Partition Alignment Dramatically Increases System Performance

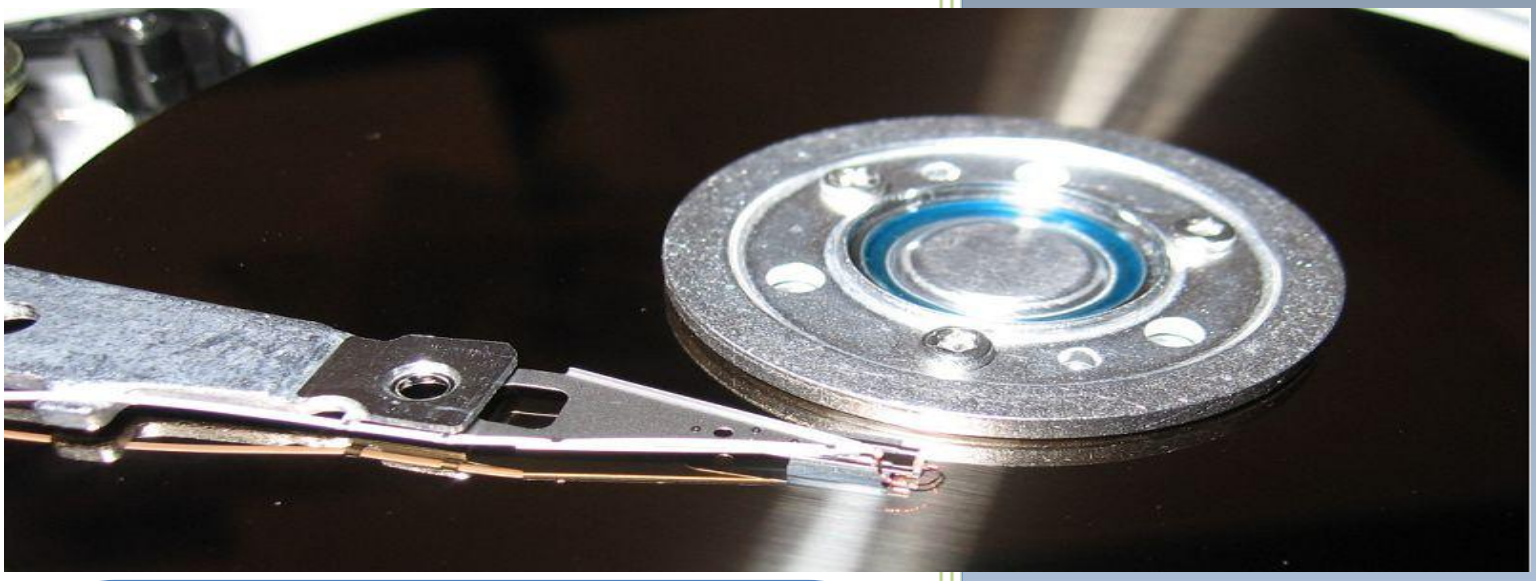

Information for anyone in IT that manages large storage environments, data centers or virtual servers.

Paragon Software Group

Paragon Alignment Tool - How to easily increase your system performance

Contents

Introduction

Storage subsystem performance depends on many factors. One of them is a properly configured partitioning scheme. How it can treat file operations performance is described in this article.

The introduction of 4K sectors in several modern disk drives makes the partitions alignment (PA) problem important for a mass segment of the IT world and for home users. Implementation of enterprise virtual environments and the use of different RAID technologies are yielding more applications for PA. In all these cases a mismatch between actual physical and logical representations of data may cause a significant decay in overall system performance and hardware longevity.

What is Partition Alignment?

To understand what partition alignment is we should first see how data is stored on the drive.

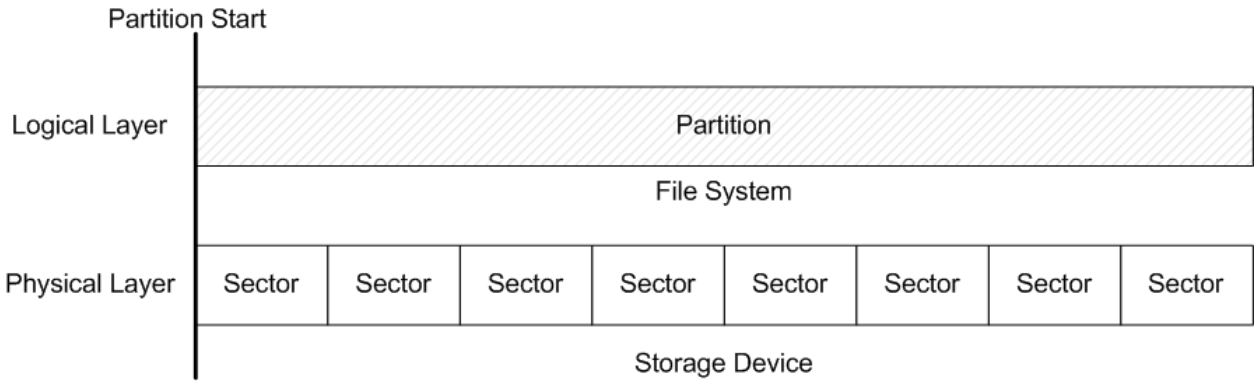

Here you can see a simplified scheme of partitioning, where one partition consumes all drive space and the sectors. It is also aligned with the first sector that is the partition start corresponds with the first sector start. This scheme is oversimplifiedThe partition consists of some logical units, the clusters. See how they connected with sectors.

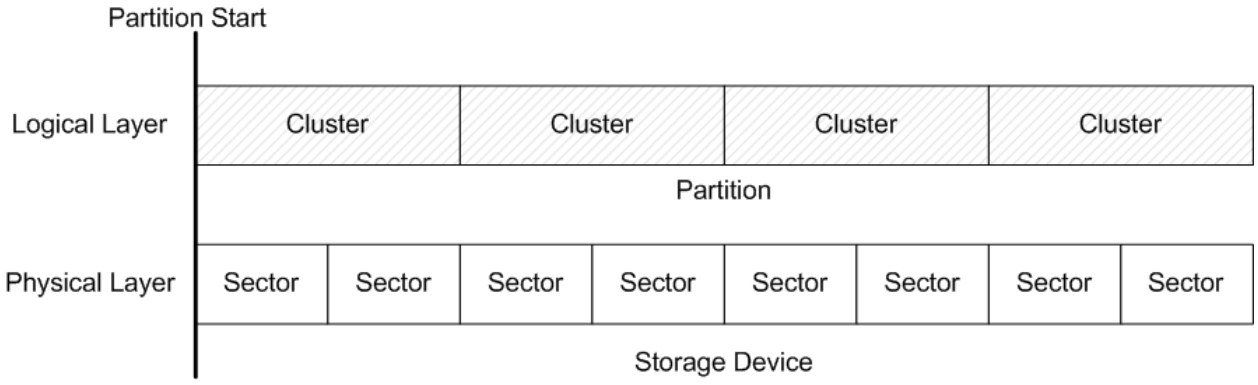

Now you can see that on the scheme one cluster corresponds to two sectors. It means that when data is being read from one cluster, for example a small text file, the storage device reads data from two

sectors. Despite that on the scheme one cluster corresponds to two sectors usually a cluster is connected with much more sectors. Also notice that because of the partition start position and the first sector start corresponds to each other – all clusters are aligned with sectors, the partition is aligned. Data operations are as fast as possible.

Let's discuss why partitions can be misaligned and what this can cause.

Why misaligned partitions are the problem for hard disk drives?

Partitions can be misaligned because the physical sector size is not 512 bytes and software does not know about it. Modern hard disk drives, for example Western Digital, Seagate, etc., have an internal 4096 bytes physical sector size, their logic operates 4K chunks of data, but for outside hardware and software they appear as "traditional" drives with 512b sectors. This emulation is needed for old-software compatibility. Thus another level of abstraction is being added. See the scheme.

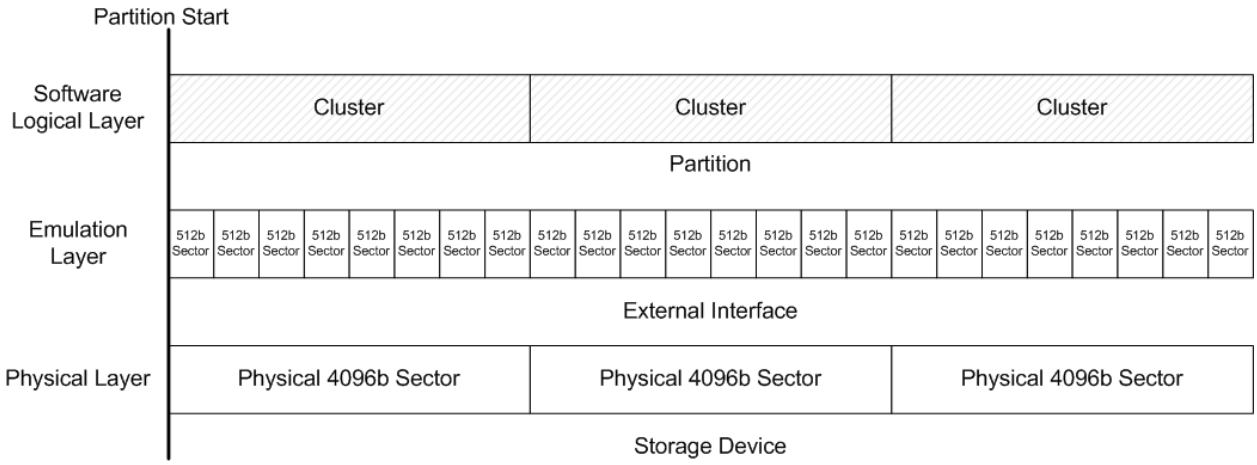

Here you can see three layers of data abstraction. The ground level is the hard disk drive itself and 4K physical sectors. Second level is external IO abstract sectors with 512 bytes size. And above all, the actual file system with clusters, which the size is also 4K that equals to eight abstract sectors and one actual physical sector. So 1 cluster = 8 abstract sectors = 1 actual physical sector.

Also notice that all tree levels are aligned in relation to each other and disk start. So when you read or write data from one cluster, it will be actualized in four 512b sectors and one 4K sector. The amount of read-write operations is minimal; all disk operations are at maximum performance.

But things get worse if logical clusters will be shifted relative to all underlying levels. Look at the picture.

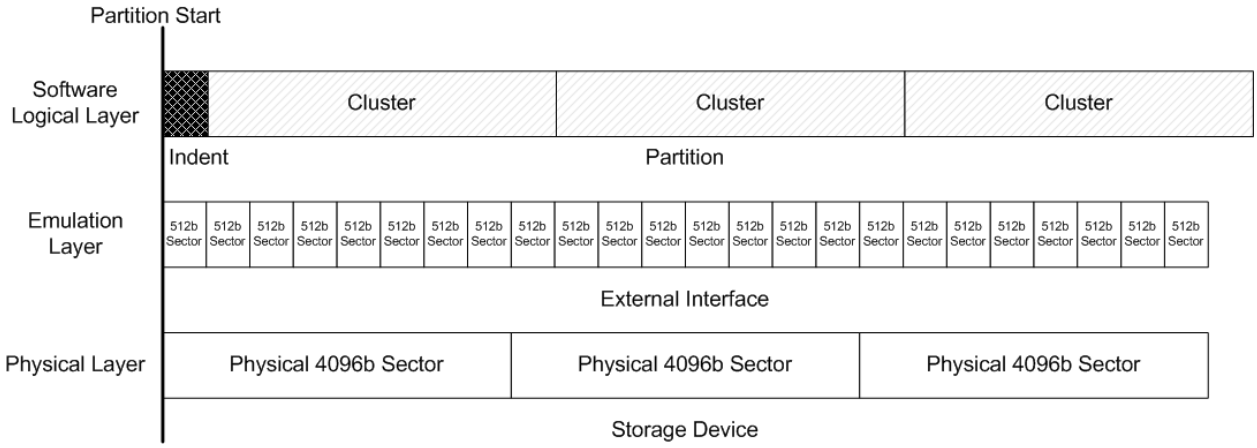

Here you can see that the partition is shifted from the disk start on one 512b sector. As the result all logical clusters now linked with two actual physical 4K sectors, and all read-write operations will be multiplied twice. Whole system performance now *becomes degraded as a hard disk has to make two operations with two sectors to manage your data* instead of one as if the partition was properly aligned.

What causes this shift? All Windows operating systems before Vista use a factor of 512 bytes to create volume clusters. Thus they place a partition start aligned to 512b sectors and not to 4K sectors, as it is shown on the previous picture, where the partition start is indented on one 512b sector.

Usually the partition start is indented on 63 sectors, because it is an old measure of a disk “cylinder” and some old versions of DOS or Windows demand that the partition has to be aligned to the “cylinder” for correctly sectors addressing and accessing. It is an old compatibility issue and all modern operating systems do not use this archaic CHS (cylinder/head/sector) addressing scheme. Instead, the Logical block addressing (LBA) scheme is used, where there is no any “cylinders” or “heads”, sectors are addressed continuously over a whole disk drive. But by legacy reasons all versions of Windows before Vista creates partitions according to this “cylinder alignment” rule.

There was no problem with this rule and partitions alignment in the home users segment before the appearance of mass 4K hard disk drives. Partitions aligned accordingly to 63 sectors start are not aligned with 4K sectors by default. This is a matter of simple math. Here is the equation:

$$\frac{\mathbf{63\ sectors}\ \times 512\ bytes}{1\ sector\ \times 4096\ bytes} = 7{,}875$$

As you can see 63 sectors by 512 bytes do not match to integer number of 4K sectors, thus this partition and all others following on the disk are misaligned.

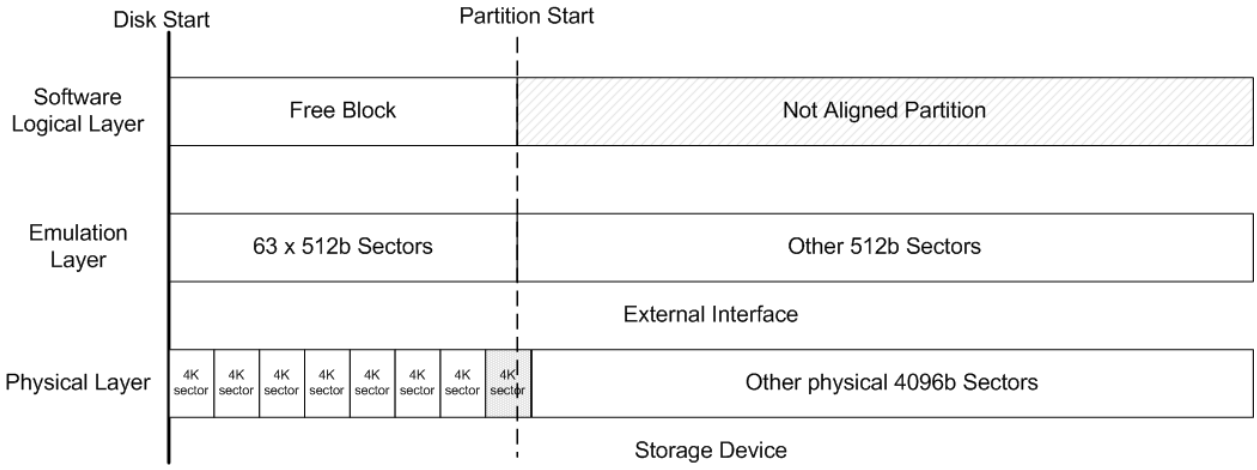

How Paragon Alignment Tool can help?

As you can probably suspect PAT does nothing more than move partitions on some amount of 512b sectors so that all volumes become aligned. For example we need to move a 63 sectors partition forward only to one 512b sector and it will become aligned.

$$\frac{\mathbf{64\ sectors}\ \times 512\ bytes}{1\ sector\ \times 4096\ bytes} = 8$$

Now the partition start corresponds to the 4K sector start. This partition and all others following on the disk are properly aligned.

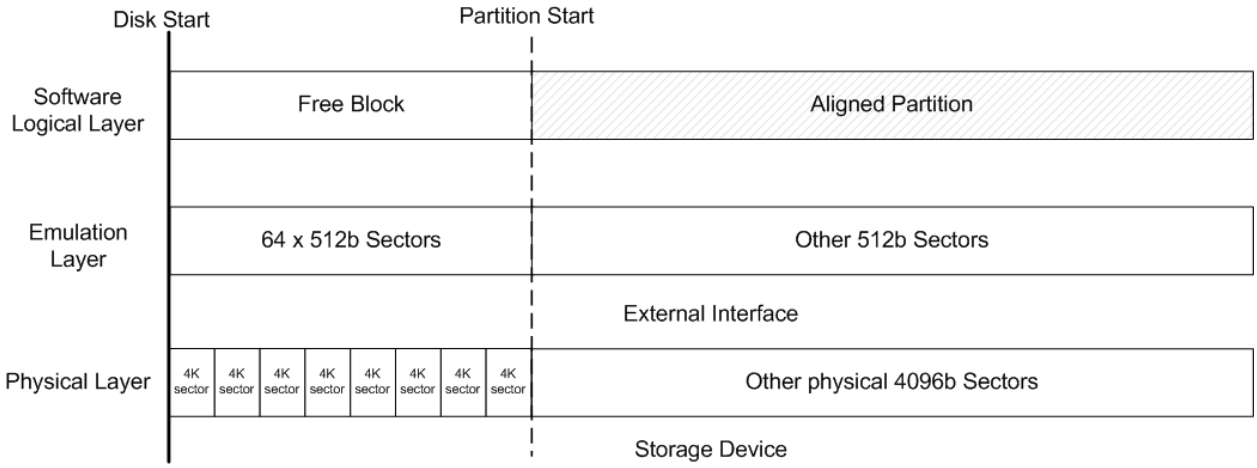

PAT moves the partition start on 2048 512b sectors from the disk start; it equals to 1MB and also perfectly fit to 4K sectors boundaries rule.

$$\frac{\mathbf{2048\ sectors}\ \times 512\ bytes}{1\ sector\ \times 4096\ bytes} = 256$$

There is another one important issue. Some hard disk drives with 4K sectors have a trigger to match PA problems. Their internal controllers can be adjusted to shift emulating level addressing to an offset of

one sector. Thus #63 sector becomes #64 and all partitions above become aligned. Please look at the scheme.

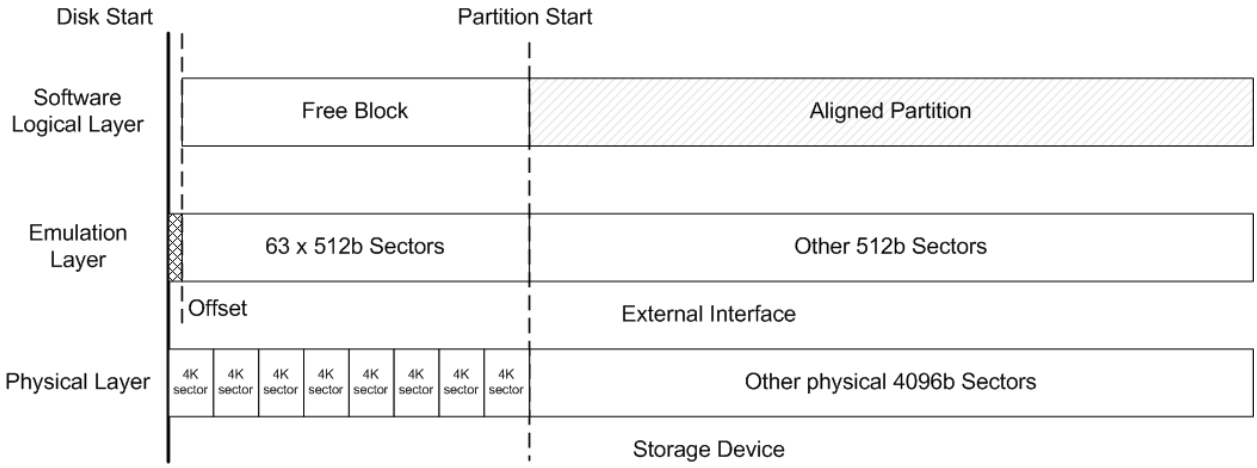

As you can see software do not even know that it works on the shifted volume, but internally all partitions are aligned. However PAT is an offset-aware tool. It is able to recognize the described situation and skip alignment of partitions on such a drive. If it was not able to notice the logical offset, the further alignment operation would cause appearance of misaligned partitions and the effect of using of the tool would be opposite negative.

Our performance tests show that partition alignment boosts file system operations with the 4K hard disk drive by 2X.

Why misaligned partitions are the problem for SSD?

Misaligned partitions problem is even more important for SSD drives than for traditional hard disk drives. Many modern SSD drives have an internal memory page size 4096 bytes or larger accordingly to 4K size, which are some analogue for 4K sectors. Thus all previously mentioned problems are the same for SSD partitions alignment.

There is one crucial SSD issue besides file system speed decline (which is not so noticeable in comparison to traditional HDD). It is the SSD memory cells degradation after some amount of write operations. So if partitions on SSD are misaligned beside downgraded system speed you put your solid state drive in danger. After partitions alignment PAT eliminates all redundant read/write operations and thus provides speed boost and grants *SSD a longer lifetime*.

Why misaligned partitions are the problem for SAN and RAID?

RAID is used to compose many hard disk drives or other storage devices into one large array of data. This array is seen as one large storage device in the system and data is striped across it. The granularity at which data is stored on one drive of the array before subsequent data is stored on the next drive of

the array is called the stripe-unit size. Stripe-unit size may be different; you can set the stripe-unit size for example to 8 KB, 16 KB, 32 KB, or 64 KB.

System performance may slow when you use a hardware-based redundant array of independent disks (RAID) or a software-based RAID and if the starting location of the partition is not aligned with a stripe unit boundary in the disk partition that is created on the RAID. In this case one data operation will be multiplied over several RAID disks.

To resolve this issue PAT aligns partitions to 2048 sectors. A starting offset of 2,048 sectors covers most stripe unit size scenarios. All data operations become faster as there are no redundant disk operations anymore. Please look at the scheme.

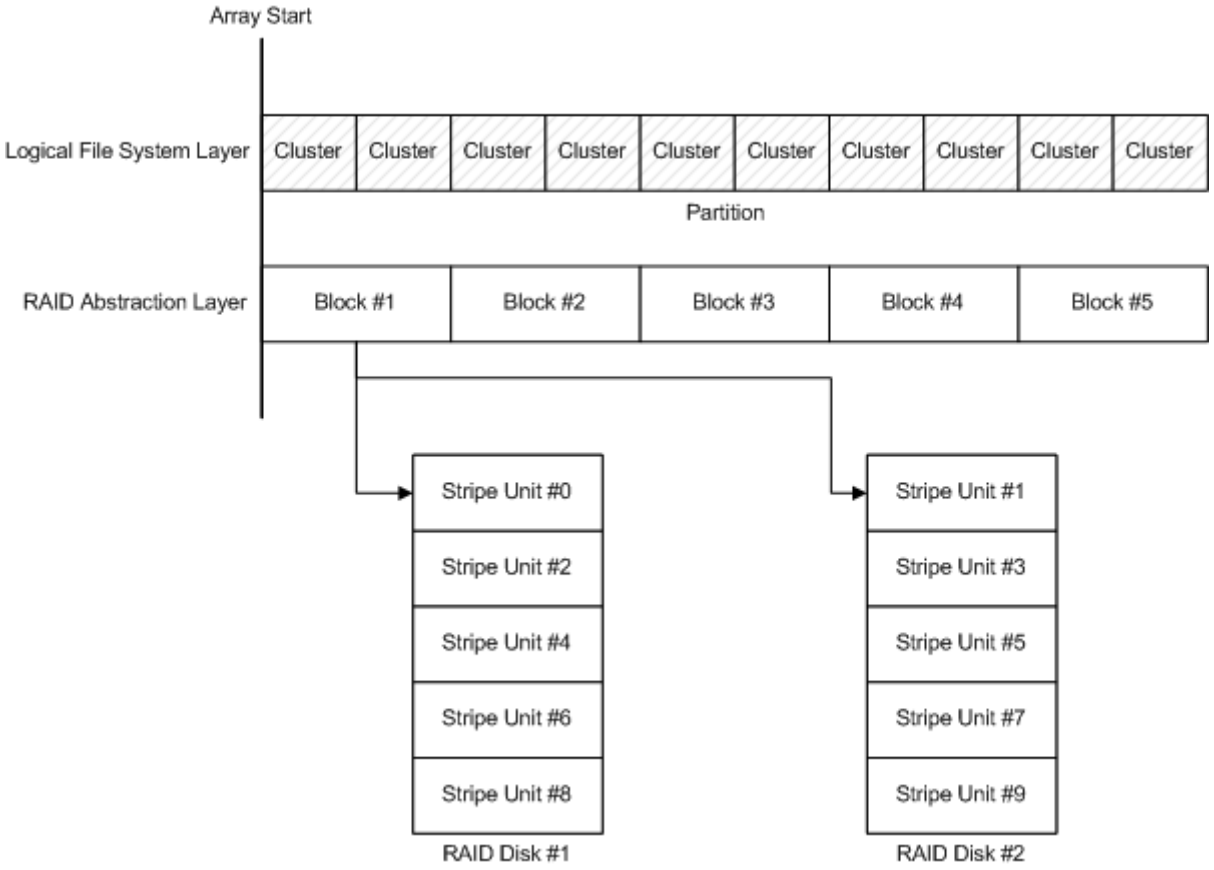

Here you can see a properly aligned partition and RAID-0 array where each block is connected with two stripe-units (it is shown for the first block on the scheme). In this situation access to one logical cluster will be followed by operation with one block and two stripe-units.

Naturally, SAN (Storage Area Network) is just a large RAID massive distributed over a local network or by Fibre Channel. Thus all issues of RAID partition alignment are the same for SAN.

Why misaligned partitions are the problem for virtual environments?

Alignment in a virtual infrastructure is critical to performance, hardware life cycle, and storage efficiencies. Misalignment results in retrieving more data from underlying array than what the virtual machine is requesting. This results in inefficiencies leads to requiring more storage hardware resources to serve a workload.

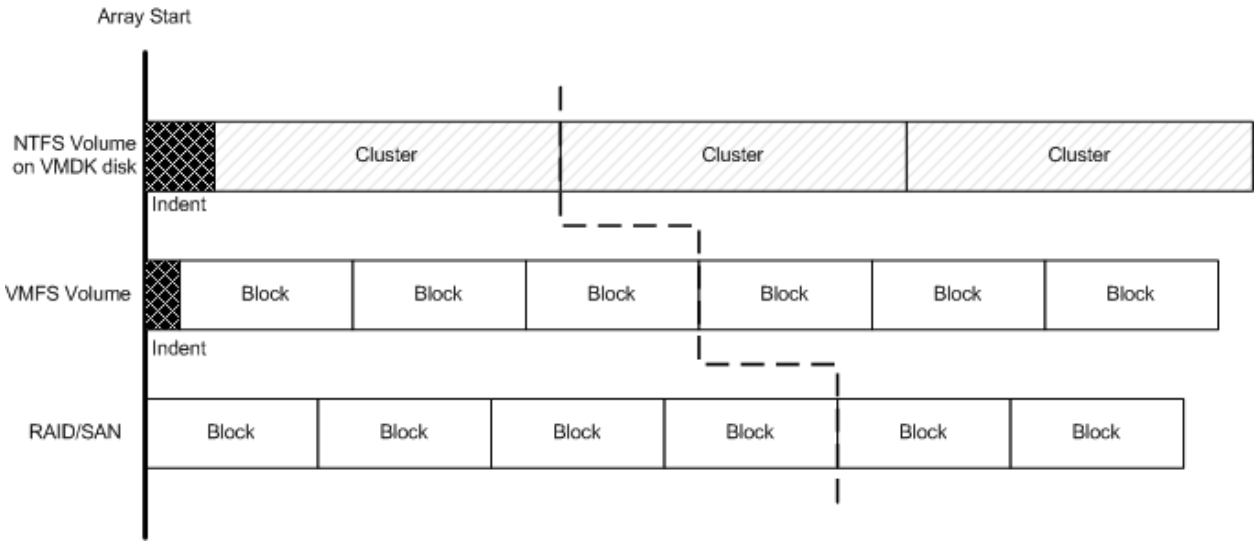

On the previous picture you can see a VMware Server configuration, which uses RAID/SAN storage. Here are even two levels of misalignment that cause a severe performance decline. For example, to read data from the first cluster the system has to read three VMFS blocks and four RAID blocks (and consequently more stripe-units).

PAT cannot move VMFS volumes itself at the present moment, but can align partitions on virtual drives.

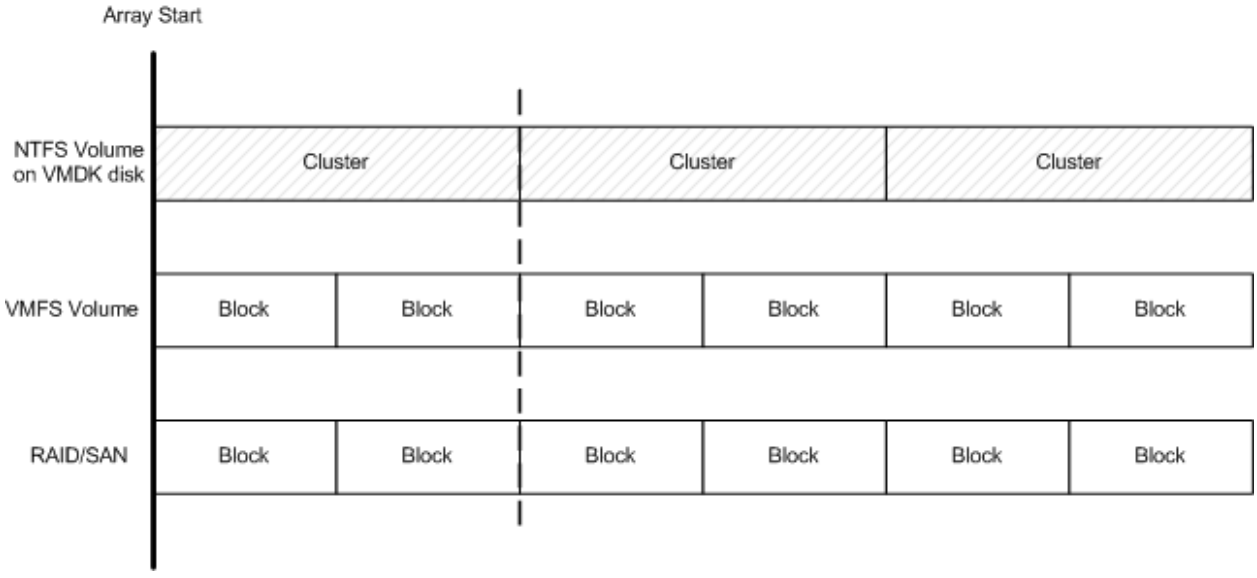

In this configuration all volumes and partitions are properly aligned and overall data operations are two times faster. Please notice again now to get data for one cluster the system just need to read two VMFS blocks and two RAID blocks.

Conclusion

Using Paragon's Partition Alignment Tool is very simple. You just run the program and it will automatically detect all misaligned partitions. You can accept these settings and run the tool. After running PAT all partitions become aligned.

In this document all major issues regarding partitions misalignment on 4K hard disk drives, solid state drives, RAID, SAN and in virtual environments were described. The reasons why partitions become misaligned and the consequential disadvantages were also discussed.

Despite that in the nearest future all AP issues for 4K HDD and SSD go away as 4K physical sectors and memory pages become visible and accessible on the operating system level and the need in emulating become unnecessary, misalignment will persist the problem for RAID/SAN and virtual environments. PAT is the best tool to beat all these problems now and in the future.

About Paragon Software Group

Paragon Software Group is an innovative software developer focused on two dynamic growth markets. The company's comprehensive product line for the data storage market addresses the needs of data security, virtualization solutions, storage and management for PCs, servers and networks. Founded in 1994, Paragon Software has offices worldwide, delivering its solutions to consumers, small business and enterprise clients through a network of Value Added Resellers, distributors and OEMs as well as online through the company website. Paragon Software provides technology to a host of world class companies and partners including Cisco, Dell, Toshiba, NEC, Siemens, Microsoft, Motorola, Nokia, and more. For more information please visit the company website at www.paragon-software.com.