Windows IT Pro®

Maximize the Performance of your Windows SAN Infrastructure

by David Chernicoff

Diskeeper corporation

→ Contents

Maximize the Performance of your Windows SAN Infrastructure

By David Chernicoff

As businesses grow and find that the demands for network storage keep increasing, they follow a natural progression of storage environments. As direct-attach storage (DASD) on their Windows servers becomes insufficient, businesses frequently move on to network attached storage (NAS); or, if and when their storage needs mandate it, storage area networking (SAN).

The full promise of storage that is not directly attached to the server is best fulfilled by implementing a SAN, which has advantages in performance, reliability, availability, and provisioning. But the overwhelming benefit of SAN storage often gives storage administrators the false impression that by simply implementing a SAN and following the vendor's instructions they will get the best possible performance and reliability from the SAN environment in their Windows Server-based network.

But at some point the SAN administrator realizes that the storage no longer performs as well as it once did. Investigation into the cause of the performance slowdown usually indicates problems with free space or available storage. The administrator may wonder what is causing all the additional unexpected I/O; but the easy answer, increasing the amount of storage available, is usually the one accepted. However, in many cases adding additional storage is unnecessary because the amount of storage on the SAN is not the problem.

The reality is that although SAN vendors do an excellent job of optimizing the performance and reliability of the networking component that they control (the storage), the SAN, by the nature of its design, is server operating system agnostic. This means that the SAN has no control over how the operating system treats its storage. The SAN's role is to provide the storage; the operating system's job is to deliver that data in a manner that works best. And because—once

again by design—the SAN appears to the operating system as locally attached storage, the way that the operating system writes files to storage is not optimized for the SAN environment. The very nature of the entire decoupled storage model precludes the operating system from being aware of the type of storage it's using. This limit's the operating system's ability to optimize for a particular storage model.

Often, storage administrators are unaware that knowledge acquired from managing DASD storage is, in many ways, still applicable to issues that occur with SAN storage. One of the most significant potential issues, and probably the most unrecognized, is defragmentation in a SAN storage system. With the implementation of a SAN in their storage environment, many Windows Server administrators believe that defragmentation, which they accepted and dealt with when using DASD storage, has gone away.

Don't Forget About Fragmentation

Due to this belief, the storage administrator frequently fails to consider fragmentation when evaluating the problem of reduced SAN efficiency. When SAN performance problems are noticed, the knee jerk reaction is to add additional storage, with the presumption that the lack of free space is causing the SAN to slow down. However, both overall performance and storage efficiency can be affected by fragmentation that occurs when Windows Server writes data out to storage, regardless of whether that storage is DASD, NAS, or SAN.

Much like the behavior that has been seen when disks are fragmented in DASD storage, the effects of fragmentation in SAN storage often manifest in reduced application performance and inefficient use of storage. Application response time begins to increase, the time necessary to load large files grows longer, and the overall user experience is impacted. Applications themselves, if loaded from the SAN storage, take longer to launch. End-users begin to feel that their computer is slowing down, leading to help desk calls with complaints about network performance or some other imagined problem. The reality, however, is that the fragmentation is simply causing data manipulation times to increase to the point where the delay becomes perceptible to the end-user.

Because the SAN is unable to manipulate the data efficiently, this results in additional overhead on the SAN when data is being manipulated; drives are written more than necessary and the amount of work that the SAN must do to service the write requests increases.

An unexpected side effect of reduced performance is a result of the way that performance issues are often handled in a SAN environment: When performance lags, additional spindles are added to the storage device. This approach addresses the performance problem by spreading the load over a larger number of high-performance drives, but adding hardware to resolve SAN performance problems means that more money is spent masking the performance problem rather than addressing the root cause of the performance hit: fragmentation.

As the additional storage is filled with fragmented files, the problem manifests once again, continuing the ongoing cycle of throwing money at the problem and only masking the symptoms for the short term.

The behavior that manifests due to data fragmentation is why you should implement defragmentation software on the Windows servers that use the SAN. Each SAN vendor has an approach to ensure that data is written to SAN storage in the most optimized fashion, but operating system-caused fragmentation is not something that the SAN is able to deal with. It's important to note that fragmentation is not the result of the SAN's behavior; the problem lies at a higher level in the storage stack.

For every I/O request an application makes, the following process takes place:

- The application requests a file read
- The request is passed to the operating system
- The file system maps the file clusters to a logical block address (LBA) and passes that information to the driver for the host bus adapter (HBA)
- The physical disk controller maps the request to a specific block
- The specific blocks are acquired from the physical disk
- The disk controller acquires the blocks
- Blocks are mapped to the LBAs and passed to the file system
- The file system maps the LBAs to the file clusters and passes the data to the requesting application
- The application receives requested data

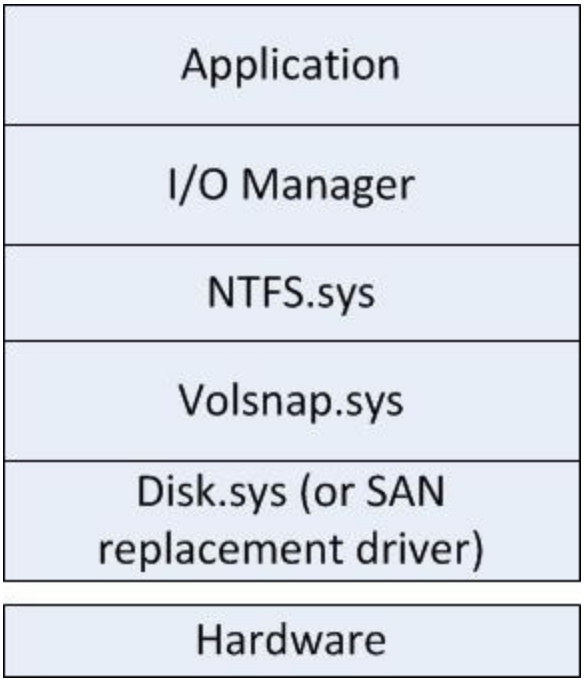

Figure 1: The Windows Storage Stack

As you can see, within the server operating system, there is significant interaction with the data for every simple I/O request. A number of disk I/Os take place for even the simplest request. Fragmented data increases the amount of disk activity necessary to service even that basic file read request, slowing down the apparent performance of the application.

As Figure 1 shows, a number of operating system components touch the data before it makes it to the SAN to be written. These are the same components that touch data as it's written to disk regardless of whether the storage is DASD, NAS, or SAN. The way that the data is finally made available to the disk to be written is one of the major causes of fragmentation, regardless of the eventual storage medium.

The Benefit of Defragmentation

There is little argument that running a defragmenter on a server with locally attached storage is critical. In the case of NAS storage—specifically NAS appliances based on Microsoft Windows Storage Server—the NAS-provided storage is basically another Windows Server in your network environment. As such, from the perspective of defragmentation, it should be treated as just another Windows Server on the network with DASD.

Like any other Windows Server that has its own storage, the benefits of installing a defragmentation utility on that storage server are the same as they would be for any other Windows Server—including improved server uptime and optimized performance. But because the NAS storage is more broadly used, the benefits are actually greater when considered in terms of the overall performance of network-attached storage.

With the SAN environment, a number of benefits accrue by using a defragmentation utility on your Windows Servers that write to SAN storage. Application consolidation and virtualization mean that servers are going to be placing greater demands on their storage, and it becomes even more critical for data to be written to the SAN in the most efficient way possible. On a single physical server a dozen VMs might all be running their own copy of Windows Server and treating the SAN as their own local storage, unaware of other clients operating on the same physical host. Each VM thinks that the storage is in a quiescent state, when the reality may be that intensive disk I/O operations are being generated by another VM.

With the virtualized storage environment provided by your SAN you need to consider three key points:

- The virtualized storage environment is treated as local storage by all servers and clients that have direct access to the storage
- The SAN hardware vendor provides its own tools and techniques for optimizing the performance of the SAN hardware and SAN file system
- The SAN file system is not the same as the Windows Server file system

Because the SAN file system is managed at the hardware level by the SAN, and any client that attaches to the SAN treats the SAN as local storage, the SAN manages its file system without regard for the way that the server operating system writes data. This means that the source of the fragmentation needs to be addressed; the data needs to be written to the SAN in a fashion that minimizes its fragmentation as it's written, allowing for fewer, yet larger disk I/Os to be passed to the SAN.

This results in reduced operating overhead and increased I/O efficiencies. Failure to address fragmentation at the source leads us back to the

performance problem discussed earlier: reduced performance of applications that need to read and write data to the SAN and the impact on the end-user. And we circle back to IT being asked to solve problems that are, at the bottom line, caused by the failure to properly optimize data being written to the SAN at the operating system level.

David Chernicoff is a technology consultant with a focus on the mid-market space and a *Windows IT Pro Magazine* Senior Contributing Editor. David has been writing computer-related feature and product reviews for more than 20 years and is coauthor of a number of operating system books, ranging from the *Windows NT Workstation: Professional Reference* (New Riders Publishing), to the *Microsoft Windows XP Power Toolkit* (Microsoft Press), as well as over a dozen eBooks on topics ranging from network switching topologies to production FAX technology.