

WHITE PAPER

Guide to 50% Faster VMs – No Hardware Required

Executive Summary

As much as everyone has bought into the ideology of virtualization, many are still only seeing half the benefit. Promises of agility and elasticity have certainly prevailed. However, many CFOs are still scratching their heads looking for cost-saving benefits. In many cases, IT departments have simply traded costs. The unexpected I/O explosion from VM sprawl has placed enormous pressure on network storage which has resulted in increased hardware costs to fight application bottlenecks. Unless enterprises address this I/O challenge and its direct relationship to storage hardware, they cannot realize the full potential of their virtual environment. ConduSIV Technologies' V-locity VM accelerator solves this I/O problem by optimizing reads and writes at the source for up to 50% performance boosts without additional hardware.

Virtualization, Not All Promises Realized

Not all virtualization lessons were learned immediately; some have come at the expense of trial and error. On the upside, we have learned that virtualizing servers has led to recapturing vast amounts of server capacity once sitting idle. Before server virtualization, the answer to the need for the ever-growing demand for processor capacity was simply to buy more processors. Server virtualization changes this by treating all servers in the virtual server pool as a unified pool of resources – thus reducing the overall amount of servers required while improving agility, scalability and reliability. IT managers can respond to changes businesses need quickly and easily while business managers spend less money on capital and operating expenses associated with the data center – everyone wins.

The Problem: An Increased Demand for I/O

While virtualization provides the agility and scalability required to meet increasing data challenges and helps with cost savings, the very nature of server virtualization has ushered in an unexpected I/O explosion. This explosion has resulted in performance bottlenecks on the storage layer which ultimately hinders the full benefit organizations hope to achieve from their virtual environment. As IT departments have saved costs on the server layer, they have simply traded some of that cost to the storage layer for additional hardware to keep up with the new I/O demand.

As much as virtualization enables enterprises to quickly add more VMs as required, the increase in VM density per physical server is rising, and with it increasing amounts of I/O, further constraining servers and network storage. All VMs share the available I/O bandwidth, which doesn't scale well from a price/performance point of view. Unless enterprises can address this challenge by optimizing I/O at the source, they won't be able to realize the full promise of virtualization due to the dependence on storage hardware to alleviate the bottleneck – a very expensive dependence, as many organizations have discovered.


While storage density has advanced at 60% CAGR, storage performance has advanced at just a 10% CAGR. As a result, reading and writing of data has become a serious bottleneck, as applications running on powerful CPUs must wait for data, forcing organizations to continually buy more storage hardware to spread the load.

This purchase practice is rooted in the unanticipated performance bottleneck caused by server virtualization infrastructures. What has not been addressed is how to significantly optimize I/O performance at the source. Without a solution, organizations are left with an inefficient infrastructure impacting how servers store and retrieve blocks of data from the virtualized storage pool.

What Changed?

Storage hasn't always been a bottleneck. In a traditional one-to-one server to storage relationship, a request for blocks of data is organized, sequential and efficient.

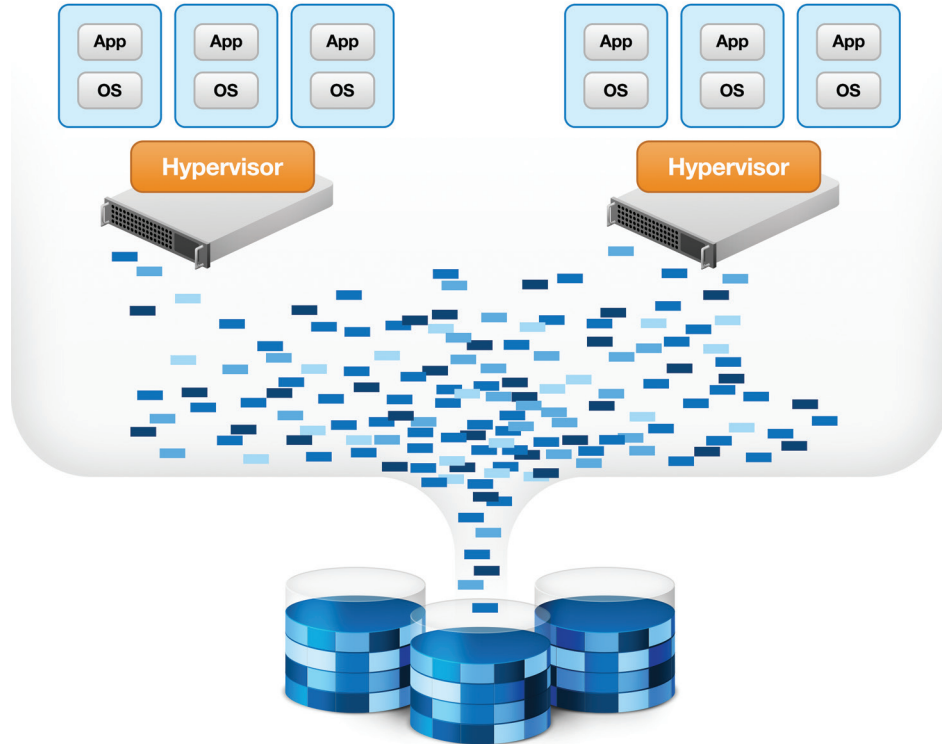
One-to-One Server to Storage



However, add multitudes of virtualized servers all making similar requests for storage blocks, both random and sequential, to a shared storage system, and the behavior of blocks being returned take on the appearance of a data tornado. Disparate, random blocks penalize storage performance as it attempts to manage this overwhelming, chaotic flow of data. In addition, the restructuring of data blocks into files takes time in the server, therefore impacting server performance. Because virtualization supports multiple tasks and I/O requests simultaneously, data on the I/O bus has characteristically changed from efficiently streaming blocks to having the appearance of data that has gone through an I/O blender that is being fed through a small bottleneck.

The issue isn't how untidy the landscape has become. The issue is performance degradation caused by inefficient use of the I/O landscape, thereby impacting both server and storage performance and wasting large amounts of disk capacity.

Multiple VMs without V-locity



I/O Bottlenecks

I/O bottlenecks are forcing IT managers to buy more storage, not for the capacity, but to spread I/O demand across a greater number of interfaces. These bottlenecks are due to unnecessary I/O, growing numbers of VMs, storage connectivity limitations and inefficiencies in the storage farm.

Unnecessary I/O

Inherent behaviors associated with x86 virtual environments cause I/O bottlenecks. For example, files written to a general purpose local disk file system are typically broken into pieces and stored as disparate clusters in the file system. In a virtualized environment, the problem is compounded by the fact that virtual machines or virtual disks also can be fragmented, creating even greater need for bandwidth or I/O capacity to sort out where to put incoming data written to the volume. Each piece of a file requires its own processing cycles – unnecessary I/Os – resulting in an unwelcome increase in overhead that reduces storage and network performance. This extra I/O slows multiple virtual machines on the same server hardware, further reducing the benefits of virtualization. A better option is to prevent files from being broken into pieces to begin with and aggregating these pieces into one sequential file to eliminate unnecessary I/Os and increase overall efficiency of the array.

Increasing Numbers of VMs

After a file leaves the file system, it flows through server initiators (HBAs). As server processor performance grows, organizations tend to take advantage by increasing the number of virtual machines running on each server. However, when I/O performance exceeds the HBA's Queue Depth, I/O latency will grow, which means application response times will increase and I/O activity will decrease. Avoiding I/O throttling is another reason organizations buy extra HBAs and servers making additional investments in server hardware to handle growing I/O demand.

Storage Connectivity Limitations

Storage connectivity limitations further impact performance and scalability. When addressing this issue, organizations need to consider the number of hosts they can connect to each array. This number depends on:

- The available queues per physical storage port
- The total number of storage ports
- The array's available bandwidth

Storage ports have varying queue depths, from 256 queues to as many as 2,048 per port and beyond. The number of initiators a single storage port can support is directly related to the storage port's available queues. For example, a port with 512 queues and a typical LUN queue depth of 32 can support up to: $512 / 32 = 16$ LUNs on 1 Initiator or 16 Initiators with 1 LUN each, or any combination that doesn't exceed this limit. Array configurations that ignore these guidelines are in danger of experiencing QFULL conditions in which the target/storage port is unable to process more I/O requests. When a QFULL occurs, the initiator will throttle I/O to the storage port, which means application response times will increase and I/O activity will decrease. Avoiding I/O throttling is another reason IT administrators buy extra storage to spread I/O.

Storage Farm Inefficiencies

The storage farm presents additional challenges that impact I/O. Rotational disks have built-in physical delays for every I/O operation processed. The total time to access a block of data is the sum of: rotational latency, average seek time, transfer time, command overhead, propagation delay, and switching time. In a standard server environment, engineers offset this delay by striping data across multiple disks in the storage array using RAID, which delivers performance of a few thousand IOPs per second for small block read transfers. In a virtualized environment, however, one can't predict the block size or whether it will be a pure streaming or a random I/O workload. Nor can one guess how many storage subsystems the I/O will span, creating added delay. At the same time, some things are certain to compound the bottleneck: for example, someone will assuredly want to write to the disk, resulting in a write delay penalty; blocks coming out of the array will be random and inefficient from a transfer performance point of view; and the storage subsystem will have enough delay that an I/O queue will build, further impacting performance of the storage subsystem and the server.

While one emerging trend is to use solid-state drives (SSDs) to eliminate some I/O overhead, this is an expensive solution. A better option is to move frequently used files to server memory, intelligently, transparently and proactively based on application behavior, reducing the number of I/Os the storage array must manage.

With the current economic slowdown, failing to overcome the I/O bottleneck will hamper IT's ability to support the business in its efforts to cut costs and grow revenue.

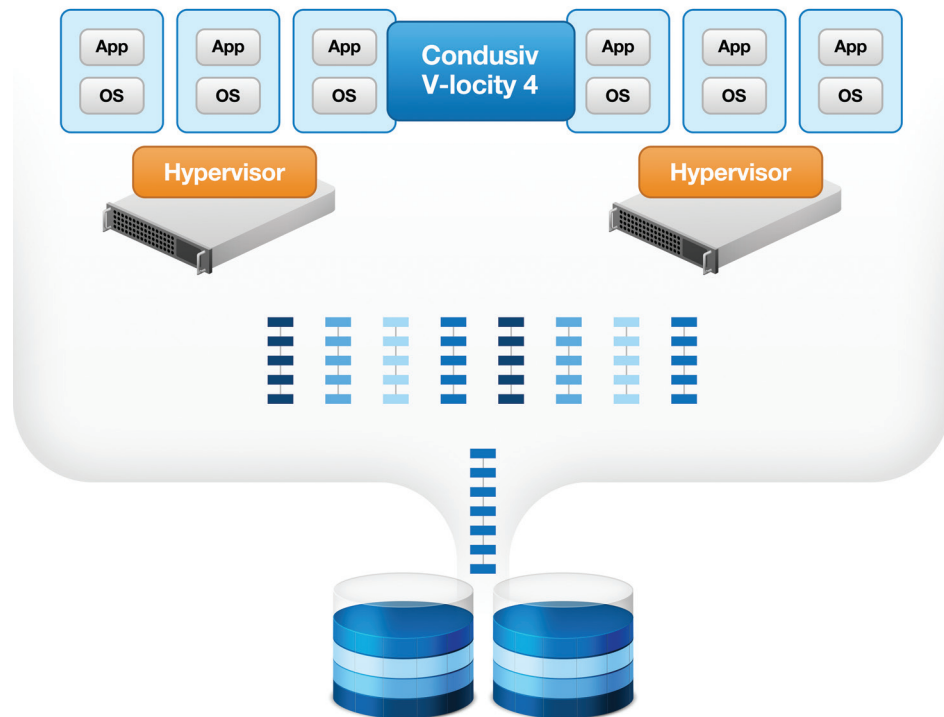
Solving the Problem

The ConduSiv solution to the problem is accomplished in two parts in a product named V-locity 4, containing IntelliMemory™ and IntelliWrite® technology.

IntelliWrite eliminates unnecessary I/Os at the source caused by the operating system breaking apart files, thus improving efficiency by optimizing writes (and resulting reads) for increased data bandwidth to the VM by up to 50% or more.

IntelliMemory automatically promotes files that are used frequently, from disks where they would ordinarily reside, to cache memory inside the server. In this way, IntelliMemory intelligent caching technology boosts frequent file access performance by up to 50% while eliminating unnecessary I/O traffic as well.

Multiple VMs with V-locity



V-locity uses these capabilities to uniquely and intelligently manage disk allocation efficiency and performance. The solution detects files that are used frequently enough to warrant performance and capacity improvements. It then reorganizes data from files that have been stored on disk in a random,

scattered fashion into sequential order (or very close to it) as a normal part of an ordinary write – without the need for IT to take any additional action. Performance improves due to a phenomenon known as locality of reference, in which the heads on the disk don't need to be repositioned to request the next block of data, eliminating additional overhead delays.

Improvements in reference of locality means when data is moved across the I/O bus, there is less overhead and wait time, reducing I/O utilization frequently by 50%. To the end user, overall system performance is seen as improving by 50% or more. The IT administrator will notice that the need to buy additional storage or server hardware has been postponed or eliminated.

Moreover, V-locity I/O optimization is specifically tailored for virtual environments that leverage a SAN or NAS as it proactively provides I/O benefit without negatively impacting advanced storage features like snapshots, replication, data deduplication and thin provisioning .

Conclusions

As a result of using V-locity, VMs can perform more work in less time, more VMs can operate in existing servers, and improved storage allocation efficiency increases performance while reducing expense.

More Work in Less Time

Organizations using V-locity have measured cache hit rate improvements in a virtual machine of 300 percent greater. A test by OpenBench Labs found an I/O reduction of 92 percent in one Virtual Environment with throughput improvements over 100 percent. This reduces the issue of VM sprawl and I/O growth while improving performance without additional storage hardware.

More Virtual Machines Operate in Existing Servers

Because organizations have seen VM performance gains of up to 50 percent, they are able to use more VMs on a single server postponing the need to purchase additional server hardware.

Improved Storage Allocation Efficiency

Tests have found improvements in available free space average 30 percent or greater. Storage in a VM environment is often purchased to improve I/O, rather than to increase capacity. By correcting the storage infrastructure I/O bottleneck, previously unused capacity that was purchased only to spread I/O across more ports can be added to the free storage pool. Since this capacity is already paid for, organizations can use it to eliminate the need to purchase new storage.