of the disks has a 273 MB partition that, given the size of each disk (154 MB), is very close to the **double** capacity of array member disk: $154\ MB * 2 = 308\ MB$. The first conclusion is that it can be either a 4-disk RAID10 or 4-disk RAID6. RAID10 is the less probable array type because in case of RAID 10 we should see two disks containing partitions in Disk Management.

Then launch ReclaiMe Pro, select the just mounted disks for the content analysis, and get the following picture:

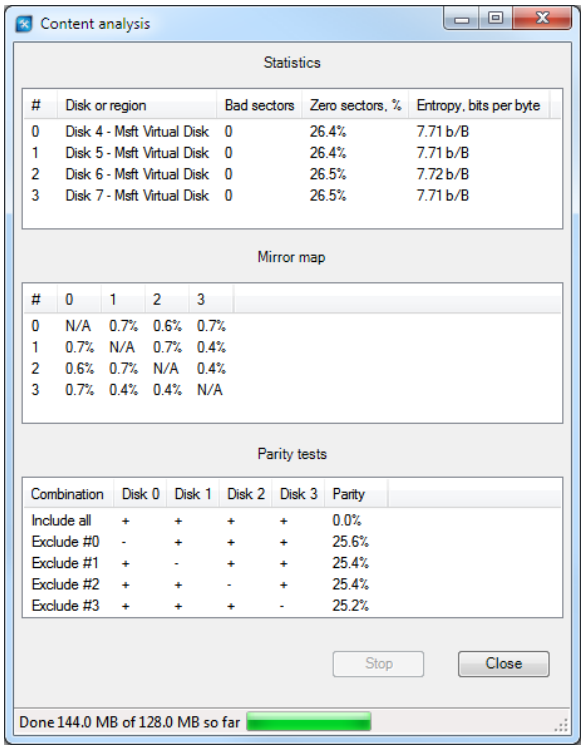

First, look at the statistics. We see that all disks have the same ratio of zeros (~26%) and average entropy (7.71 b/B) meaning that all disks belong to one array.

Based on Disk Management, we guess that we deal with either a 4-disk RAID10 or 4-disk RAID6. However, mirror map does not confirm RAID 10 layout, because there are no mirror pairs.

As for parity analysis, parity tests confirm our assumption about a 4-disk RAID6, because:

- There is 0% parity for full disk set.
- There is 1/N parity in all exclusion combinations where N – is the number of disks.

As we discussed in our lesson, such values in parity tests point to a certain RAID 6 variation.

Let's move on to the entropy analysis to get more information about RAID6 configuration. Close the window with content analysis and launch entropy analysis. With the default block size (1024 KB) ReclaiMe Pro gives you the following picture:

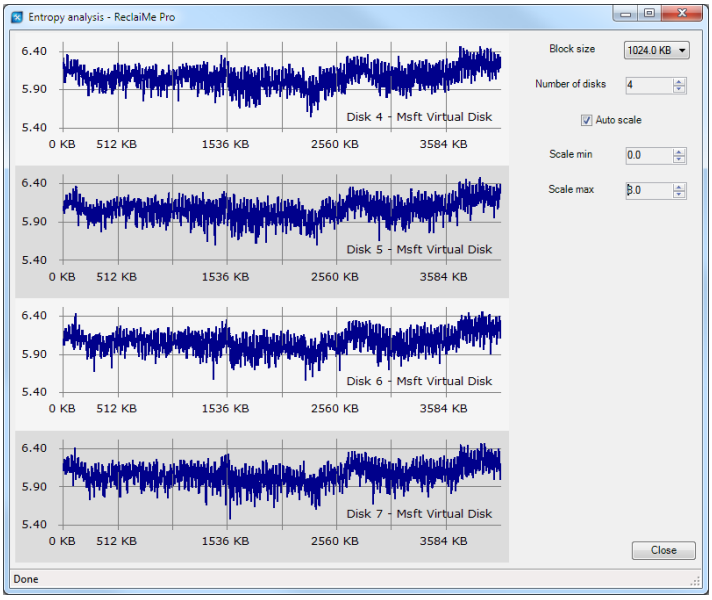

Once again, we see that the default block size gives too many peaks and so we need to decrease it. Try, say, 8 KB block size and get the following picture:

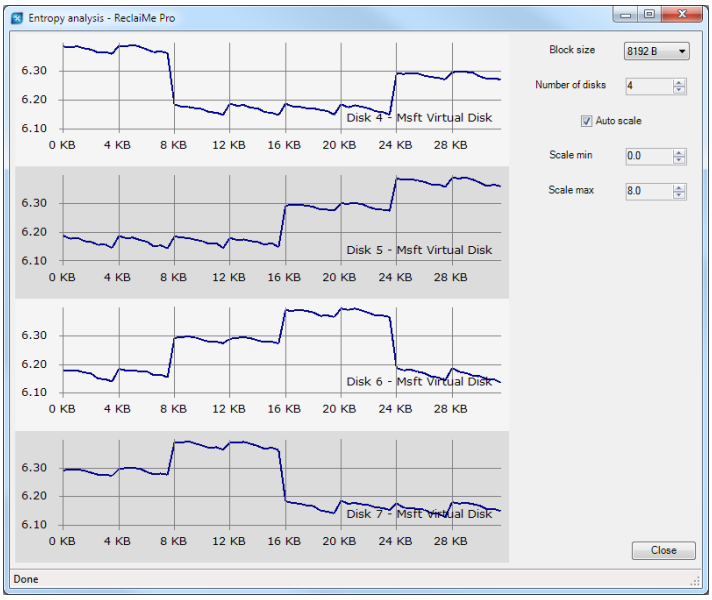

We have guessed right because we see typical 2-plateau RAID6 peaks overlapped by the half peak width. Additionally, we can get an idea about disk order – either Disks 4 - 5 - 6 - 7 or Disks 4 - 7 - 6 - 5; the first disk can be determined by looking at Disk Management.

Conclusion

The analyzed disk set is a full 4-disk RAID6, 8 KB block size, and disk order either Disks 4 - 5 - 6 - 7 or Disks 4 - 7 - 6 - 5. Next you need to recover RAID configuration automatically specifying RAID6 level and 8 KB block size.