

# Compellent Storage Center

## Understanding NTFS Fragmentation with Virtualized Storage

# Table of Contents

# Overview

Compellent is a leading provider of enterprise-class network storage solutions that are highly scalable, feature rich and designed to be easy to use and cost effective. Compellent's Storage Center is a Storage Area Network (SAN) that is designed to significantly lower storage and infrastructure capital expenditures, reduce the skill level and number of personnel required to manage information and enable continuous data availability and storage virtualization.

Utilizing Compellent's Dynamic Block Architecture, Storage Center intelligently optimizes data movement and access at the block level. Dynamic Block Architecture records and tracks specific information about every block of data that provides the system intelligence on how that block is being used. All of this metadata or "data about the data" enables Storage Center to take a more sophisticated and intelligent approach to storing, recovering and managing your data

# Fragmentation

The New Technology File System (NTFS) was developed and released with Windows NT 3.1 in 1993. The NTFS file system uses "clusters" as its means of organization of data written to disk. Cluster size can be modified by the administrator when a new partition is created, but by default the cluster size used is 4K. The cluster size also determines the smallest size of data that the operating system can write to that specific partition.

Fragmentation occurs when a file is written to disk and the data is stored in clusters that are not located next to each other on the disk. The most common reason for fragmentation is when files are deleted; however, modification of files can have the same effect.

Joe Kinsella from Diskeeper provides the following example of fragmentation:

"If you deleted a non-fragmented 40K file that occupied 10 contiguous clusters on an area of the disk surrounded by other used clusters, the disk will now have 10 free clusters available for use. If you then saved an 80K file, which requires 20 clusters, the operating system may choose to use the 10 recently free clusters and then find an additional 10 clusters from somewhere else on the disk. This means our 80K file is now fragmented, residing in two different locations on the disk."

# Defragmentation

Defragmentation is the process of analyzing folders and files across a partition and determining ideal placement on the disk. Particularly, the defragmentation process looks for files that have been broken up, or haven't been written to contiguous space. Defragmenting a drive will attempt to organize the data clusters that belong to a single file as close to each other as possible and contiguously.

# Fragmentation on a Compellent Storage Center

Storage Center stores data in 512k containers called blocks. The Storage Center is not aware of a file system or what types of files are in those blocks. Because of this, the Storage Center relies on the server to ensure that data is not fragmented to a point that results in performance degradation. It is possible with a heavily fragmented server

volume that a file might reside among multiple SAN blocks and therefore require more I/O to retrieve that file than necessary.

When deciding if you should defrag a volume there are a couple of factors that should be considered from the Storage Center point of view. The Storage Centers Dynamic Block Architecture tracks all changes to all blocks on the SAN. When defrag runs it first copies the data to a new portion of the disk and then copies it back. From the SAN this looks like new writes and will consume additional space on the SAN. This can be remedied if you use Windows Free Space Recovery and have a fairly short replay retention time.

The second thing to consider is how a defrag might affect tiering on your system. For example: You have a volume that is set to use RAID10 for writes and RAID5-5 for replays. The majority of this volume sits on 5-5. When you run a defrag all of the modified or newly created blocks will reside in RAID10; that is until the next replay is taken. This is the same for if you had a volume that used Tier 1 RAID10 for writes and Tier 3 RAID5-9 for replays. All modified or new data would be written to Tier 1.

The third concept to understand is how defragmenting effects replication. SAN-based replication is performed at the block level. Since defragmenting a drive creates a large amount of block level changes, this can greatly affect the size of data that needs to be replicated.

# Maintenance

Defragmenting a volume will most likely not increase performance of a volume from Storage Center's point of view. Defragmenting a volume may however greatly decrease the amount of IO that a server sends to Storage Center for a given operation. In this sense, a volumes performance can increase. Imagine a 20KB file that is fragmented in to five non-contiguous 4KB clusters. It will take five separate read IOs to retrieve that fragmented file. After defragmenting that file, only one read IO would be needed.

If you do want to defrag it is important to understand what that will mean on the Storage Center side of the equation. There is no need to choose the 'optimize' option when defragmenting a Storage Center volume as the location of the defragged blocks on the disk doesn't matter. Optimize will not only put the pieces of a file back together it will also move all files to the beginning of the disk and next to each other. On the Storage center this only creates more I/O on the volume but does not move any of the files together, therefore it is not recommended. If you are going to defrag then you should also run Windows Free Space Recovery to help reduce storage consumption overhead as a result of defragging a volume. If you are going to defrag, it makes sense to do it on a regular basis like once a month or every six months.

# Revision History

1.0 – Initial Release (10/28/2009) by Copilot Services & Technical Solutions.