Realtime publishers

Tactics in Optimizing Virtual Machine Disk IOPS

The Essentials Series

Defining Requirements for a VM Disk Optimization Solution

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

Copyright Statement

Defining Requirements for a VM Disk Optimization Solution

This series has highlighted the design tactics for best optimizing VM disk IOPS. Following the principles presented in the first two articles will ensure your hardware best meets the demands of your VMs. This series has also explored the specific issue of fragmentation that exists irrespective of how your design is ultimately constructed. Getting to maximum IOPS for VMs requires seeking that best design. It also requires a careful look at the configurations you apply to VMs and virtual hosts so that their activities don't foment poor performance. Fragmentation and its elimination are components of those configurations. More importantly, however, is the recognition that solutions for disk optimization may be a necessary part of your overall design.

Disk optimization, particularly with virtual environments, is more than just defragmentation. Performing disk optimization correctly also requires optimizing free space—those proverbial "holes" inside disks. It requires fragmentation prevention, stopping the problem before it happens. It also demands an orchestration of activities across host and collocated VMs, ensuring that optimization activities themselves don't become an impact on performance. As you look towards options for disk optimization, consider the following four important requirements as your specifications for a virtualization-friendly solution.

Requirement #1: Fragmentation Prevention

Defragmentation is by nature a reactive solution to a naturally-occurring problem. The "de" in defragmentation highlights the fact that such solutions must first wait for the problem to occur. Only after fragmentation has occurred can a defragmentation solution begin cleaning up the mess.

A central problem with the reactive approach lies in the effort required to reverse the damage once done. Run a defragmentation pass weekly, and you've got a week's worth of harm to undo. Just cleaning up the mess requires disk attention that impacts IOPS supply. That attention will get in the way of regular VM operations.

Unfortunately, *timing isn't everything with the reactive approach*. Run it hourly, and the problem's scope might grow smaller, yet the effort requires more regularity. Those are still resources lost. The tradeoff between the amount of damage done and the period between resolution can never really find a functioning balance.

These inefficiencies in balancing time period and effort suggest that a proactive approach might be superior. In layman's terms, a proactive approach is akin to running defragmentation constantly, with the elimination activities occurring at the very moment data is changed. This fragmentation prevention approach reduces the extra effort placed on storage by simply laying down data correctly the very first time. Your selected optimization solution will benefit from being proactive.

Requirement #2: Virtual Environment Orchestration of Activities

Alas, fragmentation prevention alone cannot eliminate every fragment. Even the most-intelligent software solution can never know exactly what "holes" will be necessary at which locations every time. Computers are deterministic, and sometimes users find themselves creating large files or making unexpected changes. Although fragmentation prevention by itself should resolve many issues, there occasionally comes the need for a small amount of extra effort fostered through classic defragmentation.

That classic defragmentation will always have an impact on server operations, whether physical or virtual. Rare, however, is the server that finds itself at 100% utilization all the time. Rarer still is the virtual host with VMs doing the same. It stands to reason, then, that a second solution requirement mandates an intelligent orchestration of activities across virtual host and collocated VMs.

Such a solution should analyze existing server activities to find the time periods of least use. That same approach can analyze activities among all VMs on a host, ensuring that optimization activities in one VM won't cause impacts across others. Physical resources are finite. As such, the best optimization solution will perform its job with a holistic awareness of activities across every part of the virtual environment.

Requirement #3: Free Space Optimization

Recall from the previous article's Figure 1 that fragmentation happens when available free space isn't large enough to support a file's expansion or in writing a new file. When the "hole" on disk hasn't been properly sized, a fragment occurs along with the subsequent need for defragmentation.

A third solution requirement recognizes that free space optimization doesn't necessarily mean creating large unused areas at the back of the disk. It means intelligently sizing holes on disk so that files can naturally expand and be added without automatically fragmenting.

Requirement #4: Support for Special Disk Types

Virtual environment disks also have special needs beyond the physical. Virtual disks come in many forms, each of which requires additional attention if they are to be fully optimized. Two of these forms merit special attention, as lacking compensation for their behaviors can create performance issues no less problematic than fragmentation.

The first disk type requiring attention is the thin provisioned disk discussed in the first article. Thin provisioned disks are designed to start small, only growing when new data requires it to expand. They're great for conserving storage space, at least at first. Yet one problem not well understood by many administrators relates to *what happens when data is removed from these disks*. Thin provisioned disks are indeed designed to grow, but they're not designed to shrink when data removal occurs. Lacking a solution for shrinking such disks, your thin provisioned disks will only keep getting bigger over time. Thus, the first half of Requirement #4 suggests seeking a disk optimization solution that resolves this gap. Such a solution will compact virtual disks after data is removed, ensuring the lowest quantity of wasted space on expensive SAN disks.

A second special disk type is the linked clone mentioned earlier, sometimes also called a differencing disk. These special disks aren't for every application, but they do provide specific benefit in certain circumstances. A common use is for hosted virtual desktops. Being able to provision hosted virtual desktops as linked clones of a central reference image enables the rapid deployment of similar VMs.

Linked clones indeed begin their lives as extremely small files. How different really are two computers that are similar in everything but name? What's not well known is that these disks can quickly grow in size, sometimes even growing to equal the size of their parent disk. This rapid sizing of linked clones stems from many factors, including temporary file storage, creation and deletion of profiles, in addition to fragmentation. Thus, for environments making use of these special disks, the second half of Requirement #4 advises seeking a disk optimization solution that compensates for the rapid sizing behavior of linked clones.

VM Disk Optimization: Often Forgotten, Always Necessary

Disk optimization in virtual environments is absolutely a necessary activity. That optimization comes in many forms. A proper design goes far in ensuring hardware is ready to support the IOPS demand of needy VMs. Correctly configuring those VMs during operations represents another facet.

Incorporating disk optimization tools and tactics comprises the third—and too often forgotten—piece to this story. Lacking disk optimization tools, a virtual environment can find itself losing performance and *you might not even realize it.*