Technical Advisory Letter

Running Diskeeper and V-locity on SAN Devices

By Gary Quan, SVP Product Strategy

This technical advisory letter addresses the concern of running the Diskeeper® or V-locity® products on SAN devices. This stems from some who have recommended not running a disk defragmenter on SAN devices. I actually agree with that in most cases, a "standard disk defragmenter" should not be run on a SAN device, but Diskeeper and V-locity are not "standard disk defragmenters." They are storage optimizers that contain technology and innovations that are designed for these specific types of environments.

I have worked with many customers that have this same concern, including representatives from VMware and SAN vendors, but after explaining this technology and how the products work, they all saw the benefits. This general concern broke down to these three specific concerns:

1. Defragmentation will not provide performance gains because it cannot control the exact placement of data at the SAN level.
2. Defragmentation will cause extra data movement overhead at the SAN level.
3. Defragmentation will cause Thin Provisioned disks to grow excessively.

After understanding how the technology in Diskeeper and V-locity work, it can be seen how these beliefs are not true with these products and how the products actually benefit this type of environment. There is actually a fourth concern which most users did not think about, but is one the two products addressed.

4. Defragmenters cause I/O resource conflicts on SAN Storage.

The rest of the letter will go into more details on these specific concerns and how the Diskeeper products handle them and how the systems and users ultimately benefit from it.

1. Defragmentation will not provide performance gains because it cannot control the exact placement of data at the SAN level.

The first part of this statement is false. In-house testing and use cases from customers have empirically shown significant I/O performance gains on the systems that are using the SANs. Here are some of the benchmark results.

figure a. *From in-house testing on a Windows Server 2008 R2 system with the test volume on a SAN.*

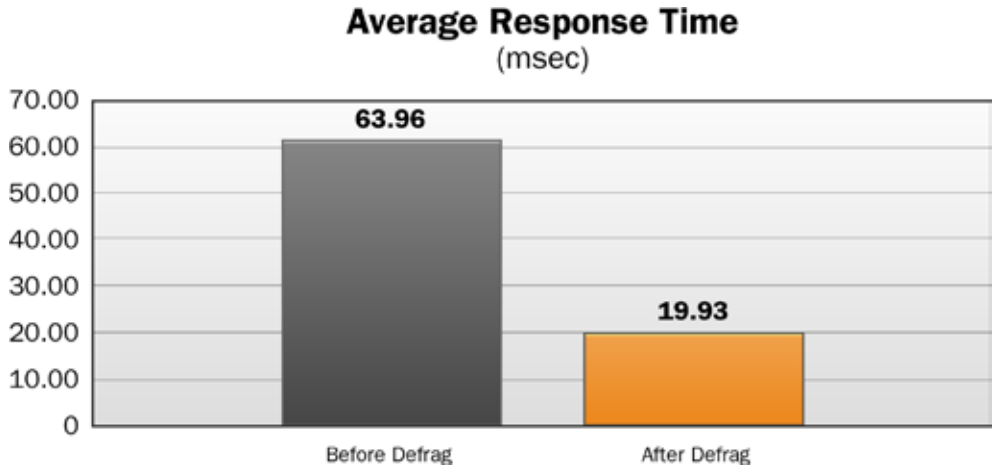

220% Improvement in Average Response Time

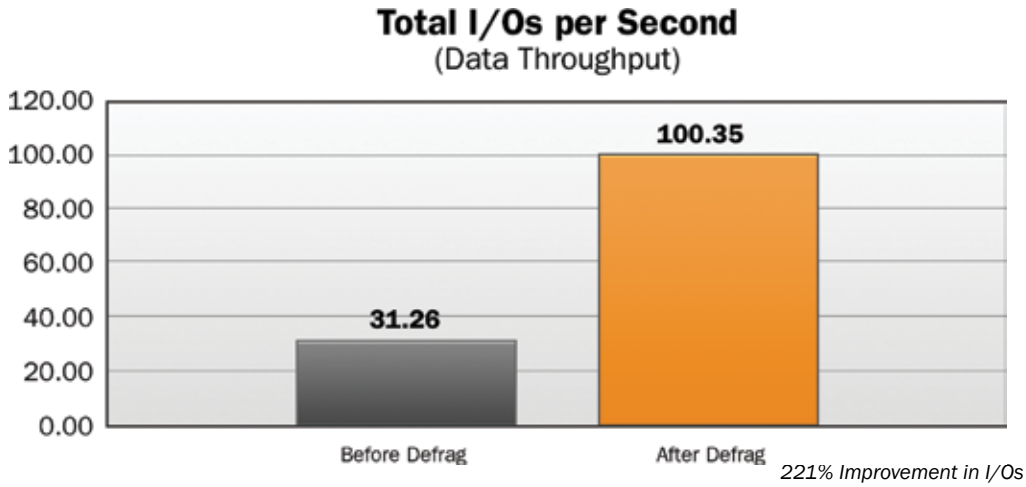

221% Improvement in I/Os per Second

This shows that fragmentation does degrade performance and it needs to be handled to get that performance back. There are more test results and testing details available in the Diskeeper white paper on Improving the Performance of SAN Storage.

figure b. *Customer Use Case on a SAN attached system, using Microsoft Performance Monitor showing I/O average performance with and without V-locity on the system:*

Performance Monitor Counter	Before V-locity Installed	After V-locity Installed	% Improvement
Split IO/Sec*	2.858	0.251	91.22
Avg. Disk sec/Transfer	0.059	0.027	54.24
Avg. Disk sec/Write	0.02	0.01	50.00
Avg. Disk sec/Read	0.118	0.097	17.80
Avg. Disk Queue Length**	6.037	1.942	67.83

In the latter part of the statement where it states that it is not able to control the exact placement of data on the SAN device, that is correct. From the O/S level, exact placement of data on the SAN device cannot be done at this time as the SAN controller will determine where to place the data. What can be controlled is the number of I/O requests from the O/S system level and this is what Diskeeper & V-locity optimize. As shown in figure c, by decreasing the amount of I/O requests at the O/S system level, it will decrease the amount of I/O requests being performed at the SAN level that result in performance gains.

* Split IO/Sec – Rate at which I/O requests are divided into multiple requests. This can occur from fragmentation since multiple requests must now be performed to satisfy the original request. A lower number is better.

** Avg. Disk Queue Length – Average number of disk requests being and waiting to be serviced. A lower number indicates better performance as there are fewer requests having to wait to be serviced.

figure c.

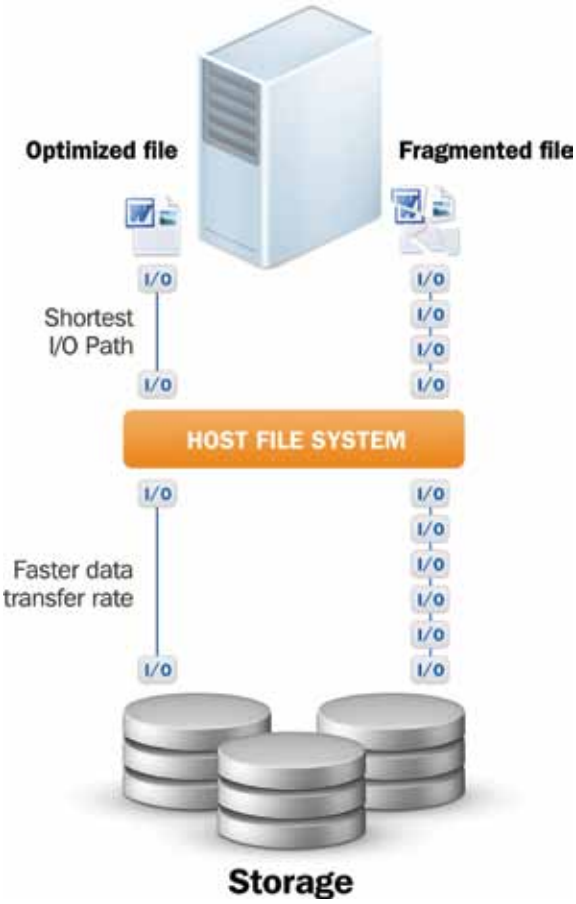

2. Defragmentation will cause extra data movement overhead at the SAN level.

This is actually true in many cases using “standard defragmentation” technology, especially in SAN storage environments using SAN-specific technologies such as Snapshots, CDP (Continuous Data Protection), or Deduplication. Whenever the system detects changes in the data, these SAN-specific technologies will need to log or copy the new changed data, which is extra data movement.

Standard defragmentation methods normally defragment the files after they are fragmented and many will try to defragment fragments that do not provide significant performance gains, such as a large file (i.e., 500mb) that is in two large pieces. The problem with moving the data to defragment, is that these SAN specific technologies believe that the data has been changed (when it has not just moved to a new location) and it will then do the logging or copy process.

This is where the IntelliWrite® technology found only in Diskeeper and V-locity really shines. This technology prevents the majority of the fragmentation from occurring in the first place which prevents the data having to be moved later. The user gets the benefits of an optimized volume without this extra data movement. Also, on the data that does need to be moved, it will only move the data that is causing significant performance degradation.

3. Defragmentation will cause Thin Provision disks to grow excessively.

Again, this can be true using 'standard defragmentation' technology which does not take into consideration what free space to use to defragment files. Along with IntelliWrite that prevents most of the fragmentation from occurring in the first place, Diskeeper and V-locity contain over 20 specialized engines, including those that when file data movement is performed, it attempts to use free space that has already been allocated which minimizes thin provisioned disk growth.

Besides Diskeeper and V-locity minimizing any thin provision growth itself, V-locity 3.0 also contains a feature to help reclaim unused data blocks. It does this by automatically zeroing out the free space which allows the compaction utility to be more efficient. The below graph in figure d shows the effectiveness of it.

Figure d.

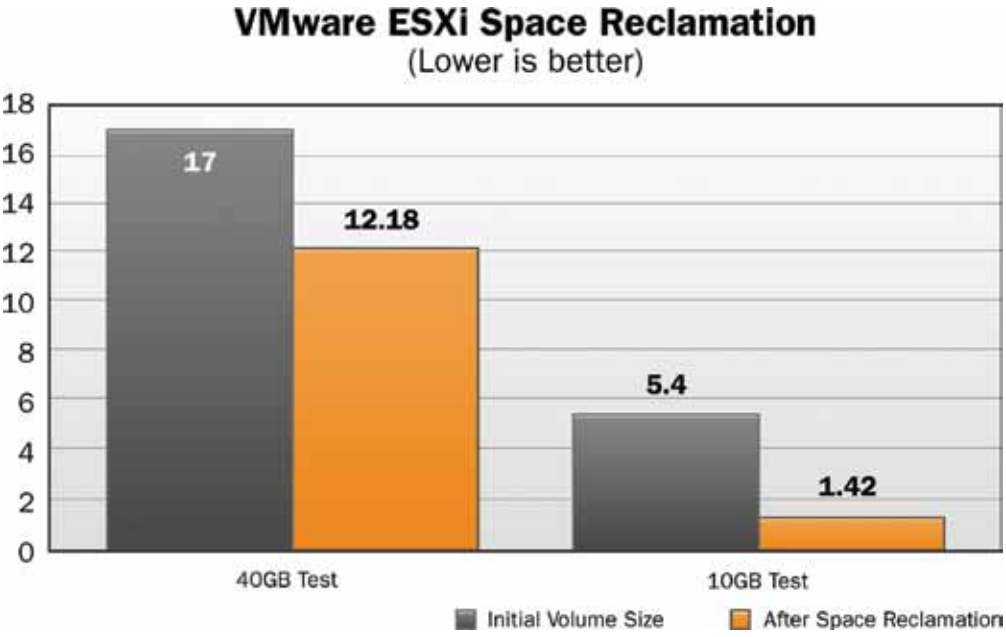

4. Defragmenters cause I/O resource conflicts on SAN Storage

Take the case where there are 10 VMs (Virtual Machines), all having LUNs (Logical Unit Numbers), all mapped to the same SAN. Well if all of these VMs are optimizing the volume without knowing what resources the other VMs, including system and user applications, are using, it can cause a conflict for resource usage. V-locity contains a new technology called CogniSAN that detects external resource usage within a shared storage system, such as a SAN, and allows for transparent optimization by never competing for resources utilized by other VMs over the same storage infrastructure.