Realtime publishers

Tactics in Optimizing Virtual Machine Disk IOPS

The Essentials Series

The Impact of Fragmentation on VM Disk IOPS

sponsored by

Diskeeper corporation

Greg Shields

Introduction to Realtime Publishers

by Don Jones, Series Editor

For several years now, Realtime has produced dozens and dozens of high-quality books that just happen to be delivered in electronic format—at no cost to you, the reader. We've made this unique publishing model work through the generous support and cooperation of our sponsors, who agree to bear each book's production expenses for the benefit of our readers.

Although we've always offered our publications to you for free, don't think for a moment that quality is anything less than our top priority. My job is to make sure that our books are as good as—and in most cases better than—any printed book that would cost you $40 or more. Our electronic publishing model offers several advantages over printed books: You receive chapters literally as fast as our authors produce them (hence the "realtime" aspect of our model), and we can update chapters to reflect the latest changes in technology.

I want to point out that our books are by no means paid advertisements or white papers. We're an independent publishing company, and an important aspect of my job is to make sure that our authors are free to voice their expertise and opinions without reservation or restriction. We maintain complete editorial control of our publications, and I'm proud that we've produced so many quality books over the past years.

I want to extend an invitation to visit us at http://nexus.realtimepublishers.com, especially if you've received this publication from a friend or colleague. We have a wide variety of additional books on a range of topics, and you're sure to find something that's of interest to you—and it won't cost you a thing. We hope you'll continue to come to Realtime for your educational needs far into the future.

Until then, enjoy.

Don Jones

Introduction to Realtime Publishers i

The Impact of Fragmentation on VM Disk IOPS 1

Fragmentation in Virtual Environments: It Only Gets Worse 1

Doesn't Windows Compensate for This? 3

Copyright Statement

The Impact of Fragmentation on VM Disk IOPS

Virtualization's early years found many an administrator focusing attention on processor utilization as primary bottleneck. "Insufficient processing power," we thought back then, "creates a shortfall condition. That shortfall translates directly to poor performance."

Those assumptions weren't necessarily incorrect. Lacking processing power, you will experience performance issues. *Today, however, we realize that disk I/O has a far greater impact than ever before realized*. You probably know that VM performance suffers when hardware doesn't supply enough disk IOPS, or when VMs demand too much. But were you aware that VM and virtual environment configurations can have an impact as well? One critically-important facet of the overall configuration story centers around disk fragmentation's impacts on IOPS.

Fragmentation in Virtual Environments: It Only Gets Worse

You've surely heard the fragmentation story before. Fragmentation as an IT problem has been around since, well, IT has been around. That's because fragmentation is a natural byproduct of normal file system operation. It occurs when a file or folder on disk must be spread across multiple, non-contiguous areas. Figure 1 shows the classic example of a disk at three units of time, starting at the top and moving down.

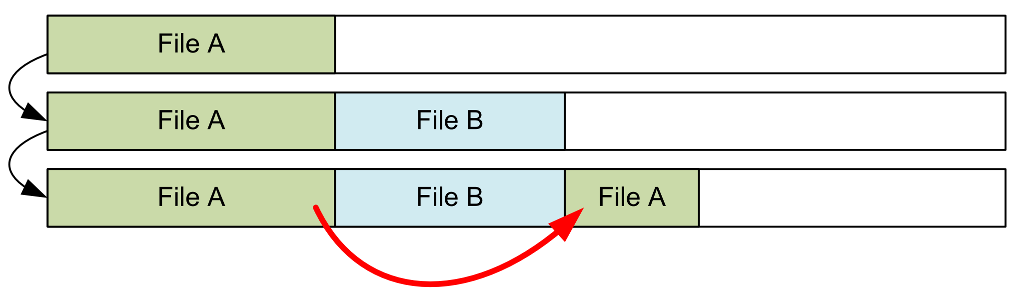

Figure 1: File A must fragment as it grows past its available space.

In this example, File A needs to grow. Perhaps additional data was added or the file was opened, modified, and then closed. File A can't grow contiguously, however, because File B happens to sit in the way. As a result, File A must fragment itself to the next available free space if it is to store its new data.

As files are read and written on disk, this process repeats itself literally tens of thousands of times every week. Files are constantly being added, modified, and deleted, creating "holes" of free space across the entire disk. A deleted file's hole gets plugged with some other file's data. When holes aren't big enough, new data fragments to the next available free space. The problem is a cascading one.

Without protections in place, data can become immediately fragmented as it is written to disk. Existing data fragments further as it evolves, creating a cascading problem where files require incrementally more effort to read and write over time. Figure 2 shows another disk representation, where data in red has grown fragmented from write, modify, and delete operations. Notice the holes in grey. They're the free space, the "holes" where data will end up next.

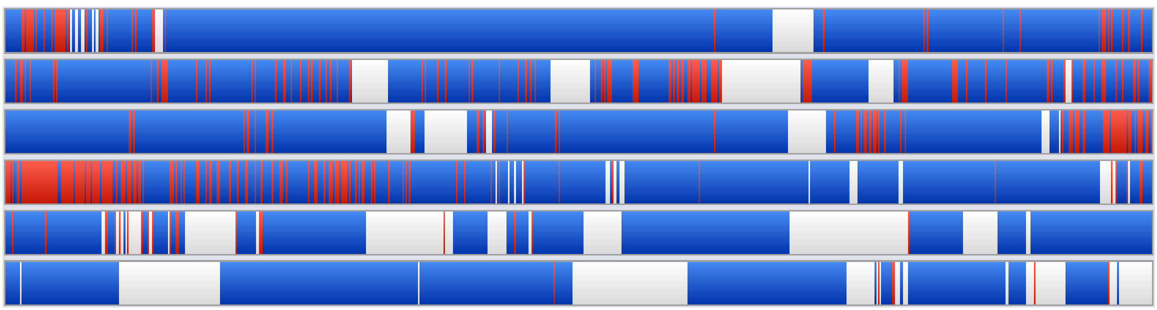

Figure 2: File fragments, represented here in red, get worse over time.

This situation is obviously problematic when experienced across a single computer system. Without compensation, its impact will grow to reduce storage performance, and eventually become a bottleneck for that server's operations. Now add virtualization to the mix. *This problem multiplies when computer systems are used to host other computer systems*, the exact configuration that defines virtualization.

A VM's virtual disks can experience fragmentation just like any physical server. The same dynamics of file creation, modification, and deletion that create fragmentation's performance impact on physical computers follow the same behavior inside VMs as well.

This situation is particularly insidious because it layers fragmentation atop more fragmentation. Figure 3 shows a representation of this multiplicative effect. There, the virtual host is experiencing disk fragmentation. Its files are written at the same time the virtual disks of other computers, combining host fragmentation with VM fragmentation. Significant performance loss is often a result, with VM files requiring extra attention by the VM's file system—which in turn requires extra attention by the host's file system.

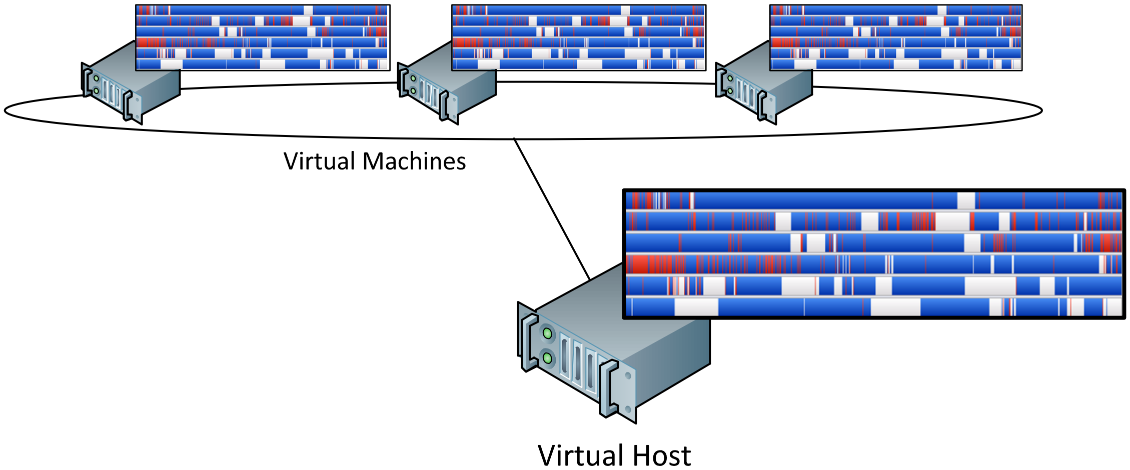

Figure 3: Fragmentation can occur at the host and inside the VMs.

Doesn't Windows Compensate for This?

Doesn't Windows operate with this idea in mind? Indeed it can. Built directly into Windows is a defragmentation feature that will schedule regular passes to reduce fragmentation. This feature is enabled by default on desktop versions of Windows such as Windows 7. It is not enabled, however, on server versions like Windows Server 2008 R2.

This disabled-by-default configuration should beg a question: If fragmentation is such a problem, why would the feature that "fixes" it be disabled by default on server OSs? A primary reason centers around the resource load defragmentation creates. *Eliminating fragmentation by focusing on defragmentation sometimes taxes server resources too greatly*. The extra resources required to even perform defragmentation can themselves impact server operations.

Now multiply these necessary resources across each VM residing on top. You can immediately see how Windows' native approach might not make sense for resource-constrained virtual environments. Of course, all is not lost. This series' final article attempts to find that best-fit solution. In it, you'll learn more about the holistic requirements a virtualization-friendly disk optimization must fulfill.