

Why SAN Defragmentation is Needed Today

By Salvatore Salamone, Executive Editor, Ziff Davis Enterprise

As data volumes continue to grow and more companies consolidate and virtualize their IT infrastructures, new demands are being placed on storage performance.

Specifically, growth in file sizes, data retention laws, and increased use of media-rich content (business video storage is growing by 50 to 60 percent per year) are straining file data storage systems. Additionally, IT equipment consolidation and application virtualization means many more applications – and their associated users – are all simultaneously accessing storage volumes to search for and read and write data. These factors and others are contributing to a data explosion and are driving the need for high-performance data retrieval rates.

As a result, many companies are adding storage capacity. For example, 2009 *eWEEK* research notes that shipped storage capacity grew 27.3 percent from 2008 to 2009. To meet these increased performance and capacity requirements, many organizations are turning to storage area networks (SANs).

However, even with the demands being placed on storage systems today, one aspect that still gets overlooked is the role SAN disk defragmenting plays in improving performance. In particular, companies overlook the performance degradation due to file system fragmentation within a SAN.

Additionally, the rapid adoption of virtualization, and to some extent the gradual adoption of cloud computing, will add to the problem since they both increase fragmentation of data and disk volumes.

How are companies addressing this? Some simply ignore the issue, believing that because of the inherently high performance of SANs, defragmenting offers little or no benefit.

Others assume that SAN or virtualization vendors address the issue with features within their offerings. However, while some SAN vendors offer tools that work within their proprietary SAN file systems, none of these tools address local disk file fragmentation.

These wrong assumptions cause many to simply ignore the problem – a tactic that is not an option in today's business environment. Companies need to defragment their SAN to achieve a noticeable improvement in system performance.

AREAS WHERE DEFRAGMENTATION CAN IMPROVE SAN PERFORMANCE

A common way in which companies improve SAN performance, particularly I/O performance, is to simply add more disks.

While SAN manufacturers would certainly not discourage this solution, it might not be the optimal solution for many companies. Instead, let's examine how performance also can be improved by defragmentation – without having to purchase additional hardware.

Companies need to defragment their SAN to achieve a noticeable improvement in system performance.

SAN defragmenting offers an alternative to costly hardware costs incurred by adding drives. Why? While SANs offer extremely efficient and high-performing data storage, it is not their job to address file system-level fragmentation. No matter how efficient data retrieval can be, and how much physical disk limitations can be mitigated, the added overhead on the operating system retrieving the file is beyond the scope of SAN technology, and is impacted by file fragmentation.

While defragmenting can improve the performance of direct-attached storage and network-attached storage, a more cautious approach must be taken with SANs because of the interaction of the various systems that come into play. These systems include operating and file systems, storage virtualization systems, and the SANs file system itself.

Let's look at an example of the interplay between these elements, and how this might complicate defragmentation efforts.

Traditionally, when a new application is deployed, storage is allocated and the requisite amount of disk space is set aside and made unavailable to other applications. This approach makes inefficient use of storage capacity especially today in light of the ease in which virtual machine instances of applications can be created, deployed, and removed.

To address this problem, many companies use a thin provisioning approach where blocks are allocated in real-time as data is written, versus the traditional method of allocating all the blocks up front. However, the efficiency of thin provision-

ing can be lessened due to the way some operating systems and file systems allocate space as the data is being written. For instance, while the storage management system may allocate space using thin provisioning, the file system might simply write data wherever it finds space. If data is written to a “high” logical cluster number (say, cluster 200), all clusters from zero to 200 are then allocated even if they are not used. As new files are added, or data is added to an old file, or data is deleted, this difference between file system disk allocation and storage system thin provisioning can contribute to fragmentation, over-allocation, and less efficient use of storage space.

Defragmenting a thin provisioned volume, if done correctly, can return deleted space to the SAN, thereby preventing over-allocation by the file system. However, when undertaking such an effort, companies should check with both the SAN vendor’s technical experts and the defragmentation software vendor to ensure proper configuration.

Similar interplays between systems must be addressed when SANs incorporate RAID. Defragmenting logical drives will improve the speed and performance of a RAID environment by eliminating wasteful and unnecessary I/Os from being issued by the file system. This occurs because the file system sees the files and free space as being more contiguous. The file system thus will spend less time checking file attributes, which means more processor time can be dedicated to doing useful work for the user and application.

DISKEEPER CORPORATION AS YOUR TECHNOLOGY PARTNER

Diskeeper Corporation has been providing defragmenting solutions since 1981. Its solutions are used to improve the performance and reliability of Windows-based laptops, desktops, and servers.

Performance gains derived from Diskeeper® solutions include:

- faster application boot times for Microsoft® Outlook, Excel, and Word applications
- anti-virus scans or backups are routinely cut in half
- routine Microsoft Exchange tasks, such as opening multiple messages and moving folders, are completed significantly faster
- times for common Microsoft SQL tasks, such as simple and complex queries, table creation, and bulk insertion of rows of data, are significantly reduced
- reliability improves, crashes are prevented, and computer up-time increases, making SAN defragging an absolute must.

Recently, companies seeking to cut electricity usage for cost savings or green initiatives have reaped numerous benefits

that Diskeeper solutions bring to the table. Notably, their computing is automatically made more energy efficient as the disk drive no longer works overtime searching through thousands of fragmented files. This translates to wattage savings due to the notable reduction in power consumption.

SAN defragmenting offers an alternative to costly hardware costs incurred by adding drives.

Additionally, Diskeeper solutions incorporate technologies that are particularly useful for defragmenting SANs. These technologies include:

InvisiTasking® technology, which guarantees that Diskeeper only uses system resources that would otherwise have been unused. This allows real-time defragmentation to continually operate while ensuring your system’s performance is not degraded.

I-FAAST® 2.0 (Intelligent File Access Acceleration Sequencing Technology), which improves file access and creation on NTFS volumes by up to 80 percent above the improvement provided by defragmentation alone. Specifically, when I-FAAST is enabled on a volume, Diskeeper runs specially engineered benchmarks on the selected NTFS volumes to learn their individual performance characteristics. Diskeeper then transparently monitors these volumes for file access frequency on an ongoing basis so as to determine which files are requested most often.

Intelligent Defragmentation, which relies on Diskeeper’s logic controller to detect various volume and system conditions. This information is then used to choose the most effective algorithm to deliver performance benefits faster and more efficiently.

Titan Defrag Engine™, which is a defragmentation engine designed specifically to work with the large volumes of data found on today’s SANs. The engine optimizes defragmentation of data volumes of 100 GB to 20 TB or more.

All of these technologies will help companies make their SANs run more efficiently and boost the performance of their applications. In today’s business world, these benefits will help IT organizations better manage the explosion of data and adoption of virtualization, and cut costs. ●

For more information and to download a free 30-day full trial of Diskeeper, visit: www.diskeeper.com

Salvatore Salamone is an executive editor in Ziff Davis Enterprise’s Strategic Content Group. He is author of three business technology books and has been writing about science and information technology for 20 years, serving as a senior editor at many major publications including *Network World*, *Byte*, *Data Communications*, *LAN Times*, and *InternetWeek*.