



I D C T E C H N O L O G Y S P O T L I G H T

The Shift to I/O Optimization to Boost Virtualization Performance

May 2013

Adapted from *Scale-Out Meets Virtualization* by Ashish Nadkarni, IDC #238534

Sponsored by ConduSiv

Enterprises are faced with managing more data than ever before, demanding more from IT. These workloads are driving an increase in virtualization, which has led to an unintended consequence: an explosion of I/O that hinders overall enterprise IT performance. This Technology Spotlight examines the rise of data-centric applications and the massive increase in I/O caused by significant technology trends such as BYOD, Big Data, and virtualization — and how the problem is compounded by the read/write inefficiencies inherent in the Windows environment. This paper describes I/O optimization and how it can eliminate the need to increase hardware resources. In addition, this Technology Spotlight looks at ConduSiv Technologies' V-locity technology, which takes a proactive approach to eliminating I/O challenges at their source.

The Changing Face of I/O: Impact on Virtual and Physical Servers

More data is moving across enterprise networks than ever before — so much so that IDC refers to today's data-centric enterprise and the rise of disruptive innovation as the "3rd Platform" of IT. Virtual environments, cloud services, mobile devices, social technologies, and Big Data have all contributed to the rise in digital information that organizations must manage. The challenge is that all of this data not only must be stored but also must be used to extract value for an organization's competitive advantage. In addition, because much of this data is unstructured, tools are needed to create data about data — metadata — to help enterprises use the information to their advantage, further adding to growth and stressing already overwhelmed IT infrastructures.

Organizations have rightly turned to virtualization to address the problems associated with this dramatic increase in data. By creating more virtual machines (VMs), IT can more efficiently use existing hardware resources and associated capital. However, unforeseen challenges have arisen from VM sprawl. Although virtualization provides added agility and scalability, it's also creating an I/O bottleneck problem. All VMs share the available I/O bandwidth, which doesn't scale well and increases processing demands, creating performance bottlenecks in applications, storage, and servers.

Unfortunately, while hardware price per performance costs continue to come down, performance improvements are not keeping up with the rate of data growth and the need to extract value from that data. This performance chasm forces enterprises to purchase ever-increasing amounts of hardware to handle the increased I/O. As IT buys more storage, not for capacity but to spread I/O demand across a greater number of interfaces, organizations cannot fully capitalize on the promise of virtualization and other important technology trends.

Unnecessary I/O: What It Is, Why It Matters

The I/O problem stems, in part, from the fact that while the number of virtual machine shipments is growing at an average of 25% annually, the number of physical servers shipped is growing at a modest 2–3%. As more workloads are put on virtual servers and heavier workloads are placed on physical servers, this can triple or quadruple the amount of random I/O generated from a single server, burdening the compute infrastructure. Increasingly, the storage controller and disk architectures cannot keep pace with this random growing I/O.

Client virtualization, or VDI, creates further issues with the effect of boot and log-in storms that result in massive storage performance issues, thus limiting the scale of a VDI deployment. These limitations and the increase in data and metadata are a dangerous combination. In this new landscape, how can IT build a smart plan around VMs, infrastructure, and the growing list of business-critical applications that run in the environment?

While enterprises can implement storage area networks (SANs) or scale-out infrastructure, it's difficult to determine I/O demands in such a fluid environment. For example, if application transaction volume increases, database access from the application server increases as well. Limited network bandwidth in peak times and lower-than-necessary database performance may also be bottlenecks.

In an attempt to solve the I/O challenge, IT adds CPU capacity and allocates more disk drives to achieve better I/O performance. However, this approach increases hardware and management costs. While IT can increase capacity or implement load balancing, inherent behaviors associated with x86 virtual environments still cause I/O bottlenecks. For example, files written to a general-purpose local disk file system typically are broken into pieces and stored as disparate clusters in the file system.

The I/O explosion brought about by technology trends and data-centric applications is well understood, but the idea of unnecessary I/O bears further examination. When it comes to I/O and its impact on servers, storage and applications, there are two performance barriers: 1) Windows creating unnecessary I/O traffic by splitting files upon write, which also impacts subsequent reads and 2) frequently accessed data unnecessarily traveling the full distance from server to storage and back.

These two behaviors create a surplus of I/O that prevents applications from performing at peak speeds. In today's enterprise, the problem is compounded as a multitude of random I/O traffic, from a mass of disassociated data access points, is making requests for storage blocks — random and sequential — to a shared storage system. All this unnecessary I/O leads to extra processing cycles that increase overhead and reduce application, network, and storage performance. This extra I/O will slow the speed of not only a single VM but also other VMs, minimizing the benefits of virtualization.

There have been other attempts to solve the I/O problem. For example, solid state disks are faster than traditional hard drives. Organizations have also increased DRAM, aggregated disk spindles with wide striping, leveraged higher spin speed HDDs, and even added solid state storage cards that reside on the host servers. But these options, besides being costly, do not solve the problem. They do not eliminate unnecessary I/O at the source, *before* it is pushed into the network and storage, impacting application and infrastructure performance. What's needed is a way to optimize I/O — eliminating the surplus of I/O at the operating system layer, where it originates.

The Need for a New Approach: Reducing the I/O Explosion at the Source

The I/O problem will continue to grow. IDC predicts that the amount of information that needs to be managed by enterprises will increase 50 times in the next 10 years, and the number of files will increase 75 times. However, with Moore's law slowing from processor speeds doubling every 18 months to doubling every three years, processor performance will grow only by a factor of eight, and storage performance will grow by a factor of four.

As a result, the time needed to find the information that is essential to business operations will increase. Because the effectiveness of server virtualization is intrinsically connected to storage performance, I/O performance between server and storage has to be optimized automatically and cost effectively to help IT deliver data anywhere while meeting service levels, especially in virtual environments. Failing to achieve these objectives will only hamper the ability of IT to support the business in its efforts to cut costs and grow revenue.

In a survey of enterprise IT executives regarding storage challenges, IDC found that organizations are feeling the pressure put upon storage architectures. In addition to archiving the vast amounts of data generated by the digital world, respondents listed supporting VMs as a challenge. The reason is that all of the activity from a virtual machine — as well as metadata stemming from the proliferation of unstructured information — is creating a huge influx of unnecessary I/O.

I/O performance can be improved, however, by changing the behavior of the virtual machine OS, which splits files into pieces unnecessarily. For example, instead of random data blocks, I/O optimization aggregates data blocks sequentially, eliminating the need to restructure these blocks during writes and reads. In addition, active data can be differentiated, enabling it to be cached dynamically within available server memory, which eliminates the need for the data blocks to travel the full distance from VM or physical server to storage and back.

Considering ConduSiv Technologies' V-locity

Burbank, California–based ConduSiv Technologies offers high-performance software to accelerate server and application performance and improve efficiency in virtualized environments. The company's V-locity I/O optimization software addresses critical I/O issues by eliminating application bottlenecks without the need to add server or storage hardware. ConduSiv's differentiator is that its software resides at the top of the technology stack, eliminating unnecessary I/O at the source, where it originates. This means the entire infrastructure derives benefit because only productive I/O is pushed through the VM OS layer. With this unique approach, ConduSiv claims significant IOPS gains and increased application response times.

V-locity is a fully certified VMware-ready and Microsoft Hyper-V–ready solution that automatically, intelligently, and transparently improves I/O performance. The software is designed to increase the throughput of all x86 platforms by eliminating unnecessary I/Os at the source and improves efficiency by optimizing writes for increased bandwidth to storage devices, especially on servers with multiple virtual machines. Designed to leverage a SAN or NAS, the software provides I/O benefit without negatively impacting advanced storage features such as snapshots, replication, data deduplication, and thin provisioning.

ConduSiv's top-of-the-stack approach eliminates unnecessary I/O at the source, before it is pushed into servers, network, or storage.

IntelliWrite Technology

At the heart of V-locity is ConduSiv Technologies' IntelliWrite write I/O optimization technology, which removes unnecessary I/Os created when the OS breaks apart files. The technology recognizes when a file is broken apart and brings the pieces back together into a single I/O and pre-allocates the file so it is written sequentially to disk. This optimizes not only write performance but also read performance when a file is read back. If a file is changed, IntelliWrite will restructure the file for continued write and read performance.

IntelliMemory Technology

In the latest version of V-locity, ConduSiv has incorporated its IntelliMemory read I/O optimization technology. A highly efficient data caching solution, IntelliMemory leverages available server memory to provide faster access to data with dramatically reduced latency. IntelliMemory can offload a significant portion of I/O read operations from a VM's logical disks, and, more important, by offloading physical I/O, IntelliMemory helps improve performance on any VMs sharing the same storage resources via a SAN or NAS. This intelligent caching approach is based on a self-learning, behavior-based algorithm that maintains optimum cache hit rates on idle server resources.

According to ConduSiv, V-locity complements other performance solutions (e.g., SQL caching, server-side solid state cards, and tiered storage) and can significantly increase IOPS by as much as 50%. Because of the improved I/O performance, existing hardware is able to process more workloads in the same amount of time, eliminating the need to purchase more infrastructure to solve bottlenecks. Latency is minimized as fewer data blocks need to travel from server to storage and back.

ConduSiv emphasizes that because V-locity optimizes I/O from the top down, it is specifically tailored for environments that leverage a SAN or NAS. Since less I/O traffic is being pushed through the storage, the SAN/NAS is no longer processing unnecessary I/O requests. SSDs, flash I/O, and other technologies will also see performance benefit because the entire infrastructure is now processing only productive I/O.

ConduSiv's top-of-the-stack approach differs from hardware-based technologies, resulting in a proactive approach to solving the I/O bottleneck issue.

Challenges

ConduSiv Technologies does face market challenges, however. Given the continued rapid data growth and the increasing role of rich content in enterprises, the potential for storage and information bottlenecks caused by virtualization will only increase in the coming years. ConduSiv must continue to enhance V-locity for physical and virtual servers to ensure that the software's capabilities match the exploding data growth.

As virtualization and cloud computing continue to become mainstream enterprise technologies, the problem of I/O bottlenecks will increase, creating more demand for solutions, which will lead to increased competition. ConduSiv should prepare for major virtualization and storage players to offer more solutions by maintaining its technological advantage. The company needs to work at changing the market perception that adding more hardware is the only way to solve the I/O explosion and its resulting bottlenecks.

Conclusion

The continued expansion of information and rich content within public and private clouds and content-centric organizations is changing the business dynamic of enterprises. Extracting value from an expanding universe of digital information is becoming a core business mandate. With the explosive growth of data, much of it unstructured, enterprises continue to face major challenges in how they capture, store, and use information. Although virtualization has been tapped as a key approach to increasing IT performance while maintaining costs, it raises key challenges as well.

Both trends are creating performance bottlenecks by generating unnecessary I/O.

IDC recommends that instead of throwing new hardware resources at performance problems, organizations consider multiple approaches to gain higher I/O performance from existing resources. For example, optimizing I/O at the top of the stack can eliminate many of the issues associated with

extra I/O overhead. Optimizing performance at the source can actually reduce I/O instead of treating the symptoms of excess I/O by simply adding resources.

With higher performance per server, IT can support more users and handle demanding and growing workloads with existing virtualized environments.

To determine the best solutions for eliminating bottlenecks and increasing application performance, organizations should seek vendors that have experience not only in I/O optimization but also in providing solutions for virtualized environments. Naturally, solutions should not add to processing overhead and should address I/O bottlenecks at the source. Indeed, solutions that increase performance in virtual and physical environments will help enterprises realize the full value of their infrastructure investments.

As information continues to grow not only in volume but also in value, virtualization and cloud computing will provide a way for enterprises to manage information more cheaply and efficiently, but self-restricting challenges are associated with these technologies. Organizations must realize that these technologies have challenges and take steps to mitigate these issues by eliminating unnecessary I/O and optimizing I/O performance.

To the extent that ConduSiv Technologies can address the challenges described in this paper, the company has a significant opportunity with organizations looking to regain value lost as a result of I/O bottlenecks and hardware purchases.

A B O U T T H I S P U B L I C A T I O N

This publication was produced by IDC Go-to-Market Services. The opinion, analysis, and research results presented herein are drawn from more detailed research and analysis independently conducted and published by IDC, unless specific vendor sponsorship is noted. IDC Go-to-Market Services makes IDC content available in a wide range of formats for distribution by various companies. A license to distribute IDC content does not imply endorsement of or opinion about the licensee.

C O P Y R I G H T A N D R E S T R I C T I O N S

Any IDC information or reference to IDC that is to be used in advertising, press releases, or promotional materials requires prior written approval from IDC. For permission requests, contact the GMS information line at 508-988-7610 or gms@idc.com. Translation and/or localization of this document requires an additional license from IDC.

For more information on IDC, visit www.idc.com. For more information on IDC GMS, visit www.idc.com/gms.

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com