Disk Performance

Diskeeper was kind enough to contribute material on disk performance issues in a virtualized environment. Managing your disk performance ongoing is key to a successful implementation with virtualization. Earlier, we'd created a configuration for the initial investment in performance. Let's now discuss managing that performance ongoing.

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 earlier, disk fragmentation slows disk performance.

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 (left picture). 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 (right picture), which still has typical file fragmentation contained within it.

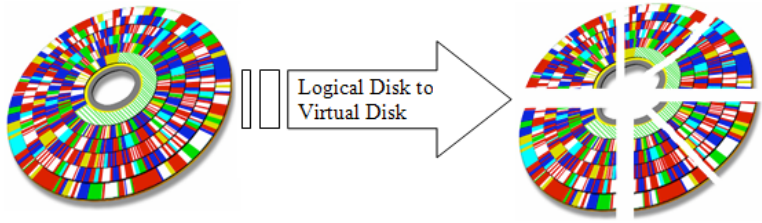

The picture represented in on the right would appear in a defragmentation analysis report's "Fragmented Files" list run from a host Windows operating system as "VirtualServer1.vhd, 4GB in size, in 6 pieces".

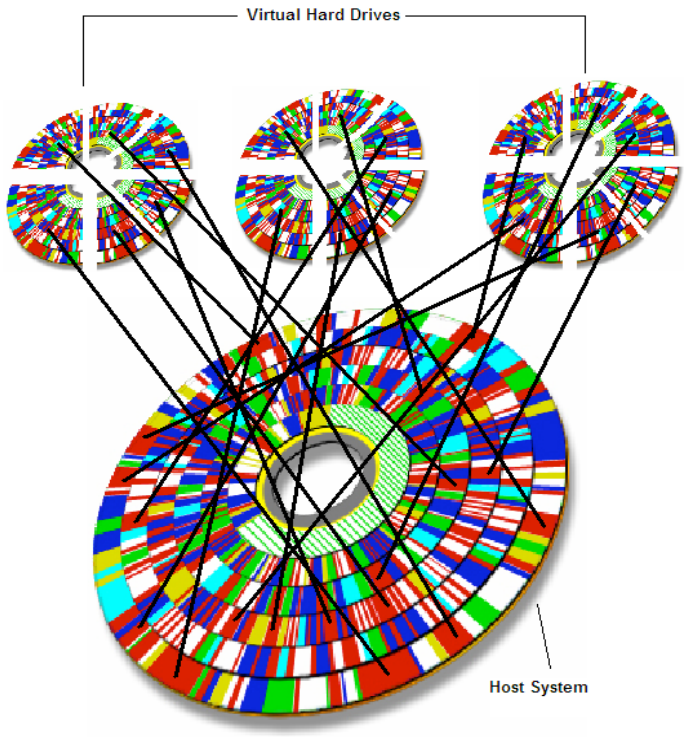

This equates to hierarchical fragmentation or more simply fragmentation-within-fragmentation. The black lines above represent disk I/O mappings of the virtual disk file fragments to the host system in Hosted Architecture. The smallest unit of data access in a virtual machine is typically 128 sectors,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

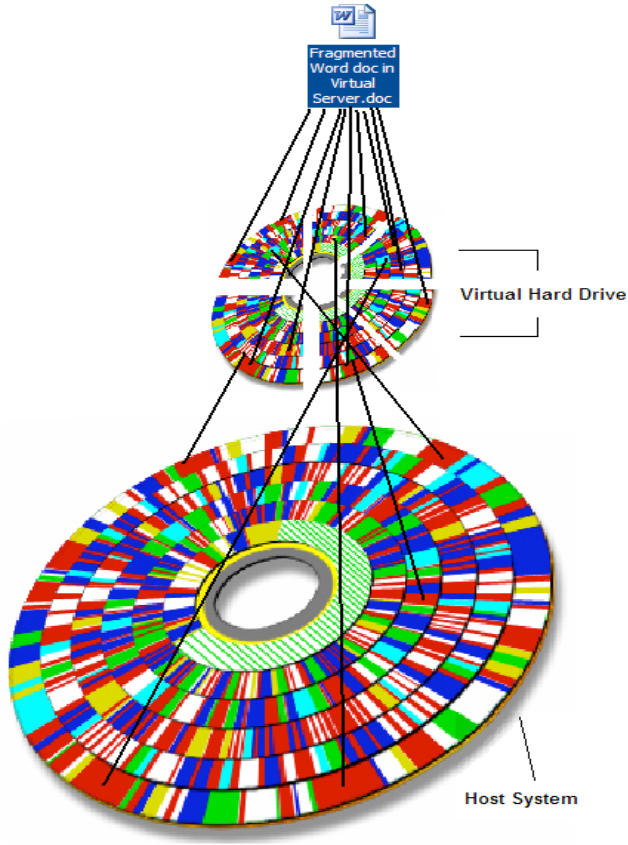

The image above depicts a fragmented file ("Fragmented Word doc in Virtual Server.doc") residing on a virtual disk, which in turn exists as a fragmented file on the host operating system. The current design of software-based server virtualization requires the host system capture and process any disk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

Machine Virtualization Architectures and I/O:

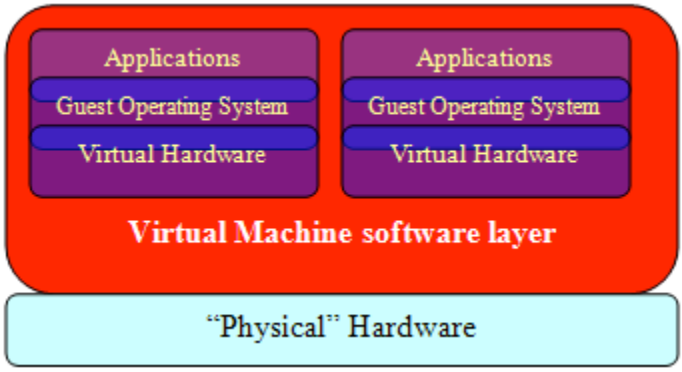

Given either of two predominant virtualization architectures (Hosted or Hypervisor) remember that the virtual machines are emulating hardware and may not emulate the exact specifications. For example, a high-end video card may not be emulated in a host system with all the advanced capabilities.

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

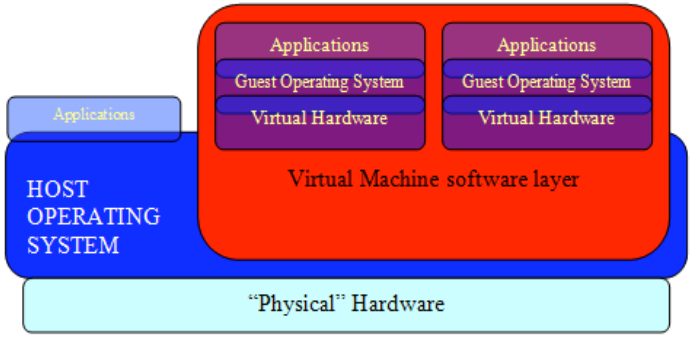

As demonstrated earlier, Disk I/O's generated from virtual systems (Hosted Architecture) can suffer from increased software stack processing. This means that disk I/O has to go up and down software layers that abstract the physical hardware. In Hosted Architecture, a low-level disk request in a guest 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.

In summary, server virtualization establishes a symbiotic relationship, so it is important to remember that generating disk I/O in one virtual machine slows I/O to the disk from other virtual systems, no matter the architecture. Fragmentation is both increasingly substantial in virtual machines environments (hierarchical in Hosted Architecture) and compounds the disk bottleneck more so than on conventional systems (shared resource).

For the future, with the opening of proprietary formats for third party development, virtualization-ready hardware from Intel and AMD (improved hardware support and access), operating system advancements (Hypervisor will be an integral part of Windows Longhorn) and technology partnerships such as that between Diskeeper and Microsoft, look for continuing improvements to disk performance as virtualization gets further entrenched in everyday IT.