Diskeeper And Virtualization

Introduction:

Given that virtualization technologies have many specific applications this paper will begin by first presenting definitions.

Definition: Virtualization

Essentially to virtualize something means to make something that doesn't actually (physically) exist *appear* to exist. Think of the context in *virtual reality*. Let's make a quick example of something everyone in IT is familiar with, a PC with 4 logical volumes (C, D, E, and F). In reality that desktop has one physical disk drive partitioned into 4 volumes. A logical volume is in this case a virtual drive.

Next we'll define two popular modern applications of virtualization technology.

Definition: Server/Machine/Hardware Virtualization

Machine virtualization describes the creation of one or more isolated virtual instances of a "guest" operating system either on top of a "host" operating system (Hosted Architecture) or directly on top of a specialized thin software layer called a hypervisor (Hypervisor/Bare-Metal Architecture).

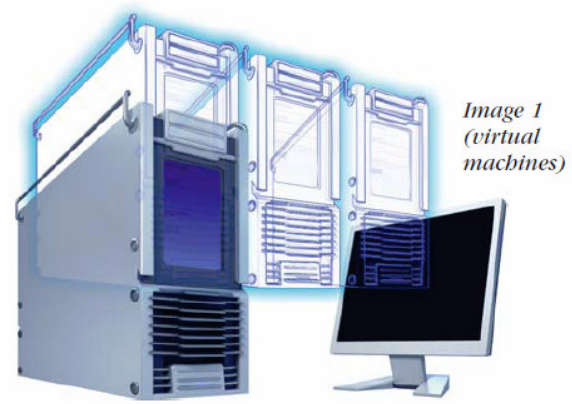

Image 1 (virtual machines)

In either architecture, the host system's virtualization of other operating systems is accomplished by software, proprietary to the vendor (e.g. Hyper-V™, ESX™), which resides between the physical hardware (CPU, memory, etc) and the "guest" operating systems. Each guest or host operating system runs its own applications independently, as if it were the only system operating on the hardware.

Definition: Storage Virtualization

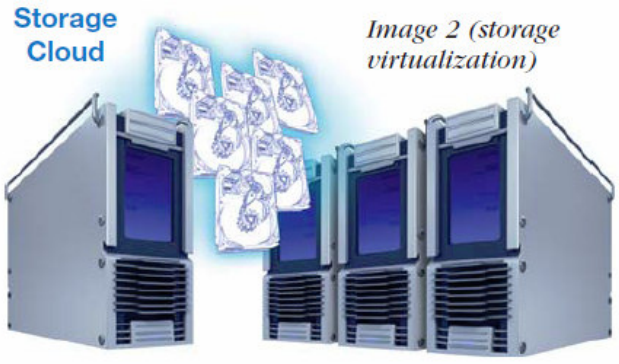

Image 2 (storage virtualization)

Storage virtualization involves the creation of a usually very large, logical-pool of data. Via software, that pool appears to be physically located all on one server. In actuality, that data may be located across hundreds of physical disks spread across dozens of servers. This is the concept implemented by Storage Area Networks (SAN). For peak performance these storage pools require automatic disk defragmentation just the same as a single hard drive would. Automatic defragmentation is implemented from server(s) that manage the respective logical disk volumes.

Our last definition is a broad explanation of disk fragmentation.

Definition: Disk Fragmentation

Disk fragmentation, is the condition in which pieces of individual files and free space on a disk are not contiguous, but rather broken up and scattered around the disk. This requires the hard drive to locate all the fragments of a file. The collection of file fragments from numerous places instead of just one causes file access to take significantly longer than it should. File writes into fragmented free space, also take longer and can increase the likelihood of newly created files fragmenting.

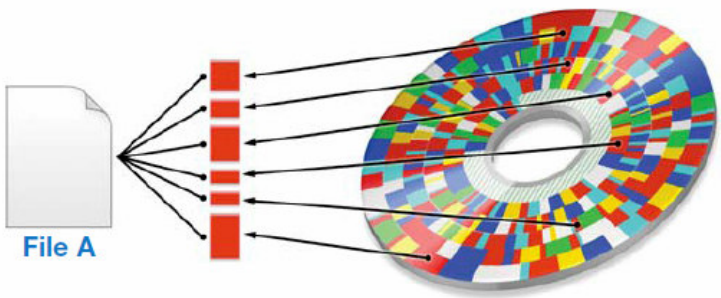

Image 3 (disk fragmentation)

The affect of disk fragmentation is slower system performance, increased I/O overhead, and more severe cases, compromised reliability resulting in phenomena such as application and system hangs and crashes.

The purpose of *defragmentation*, on the other hand, is to consolidate file fragments into a single extent, increasing access speed, and to reduce free space fragments to a small handful of larger chunks. The benefit of defragmentation is reduced I/O overhead and improved operational performance.

The Disk is the Weak Link:

CPUs and memory operate orders of magnitude faster than mechanical hard drives. The slower the disk, the slower the entire system will be. While these facts are well known to industry professionals, it deserves re-iteration as the issue becomes manifest when those *disks* are asked to do more. Such is the case with virtualization, where the given hardware has to support numerous simultaneous operating systems.

Another vital factor to consider is that server virtualization can compound disk fragmentation; and as we covered, disk fragmentation slows disk performance.

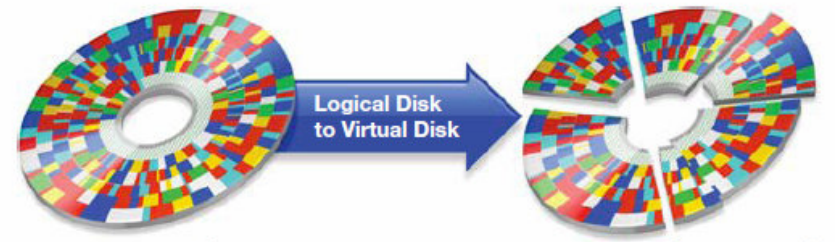

Image 4 (fragmented drive)

Image 5 (fragmented virtual disk)

Typically fragmentation occurs on logical disk drives, and by device drivers is translated to physical sectors on a disk. It can be demonstrated as pieces of a file located in a non-contiguous manner (Image4). In the case of virtual systems, the logical volume is masked by the technology; known as a virtual disk. These virtual disks reside on logical disks in the form of container files. Those virtual disk files can fragment just as any other file can resulting in what amounts to a "logically" fragmented virtual hard disk (Image5), which still has typical file fragmentation contained within it.

The picture represented in Image5 would appear in a defragmentation analysis report's "Fragmented Files" list run from a host Windows operating system as "VirtualServer1.vhd, 4GB in size, in 6 pieces".

This equates to hierarchical fragmentation or more simply fragmentation-within-fragmentation. The black lines in Image6 represent disk I/O mappings of the virtual disk file fragments to the host system in a Hosted Architecture. The smallest unit of data access in a virtual machine is typically 128 sectors, or 64KB. Therefore if these access units (called *grains* in VMware) are fragmented, performance suffers.

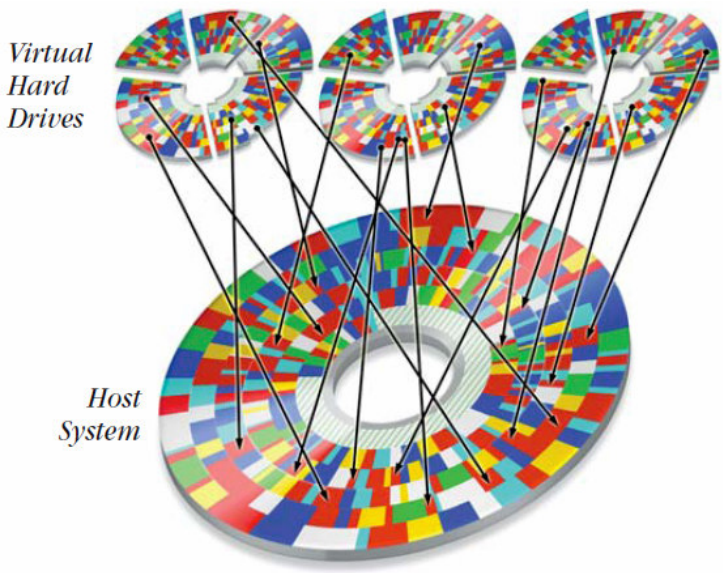

Image 6 (hierarchical disk fragmentation)

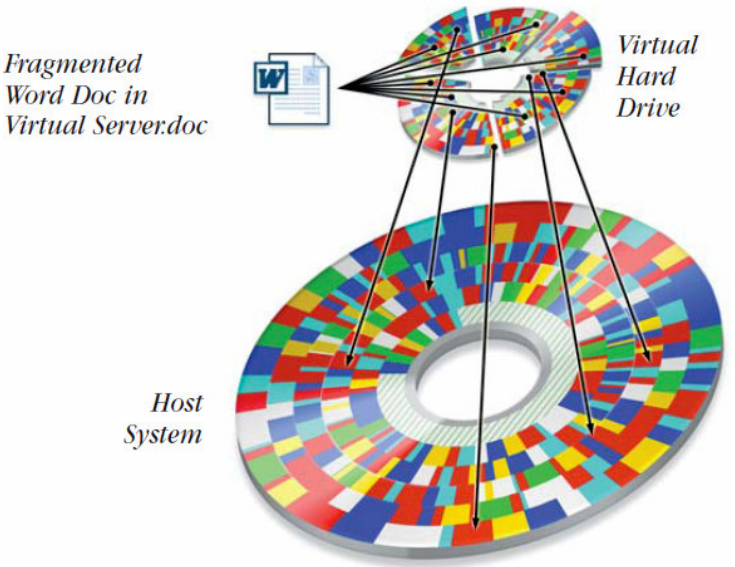

Image 7 (tiered disk I/Os to a fragmented file in a fragmented virtual disk)

Image7 depicts a fragmented file ("Fragmented Word doc in Virtual Server.doc") residing on a virtual disk, which in turn exists as a fragmented file on the host operating system. The current design of software-based server virtualization requires the host system capture and process any disk I/O generated by guest operating systems, adding an additional layer in the I/O processing stack.

Machine Virtualization Architectures and I/O:

Given either of two predominant virtualization architectures (Hosted or Hypervisor) remember that the virtual machines are emulating hardware and may not emulate the exact specifications. For example, a high-end video card may not be emulated in a host system with all the advanced capabilities.

The Hypervisor architecture removes the requirement for a full host operating system and improves overall virtual systems performance.

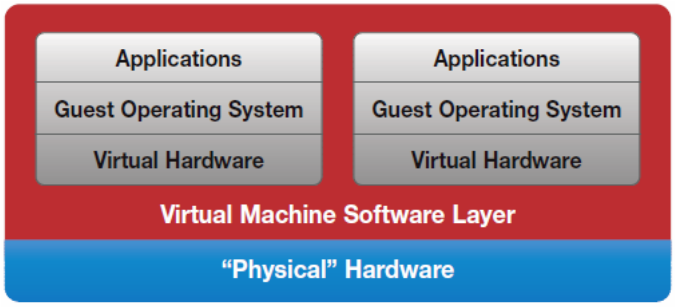

Image 8 (hypervisor virtualization)

As demonstrated earlier, Disk I/O's generated from virtual systems (Hosted Architecture) can suffer from increased software stack processing. This means that disk I/O has to go up and down software layers that abstract the physical hardware. In a Hosted Architecture, a low level disk request in a guest system is translated into a user–level call in the host system. With the likely loss of disk caching at the guest level (hardware support consideration), and limited queuing ability, this process will not be as speedy as a direct physical hardware call by the host system.

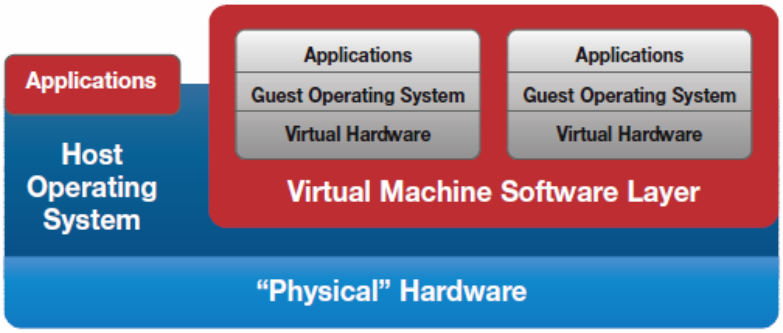

Image 9 (Hosted Virtualization)

In summary, server virtualization establishes a symbiotic relationship, so it is important to remember that generating disk I/O in one virtual machine affects I/O requests from other virtual systems, no matter the architecture or technology vendor. Fragmentation artificially inflates the amount of disk I/O requests which, on a virtual machine platform, compounds the disk bottleneck even more so than on conventional systems.

Regular defragmentation enabled on hypervisors such as Windows Hyper-V, and Windows hosted systems, such as VMware Server and Microsoft VirtualServer, and of course the guest operating systems that reside on ESX Server or any other machine virtualization solution, will dramatically increase performance and decrease the stress and wear and tear on the physical disk. This will enable you to get the most out of your virtual machines.

Diskeeper Corporation
7590 N. Glenoaks Blvd.
Burbank, CA 91504
800-829-6468
www.diskeeper.com