

ELIMINATING THE UNNECESSARY

UNNECESSARY I/O AND ITS IMPACT ON PERFORMANCE

While some IT professionals don't even realize the problem exists, unnecessary I/O can be crippling to application performance—and unfortunately the problem is greatly compounded in virtual environments. This paper looks at the nature of I/O and how I/O optimization software is solving today's toughest performance barriers.

Uncovering the Unnecessary

Unnecessary I/O has long flown under the radar, however it's now becoming an issue IT cannot afford to ignore. This is not only due to the volume and velocity of I/O in the modern environment, but primarily due to megatrends like virtualization, which mixes I/O streams from multiple virtual machines as it is funneled down to storage. The departure from the traditional single application on a single physical server has had a profound impact on the underlying layers that service application demand. Until now, that impact hasn't been fully understood as it relates to the surplus of unnecessary I/O—perhaps the biggest performance barrier in many environments.

In an attempt to solve the resulting performance problems it creates, IT adds more hardware to the mix, which leads to spiraling costs in data centers that consume more budget, space, energy and staff.

The ongoing I/O explosion is creating a gap between I/O demand and hardware performance growth rates, explains Robert Woolery, senior vice president of product management and marketing at ConduSiv Technologies. “Unnecessary I/O not only exists, it has a tremendous impact. While adding hardware might initially ease the pain, the cost rapidly becomes prohibitive and the bottleneck reappears,” he says. “No matter how fast or big the new hardware, you never have enough resources to keep up. We are faced with a paradigm shift in how we look at I/O.”

Gaining Control

According to Woolery, there are two challenges, each resulting in performance barriers. The first occurs when the Windows OS splits files apart when they are written, so a single file can translate into thousands of I/Os. The other challenge occurs during read operations, when even more unnecessary I/O is created as the system reads the same files or blocks of data over and over—going from server to network to storage and back—stealing storage bandwidth and creating application latency.

All this excessive read and write I/O creates an I/O tidal wave that must be addressed directly at the source, before it gets pushed into servers, networks or storage, explains

Woolery. “The majority of the I/O within the data center is unnecessary, so our objective is to prevent it altogether. We stop unnecessary write I/O and optimize what remains to dynamically and non-disruptively improve IOPS and reduce latency.” He adds, “And to optimize read operations, we cache a single copy of the hot data within available server resources then optimize the I/O stream that remains based on application behavior analytics. Optimizing read and write I/O is the only way to avoid burying the organization in hardware.” He concludes, “I/O optimization software restrains the data center footprint—saving floor space, energy, and dedicated resources to manage it.”

ConduSiv’s V-locity[®] comprises two distinct technologies that address these challenges: IntelliWrite[®] and IntelliMemory[™].

IntelliWrite resides within the OS at the top of the stack, explains Woolery. “This means we are as close to the application as possible. The benefit is that we are right where I/O is created, so we can remove unnecessary I/O and ensure I/O creation remains efficient at the source,” he says. “With IntelliWrite in place, rather than splitting a file, Windows writes sequentially—only minimal I/O. When we put data down sequentially, reading it back is faster because the storage system does not have to find all the pieces.”

To address the unnecessary read I/O, V-locity leverages its server-side caching technology, IntelliMemory. Also residing within the OS, IntelliMemory utilizes DRAM to reduce the times data has to travel from network to storage. The objective is to significantly reduce the amount of I/O going back and forth between server and storage by allowing Windows to read data straight out of the DRAM.

“Residing within the operating system allows us to request the data blocks in order, even before the request gets to memory. Our behavior analytics engine learns application behavior to optimize the data in memory based on frequency of use.” IntelliMemory differs from other server-side caching in that it does not rely on prefetching or predictive analysis. “We do not fall prey to false predictions. When you inaccurately predict that the OS or an application will need data blocks it wastes I/O,” says Woolery. “The alternative solutions are further down the stack (i.e., within the hypervisor, server, network, storage, etc.) and lack application or OS awareness.”

Bottom Line

Simply put, it’s time for IT to address I/O. With optimization software that eliminates the tidal wave of unnecessary I/O and organizes I/O traffic patterns to behave efficiently, IT has a smarter approach to performance. V-locity is an I/O optimization solution designed to help protect investments and free up budget dollars for more strategic initiatives. V-locity customers (well over 1,000 at time of print) typically see 50% performance gains through this nondisruptive, nonintrusive technology.

“This is a software-only approach to better performance,” Woolery says. “You can install it in your existing infrastructure and start receiving benefits within an hour.”

Click here to learn more about V-locity I/O optimization software, the free evaluation offer, and the built-in benefit analyzer to track results.

