openBench Labs
Executive Briefing:

December 5, 2012

Condusiv's V-locity VM Boosts Virtual Machine Performance Over 50% Without Additional Hardware

Improving Storage Efficiency by Increasing I/O Throughput and Decreasing I/O Operations

Executive Overview:

"By running V-locity VM on all Windows-based VMs, IT administrators can significantly increase the number of VMs running on a host and improve application performance without adding additional storage hardware, which enables IT to maximize the ROI associated with any VI initiative."

Why Read This Document?

If you are currently using a VMware® or Microsoft® Virtual Infrastructure (VI) and are avoiding the use of I/O optimization software, or if you are concerned over the deleterious effects defragmentation has on the advanced features of SAN and NAS based storage systems, V-locity® 4 by Condusiv® Technologies introduces an entirely new way to optimize I/O performance and reduce many datacenter operating expenses associated with storage resources, which are often resident on a SAN fabric.

In particular, Condusiv Technologies' V-locity VM optimizes I/O and boosts standard I/O performance metrics from the unique perspective of limiting unnecessary I/O. Not only does this technique afford a VM running V-locity VM with higher I/O throughput and faster latency by limiting I/O processing on SAN and NAS based devices, V-locity VM reduces I/O stress on multiple systems and improves scalability, particularly with respect to VM density and application performance within an enterprise-scale VI.

For this white paper, openBench Labs tested the ability of V-locity VM to optimize all I/O, both reads and writes on a server VM running a Windows-based guest OS. In this assessment, openBench Labs measured the ability of IntelliWrite® technology to prevent unnecessary split I/Os, using its intelligence to create new data files and extend current files as single contiguous collections of logical blocks. In addition, we measured the ability of IntelliMemory™ technology to offload I/O on reads through dynamic caching, in order to boost throughput and reduce latency.

Snapshot of Findings

- **SQL query tests reading and writing 4KB, 16KB and 32KB packets demonstrated an average performance improvement of 129%**
- **Streaming video tests demonstrated an average performance improvement of 131%**
- **SAN/NAS compatibility tests demonstrated no negative impact on advanced storage features (e.g., thin provisioning, replication, snapshots, deduplication)**
- **The embedded Benefit Analyzer in V-locity VM displayed results that matched our Iometer readings within a narrow margin of error**
- **V-locity VM's predictive intelligence compounded performance over time for repeated processes using small-block random or large-block streaming I/O, which greatly reduced network storage traffic**

Executive Briefing:

Condusiv's V-locity VM Boosts Virtual Machine Performance Over 50% Without Additional Hardware

Jack Fegreus
December 5, 2012

Improving Storage Performance by Increasing I/O Throughput and Decreasing I/O Operations

In this analysis, openBench Labs examines the ability of Condusiv Technologies' V-locity VM to maximize I/O performance by efficiently optimizing the way data is read from and written to disk for Windows-based Virtual Machine (VM) servers in a VMware vSphere or Hyper-V Virtual Infrastructure (VI) that are using high-performance Fibre Chanel (FC) and iSCSI SAN-based storage infrastructure. In particular our goal was to evaluate V-locity VM effectiveness on scaling I/O-bound applications on VMs, and scaling VM density on host servers.

Under Test: VM I/O Acceleration
Condusiv Technologies' V-locity VM

1) IntelliWrite® technology provides continuous I/O write optimization using dynamic intelligence when creating or extending files so that random I/Os are optimized into sequential behavior for improved performance on both writes and reads.
2) IntelliMemory™ technology reduces disk I/O requests by predictively caching active data within the server to keep requests from unnecessarily traveling the full distance from VM to network storage, then back.
3) By optimizing I/O processing on VMs with V-locity, more VMs can be run on a host, as each VM will generate less SAN or NAS I/O traffic on an underlying storage system. By using both IntelliWrite and IntelliMemory, we have seen overall application performance improvement over 100%.
4) V-locity provides I/O optimization at the server level to ensure no negative impact to advanced storage features like replication, deduplication, thin provisioning, snapshots, etc.

At the heart of the I/O optimization problem is the difference in the rate of performance enhancement of CPUs versus the rate of performance improvement of storage density and its relationship to the mechanics of disk devices. While storage density has advanced at a 60% compound annual rate, I/O throughput for disk drives and access speeds have advanced at just a 10% compound annual rate. As a result, the reading and writing of data has become a serious bottleneck, as applications running on powerful CPUs must wait for data.

The disparity in storage technology advancements is equally true for direct attached storage (DAS) and a Storage Area Network (SAN); however, the disparity is exacerbated in a SAN and even more so in a VI built on a SAN. Organizations have looked to an improvement in storage connectivity performance to increase the number of VMs on the server level with 8Gbps FC SANs and 10GbE iSCSI SANs. However, the I/O bandwidth requirement demands multiple storage initiators beyond its ability to process despite the availability of greater connectivity bandwidth. As a result, resource sharing, even if only at the level of sharing logical storage units (LUNs) from a common storage array, is a critical performance characteristic of a SAN fabric topology.

With shared resources, inefficient I/O by any VM on a SAN will not be an isolated event and will impact all the other servers and VMs. On a SAN, VM I/O operations cascade down over the operations of all of the other VMs, with which the initial VM shares an array.

More importantly, the rapid adoption of VI environments and the rapid growth of VM sprawl has doubled down on the notion of sharing on a SAN and is pushing the longstanding IT distributed architecture paradigm towards a critical tipping point, i.e. I/O bottlenecks impacting

the full realization of the VI investment. Hypervisors, such as VMware's ESXi, allow multiple VMs on a host to share a common datastore in much the same way that physical machines share a common SAN array as a source of independent LUNs. What's more, the ability of multiple VMs to share a datastore opens the door for multiple hosts to mount a common set of LUNs.

With multiple hosts sharing a common LUN, VMs are freed to move among hosts as a means of providing load balancing and high availability, which makes it imperative that IT is able to scale out storage resources in terms of I/O performance as easily as it does in terms of storage capacity. By concentrating the read and write I/O streams of local VMs, VI hosts present storage resources with far more intense and far less predictable I/O streams that stress both I/O latency and throughput. More importantly, as the demands on storage services intensify in a VI, storage arrays become inextricably linked to the capital and operational expenses that IT must restructure to maximize the return on investment from any corporate VI initiative.

"With multiple hosts sharing a common LUN, VMs are freed to move among hosts as a means of providing load balancing and high availability, which makes it imperative that IT is able to scale out storage resources in terms of I/O performance as easily as it does in terms of storage capacity."

Removing Unnecessary I/Os as Barriers to Performance

Optimum I/O operation involves maintaining the representation of every file within the same logical block space on a storage device as a single contiguous cluster of blocks. However, there are inherent behaviors associated with Windows-based VMs that cause a tremendous amount of unnecessary I/O operations to be pushed down into the network and storage. For example, files written to a general-purpose local disk file system are typically broken into pieces and stored as disparate clusters in different logical locations. Windows cannot issue an I/O request that spans disjointed clusters. When the I/O request does span multiple clusters, it must be split into multiple I/Os – one for each cluster.

As a result, processing all these extra split I/Os generates extra work within the VM, passes extra I/O requests to the VM's host server, pushes extra I/O traffic onto the network and creates extra work for the storage subsystem servicing the VM's host. In a virtualized environment, the problem is compounded as multiple VMs share the same storage resource resulting in highly random I/O behavior pushed down to the disk subsystem. The generation of unnecessary I/Os not only slow the speed of a single Virtual Machine, but others on the same host or those using the same shared storage as well, minimizing the benefits of virtualization.

These problems are forcing IT managers to buy more storage, not for the capacity, but to eliminate I/O bottlenecks by spreading I/O demand across more interfaces or by creating complex, expensive storage tiers using different classes of storage from SATA, to SAS, to Solid-state Devices (SSDs). With the current economic slowdown, failing to overcome the I/O bottleneck challenge will only hamper the ability of IT to support the business in its efforts to cut back costs and grow revenue.

It's this I/O barrier that V-locity VM has been architected to solve. Step 1, V-locity eliminates nearly all unnecessary I/O operations at the source. Step 2, V-locity caches frequently accessed data to keep I/O requests from traveling the full distance to storage and back. We examine the performance improvement of V-locity VM's I/O optimization in the proceeding tests.

Solving Unnecessary I/O at the Source

The first way Velocity VM solves this critical issue is with IntelliWrite® technology to prevent the OS from breaking a file into pieces as each disparate piece requires additional I/O operations. In particular, IntelliWrite adds more intelligence to the way that the Windows OS preallocates file

space and then restructures writes in a coherent manner to prevent a performance penalty. Once selected, IntelliWrite optimizes writes automatically and transparently. Moreover, if a file is accessed and changed in the future, IntelliWrite will at that time restructure the file for optimal write and read performance.

Optimized writes do not resolve all of the important data access and throughput issues when reading data. To provide full-spectrum I/O optimization, V-locity VM introduces IntelliMemory™, a highly efficient data caching solution to provide faster access to data and dramatically improve throughput. IntelliMemory can offload a significant portion of I/O read operations from the disks that underlie a VM datastore. Not only does I/O offloading improve operations on the VM running IntelliMemory, it improves performance for all VMs sharing the same storage resource. Furthermore, it reduces the work performed by the SAN or NAS supporting the datastore.

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

To assess the performance capabilities of V-locity VM by putting both IntelliWrite and IntelliMemory to the test, we set up a series of application-centric scenarios that focused on both random I/O and sequential I/O environments. In addition, we wanted to cover the gamut of application environments in terms of the size of data set. In particular, caching solutions can degrade with larger data sets. With this in mind, we set up two environments, in which the base VM was configured with a resource pool having four core processors and 8GB of RAM; however, the size of the data sets in each of these scenarios differed from the previous test by a factor of 100X.

Random Database I/O Benchmark

We began with a database query activity scenario using a 50MB file and Iometer to simulate repeated access of a large SQL database table at random data points. We ran a sequence of Iometer scripts to set a performance baseline and load the cache with elements of the data file. We then repeated the test to measure the degree to which V-locity could accelerate throughput on reads.

Iometer Query Benchmark

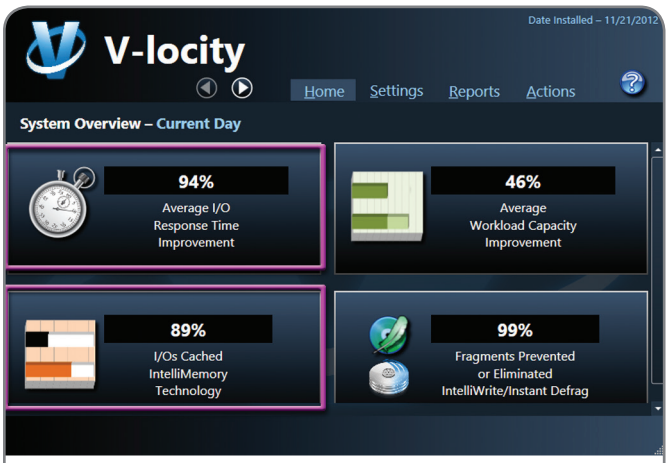

After we ran our SQL query test, V-locity VM's embedded Benefit Analyzer reported an average I/O response time improvement of 94%.

When we ran the script sequence for a second time, a high level of cache utilization, improved the average throughput for the three queries in the script by an impressive 129%. What's more, thanks to a large proportion of the I/O being redirected from disk to memory, our database query model generated very little physical I/O on the logical disk storing the test file.

More importantly, when we finished running our simulated database queries, we also examined the impact of the tests on the overall status of V-locity VM performance. IntelliMemory average caching hit rate rose to 89%. In turn, V- locity VM used this data to project that applications would benefit from a 94% improvement in response time, which closely paralleled our measured throughput rate.

Sequential Streaming I/O

In our second set of tests to capture performance improvement in a sequential I/O environment, we increased the size of our data set by a factor of 100X and changed the test environment to stream

SQL Activity Test Database Table Model (50MB)			
I/O Activity	Without V-locity VM	With V-locity VM	Percent Improvement
4KB reads (25%) 4KB writes (75%) Random Access (80%)	1.85MB per second	4.0MB per second	116%
16KB reads (50%) 16KB writes (50%) Random Access (80%)	15.7MB per second	36.0MB per second	129%
32KB reads (75%) 32KB writes (25%) Random Access (20%)	57.9MB per second	140.0MB per second	142%

video files from a VM running Windows Server 2012. The growing use of customer-facing video data has introduced a new class of video content creation applications – dubbed Non Linear Editing (NLE) – which stream large I/O blocks sequentially.

NLE software applications depend entirely on strict minimum storage performance to prevent frame drop. As a result, NLE storage configurations typically feature both SSD and HDD arrays to ensure that multiple editing streams will scale within a user pool while continuing to meet strict

Cached Video Stream

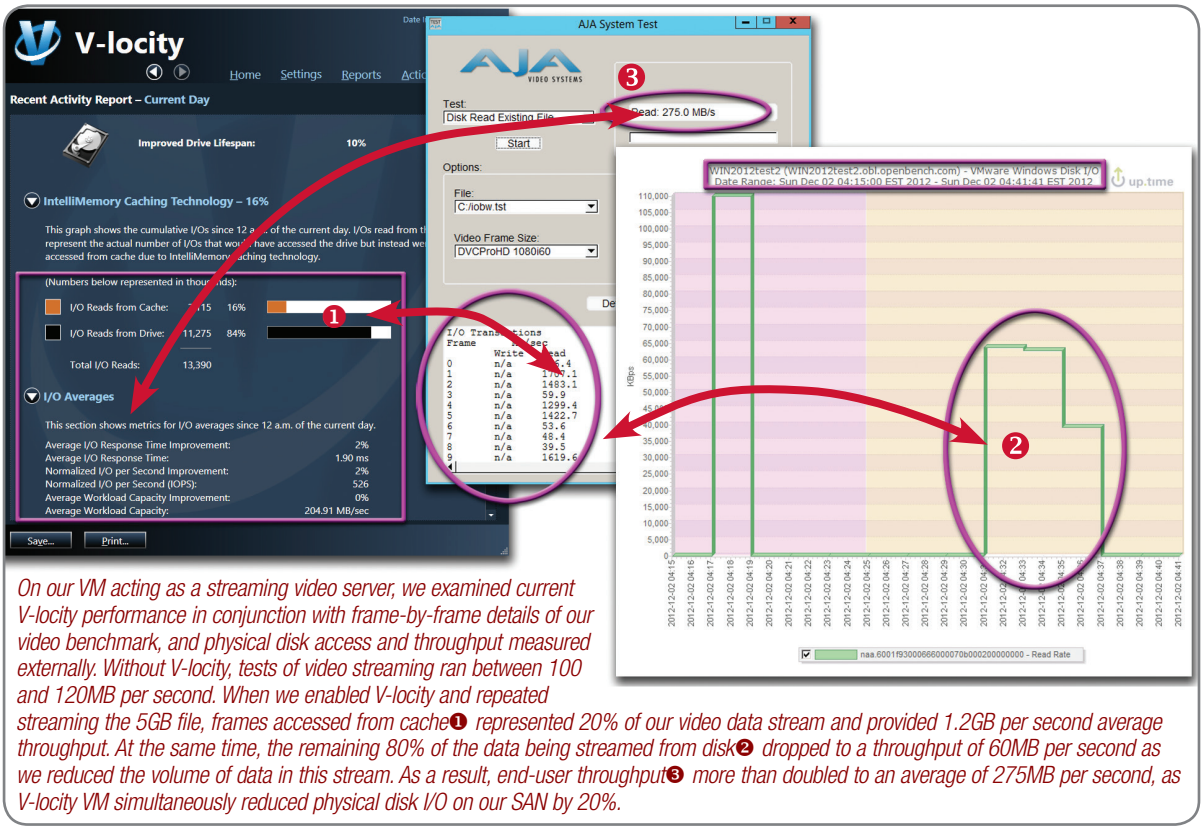

On our VM acting as a streaming video server, we examined current V-locity performance in conjunction with frame-by-frame details of our video benchmark, and physical disk access and throughput measured externally. Without V-locity, tests of video streaming ran between 100 and 120MB per second. When we enabled V-locity and repeated streaming the 5GB file, frames accessed from cache❶ represented 20% of our video data stream and provided 1.2GB per second average throughput. At the same time, the remaining 80% of the data being streamed from disk❷ dropped to a throughput of 60MB per second as we reduced the volume of data in this stream. As a result, end-user throughput❸ more than doubled to an average of 275MB per second, as V-locity VM simultaneously reduced physical disk I/O on our SAN by 20%.

performance and space requirements. In this test, we used a 5GB video file, which we were able to initially stream in standard HD 1080 x 60 frames at 119MB per second. Nonetheless, to effectively support an NLE application, we would need to achieve a 250MB per second throughput rate.

When we repeated our video streaming performance test, there was little likelihood that our 5GB test file would be entirely consumed in the local IntelliMemory cache. Nonetheless, there were enough data frames resident in the IntelliMemory cache to yield an average throughput of 275MB per second, which represented a 131% improvement in I/O throughput with V-locity VM and provided us a 10% margin of error for safely supporting an NLE application on a VM server. What's more, as we were improving end-user performance, we were also reducing SAN I/O traffic from this VM by 19%.

Advanced Storage Feature Compatibility Test	
Test System	Percent Over-Subscribed Disk Blocks
V-locity VM	0%
V-locity 3	78%

I/O Optimization Without Impacting Advanced Storage Features

Intelliwrite and IntelliMemory have been designed to solve I/O bottlenecks at the source – the VM guest OS. One new feature added in V-locity VM was the ability to provide I/O benefit without negatively impacting such advanced storage features as snapshots, replication, data deduplication and thin provisioning. With V-locity VM, since I/O optimization takes place at the server level, I/O is proactively optimized before it reaches the network or storage.

To test this proposition, we set up two server VMs running Windows Server 2012 with thin-provisioned system disks. On one server we installed V-locity VM and on the other we installed V-locity 3, which does not have the same default capabilities. We then created 5,000 new text files on both servers. As writes were optimized for the new files, the VM running V-Locity showed no Over-Subscribed growth in disk blocks automatically provisioned to support the VM's logical disk file. In contrast, the logical disk file on the VM running V-locity 3 grew 78% more than the nominal logical disk file on the VM running V-Locity.

Automated Benchmarking with the Benefit Analyzer

In a typical VI infrastructure, optimized I/O activity will have a primary effect on multiple users of a single VM and secondary effects across multiple VMs located on a common datastore. As a result, improvements brought about by IntelliWrite and IntelliMemory will improve the scaling of users and activity on a VM along with the ability to scale the number of VMs a host can support.

For many sites, the time and expense for profiling and benchmarking the potential benefits they can garner is a serious stumbling block. To provide IT administrators with a quick snapshot of the benefits that will be garnered on a VM, V-locity has an analysis facility which takes V-locity off-line for 24 hours as usage data is collected, then brings V-locity back online and collects data for a before/after comparison. To get a true before/after analysis, the Benefit Analyzer must be run during initial installation.

To test the accuracy of the Benefit Analyzer results, we ran our Iometer tests in parallel using a Windows PowerShell script. The Benefit Analyzer produced similar results within narrow margin of error to our findings after each test. For any organization seeking a proof-of-concept of V-locity VM, the Benefit Analyzer is an invaluable tool to capture performance benchmarks without impact to IT staff and intuitively display results that are presentation-ready.

Bottom Line

By running V-locity VM on all Windows-based VMs, IT administrators can significantly improve application performance, add more user processes to a VM and increase the number of VMs running on a host without adding additional storage hardware, which enables IT to maximize the ROI associated with any VI initiative.

Benefit Analysis Model

The Benefit Analyzer gives a very complete, single-pane view of all the performance metrics necessary to establish a proof of concept in any environment.

Westborough, Mass.-based openBench Labs was founded in 2005 by Dr. Jack Fegreus. openBench Labs is a trusted IT industry source, providing hands-on evaluation and certification of Information Technology products and services. openBench Labs enjoys a unique position in the information technology sector. As the premier independent test lab and provider of third-party validation services, OBL has worked with virtually every major vendor and evaluated the most important products and technologies to appear over the past decade.