

ESG Lab Review

Independent Audit and Overview of ConduSiv's V-locity I/O Reduction Software on 3,450 Virtual Servers in Production Environments


Date: October 2015 **Author:** Mike Leone, ESG Lab Analyst

Abstract: ESG Lab audited data that was collected from 3,450 systems running in production environments that evaluated V-locity I/O reduction software at over 100 sites and measured the achieved performance benefits from virtual servers to storage. The results published in this ESG Lab report shine a light on the I/O inefficiency problems commonly experienced in virtual environments and how V-locity I/O optimization software can help solve them.

The Challenges

Tight budgets and exponential data growth continue to drive organizations to utilize server virtualization technology. By leveraging existing server and storage hardware and improving resource utilization, organizations have seen significant improvements in operational efficiency by consolidating applications and workloads through virtualization. In fact, managing data growth and increasing the use of server virtualization are two of the five responses most-cited by 600+ respondents to an ESG research survey when they were asked about their top IT priorities for 2015 (see Figure 1).¹

Figure 1. Top Ten Most Important IT Priorities for 2015



Source: Enterprise Strategy Group, 2015.

Despite the cost savings and efficiency improvements server virtualization delivers, challenges still remain as resource consolidation can lead to unpredictable performance, especially in highly virtualized environments. With multiple virtualized applications sharing the same underlying storage system, the result is what is referred to as the "I/O blender effect." Though the problem can be addressed with careful planning, over-provisioning storage, and utilizing flash, these "fixes" not only add additional complexity and cost, but they also serve as a Band-Aid approach that does not address the underlying problem.

¹ Source: ESG Research Report, [2015 IT Spending Intentions Survey](#), February 2015.

Challenges in Windows Environments

When virtualization occurs in Microsoft Windows environments, numerous inefficiencies are exacerbated due to the way Microsoft handles I/O from an OS level. Because the Windows file system serves as an abstraction layer on top of the storage, it is unaware of file sizes. Therefore, data gets written at the first available address at the logical disk layer that can fit the file size. This is not a problem the first time a file is written on a fresh OS install. The performance issues begin as the file system ages as files are written, erased, and extended, creating I/O segmentation. In other words, fitting a larger-sized file into a contiguous section becomes impossible. Therefore, the file must then be broken up into smaller chunks, which penalizes performance since each of the pieces requires its own dedicated I/O operation, creating an excessive amount of split I/Os to process a given workload. The more chunks the Windows OS breaks a file down into at the logical disk and subsequent physical media, the longer it takes to write or subsequently access the file. The performance impact is obvious, causing the system to work harder over time and inflating the number of IOPS required to read or write data. And this is just talking about a single Windows install. When factoring in virtualization and consolidating multiple virtual machines or hundreds of virtual desktops all running Windows on the same host, the performance penalty is further exacerbated by randomization from the “I/O blender effect” as the hypervisor mixes and “blends” these disparate I/O streams from multiple VMs before sending them out to storage in a very random I/O pattern.

The Solution: V-locity from Conduktiv Technologies

V-locity is a software solution that reduces I/O by optimizing data placement at the Windows OS level in both physical and virtual environments. Used primarily in virtual environments, significant storage performance improvements are achieved without the need to purchase additional hardware. The software is installed directly on the Windows OS running on a guest machine as a thin file system driver to feed the OS intelligence on file sizes and help choose the best write allocations. Because the software resides at the point of I/O creation (the application), unnecessary I/O can be reduced and, in most cases, completely eliminated.

Optimizing Write I/O with IntelliWrite

The application-aware V-locity software recognizes when the Windows OS is about to break a file into pieces at the logical disk layer, resulting in excessive split I/Os. File size intelligence is fed to Windows to help reorganize the write for a more sequential pattern, leading to the absolute minimum I/O to process any given unit of data. This reduces the I/O requirement for every workload from virtual server to storage, helping systems achieve higher levels of throughput and faster response times. This performance benefit also enables performance improvement when reading that smartly placed data. Because the data is written with a single or minimum I/O, anytime that data is read, it can service the read request with a minimum amount of I/O.

Optimizing Read I/O with IntelliMemory

By leveraging a server-side DRAM caching engine that resides uniquely close to the OS in the overall software stack, V-locity uses a behavioral analytics caching engine to detect and cache active data in available memory, using a minimal amount of capacity to deliver an average 40% reduction in I/O response time based on the ESG Lab analyzed data. By using available memory, no additional cache allocation is required and since the technology leverages a dynamic cache that is automatically throttled based on application needs, memory starvation is a non-issue. This greatly reduces read I/O requests that traditionally must traverse the OS stack to get to the underlying storage and back.

Data Collection: 3,450 VMs Evaluate V-locity on Production Workloads

ESG Lab worked with Conduktiv Technologies to audit and analyze data that was automatically collected from over 100 different sites that evaluated V-locity deployed across 3,450 virtual machines in production environments. This self-auditing benchmark, called the V-locity Benefit Analyzer, analyzes performance improvements before and after the V-locity technology optimizations and highlights key performance metrics like IOPS, throughput, response times, workload comparisons, and more. A view of the V-locity Benefit Analyzer is shown in Figure 3.

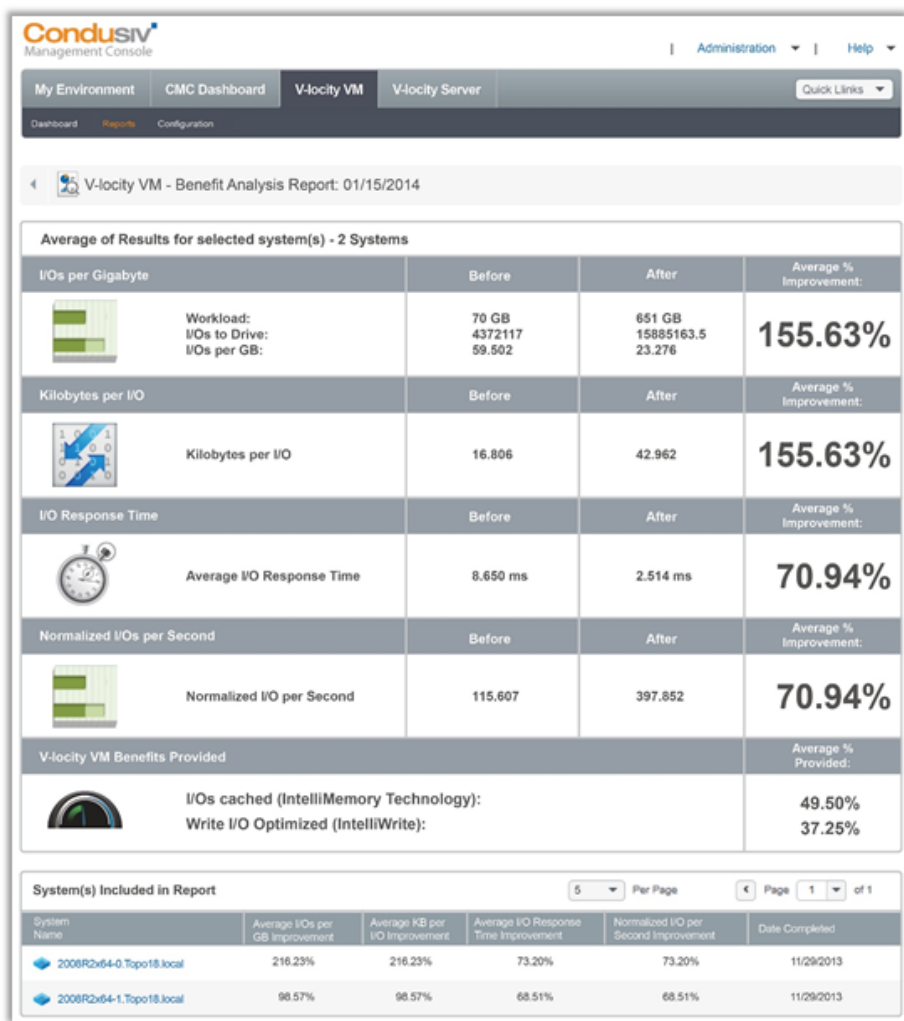
The Benefit Analyzer software contains multiple components that collect telemetry data from Windows VMs before and after V-locity software. Three periods of data are collected:

1. Baseline – The software is installed, but runs in passive mode to strictly monitor workload performance metrics before V-locity optimizations take place.
2. Initial Optimization – The software runs and completes initial optimization.
3. Normal Operation – The software runs in normal mode and collects the “after” performance results, which then get compared with the baseline results before V-locity optimization.

For all data collection to complete for each of the 3,450 analyzed systems nationwide, the average length of time was approximately 28 hours per system. The data was collected and presented via a JSON formatted file. Each file contains three basic types of data:

1. Site and system data – Basic data obtained mostly through standard Windows calls. This includes information like the version of software being used to collect the data, customer name, operating system, etc.
2. Storage I/O data – Data collected through telemetry and cache drivers by monitoring each I/O sent directly to the Windows storage layer (disk.sys) in the kernel. High-performance Windows counters are used to provide granular timing information, all of which are normalized to 1 microsecond increments.
3. Fragmentation prevention and mitigation component data – Reported data from IntelliWrite about the number of extra fragments prevented due to the V-locity software. The fragmentation mitigation engines provide the actual number of eliminated fragments. The V-locity telemetry driver detects files that cause storage performance issues due to file fragmentation.

Figure 2. V-locity’s Benefit Analyzer



ESG Lab Analysis of Performance Results

Reduced Read I/O to Storage – ESG Lab calculated that 55% of systems saw a reduction of 50% in the number of read I/Os serviced by the underlying storage. Even more, 27% of systems saw a 90% or more reduction in read I/Os. These impressive I/O reductions could be attributed to the caching advantages provided by V-locity.

Reduced Write I/O to Storage – By optimizing writes to be written in a more contiguous fashion, the size of an I/O consistently increases. In other words, instead of writing four 4Kb blocks of a 16Kb file, V-locity enables the system to write a single 16Kb write, requiring a single I/O operation. As a result of the I/O density increase, ESG Lab witnessed a 33% reduction in write I/Os across 27% of the systems. 14% of systems experienced a 50% or greater reduction in write I/O from VM to storage.

Decreased I/O Response Time – With an average available DRAM size of 3GB across all 3,450 systems, ESG Lab calculated the total time required to process all requests for each system and concluded that on average, systems achieved a 40% reduction in response time.

Increased Throughput – ESG Lab witnessed throughput performance improvement of 50% or more for 43% of systems. Further, 29% of systems experienced a 100% increase in throughput and as much as a 300% increased level of throughput for 8% of audited systems.

Increased IOPS from DRAM – Though the overall goal of the solution is to lower IOPS and improve throughput to the underlying storage, in some cases, because the working set was consolidated and serviced primarily out of DRAM, the number of measured IOPS dramatically increased. This means that the application was able to service requests faster. In fact, 25% of systems saw IOPS increase by 50%, and a small group of 25 systems achieved a 1,000% IOPS improvement. The same can be said for throughput. ESG Lab witnessed throughput performance improvements of 50% for 43% of systems, and as much as 300% increased levels of throughput for 8% of systems.

Why This Matters

Performance is a critical concern when a mix of mission-critical applications shares the same host and therefore, the same underlying storage. This is especially troublesome in Windows environments, where over time, write inefficiencies lead to a loss of throughput and an inflation in IOPS. I/Os are increasingly fractured and randomized, leading to performance penalties that directly impact every application on the same shared storage infrastructure. Over-provisioning storage can help, but serves as a short-term fix to meet peak performance requirements, while adding complexity and increasing costs.

ESG Lab validated that V-locity software optimizes read and write I/Os for physical and virtual Windows systems that consistently demand high levels of storage performance. The IntelliWrite component worked to eliminate small, fractured I/Os with more contiguous writes, leading to reduced randomization and noise from the “I/O blender effect.” IntelliMemory allowed for a high percentage of audited systems to leverage available DRAM as Tier 0 storage for their caching strategy. This combination led to a significant reduction in both read and write I/Os to the underlying storage, an increase in IOPS and throughput by more efficiently utilizing DRAM, and overall application performance improvements by reducing latency by an average of 40%.

The Bigger Truth

Microsoft Windows serves as a primary operating system running a majority of the most important enterprise IT applications in the world. Whether it's in a physical system or in a virtualized environment with hundreds of VMs, Windows applications traditionally serve as the lifeline of the business. The problem is that as the Windows file system ages, those important applications that the business relies so heavily on get impacted by write inefficiencies from the OS. Poor performance, application starvation, and lost productivity start wreaking havoc on the whole environment. A simple database query from one application turns into an unpredictable burst of random I/O activity, impacting every other system that shares the underlying storage and creates the infamous "noisy neighbor" issue for every other VM on the same host.

V-locity from Conduktiv Technologies serves as a lightweight software solution to non-disruptively optimize I/O profiles to be more storage friendly by ensuring that files are written and read contiguously. This has an immediate and long-term impact on file system performance by improving resource utilization, storage efficiency, and application performance. This is especially important in virtualized environments where multiple systems share the same host and underlying storage. Further, the software intelligently leverages DRAM to service the smallest and most random I/O that has the greatest impact on performance to free up I/O bandwidth for other systems or applications.

The V-locity software provides the most benefit to the heaviest workloads in an organization. These are the systems that can bring the entire infrastructure to its knees, especially if the application demands consistently high levels of storage performance to service an unfriendly I/O profile whose characteristics are much smaller, more fractured, and more random than it needs to be, which effectively starves every other system that shares the underlying storage.

The last thing IT organizations want to do is spend more money on additional hardware to *hopefully* improve their storage performance and operational efficiency for the immediate future. It takes time to pick the right technology, vet it, deploy it into production, and then start seeing the benefits. A free evaluation of V-locity from Conduktiv Technologies can prove that an easy-to-deploy, non-disruptive software solution can drastically improve application performance without the added cost or complexity of more hardware.

The goal of ESG Lab reports is to educate IT professionals about data center technology products for companies of all types and sizes. ESG Lab reports are not meant to replace the evaluation process that should be conducted before making purchasing decisions, but rather to provide insight into these emerging technologies. Our objective is to go over some of the more valuable feature/functions of products, show how they can be used to solve real customer problems and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab report was sponsored by Conduktiv Technologies.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of The Enterprise Strategy Group, Inc., is in violation of U.S. copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at 508.482.0188.