**September 8, 2014**
**Jeremy Li**

## V-locity Overview with Questions and Answers

Condusiv Technologies (**http://www.condusiv.com/**), formerly known as Diskeeper, was founded in 1981 and once licensed Microsoft its defragmentation utility as a light Diskeeper in every Windows Server Operating System. Today, it goes beyond its once a flagship product of defragmentation utility and offers additional product called V-locity to target many Windows customers to improve Windows application performance in a Windows Server OS virtualized environment by using its patented technologies to accelerate IOPS performance, as well as reducing the I/O blender issue. Condusiv states if V-locity doesn't boost performance of its customer system by 25 percent, the software is free.

Gartner named Condusiv Technologies as a Cool Vendor in Storage Technologies, 2014 due to its superior patented technology to reduce the I/O blender issue. Thus, many customers can avoid purchasing new hardware such as a newer storage, server and etc., while seeing the performance gain by 50% on average.

Below are the Q&As after the V-locity overview presentation with follow-up E-mails.

**Q1: Do you think the I/O blender issue solved by Gridstore (see Figure 2 for reference) will not be resolved by installing V-locity software, as illustrated in Figure 1? Let's only talk about the I/O blender issue because of server-side virtualization?**

**Figure 1**

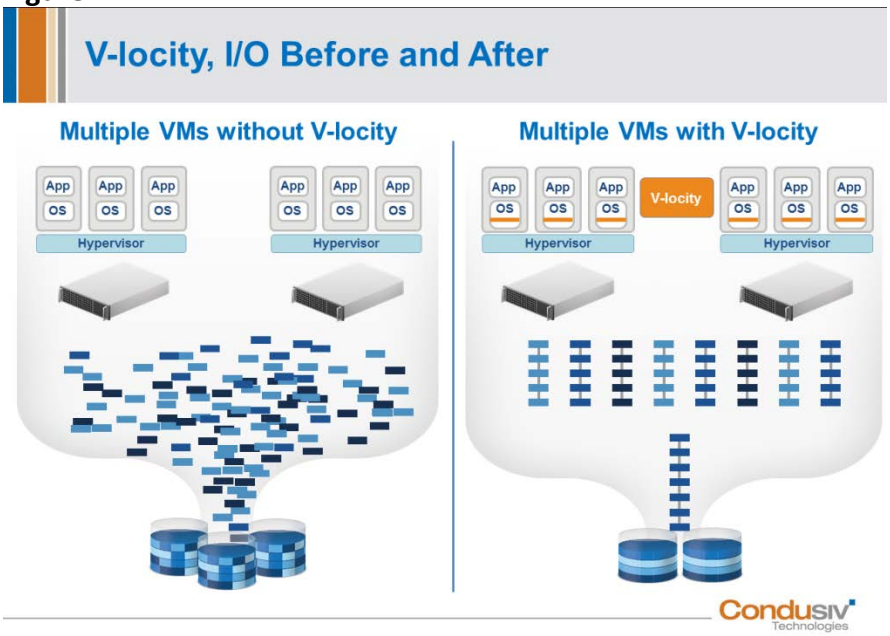


**Courtesy of Condusiv Technologies**

A1: The I/O blender issue is a complex subject. I do believe the I/O blender effect can/will be solved by V-locity. The only reason why anyone cares about the issue is because of the overwhelming amount of random yet inner related I/O data requests that saturates the storages ability to response in a time frame acceptable by users.

With more random I/Os it simply takes longer to complete the amount of data transferred by the application running on that guest system. Each Guest VM running Windows has its own abstract layer of their corresponding VHD or VMDK file that is managed by the Hypervisor & Host. When an application running on a Windows Guest VM requests to write data, it must first pass the request to Windows and the NTFS file system. NTFS sees a partition which is reference by metadata structures such as the $MFT and $BitMap. When a write request is being processed, NTFS does not scan the $BitMap looking for the best possible fit. It will allocate space based upon the next available free space segment from where it previously had allocated data. So for example, if the next several sections of free space are chopped up into many 4kb segments and you are attempting to write a 1MB size file there will be many small individual random write I/Os. But wait it gets worse… Multiple applications and users are accessing data at the same time. The behavior of applications is to process data, write, process data, write some more information, etc… As a result there are many I/O threads writing data within NTFS and these writes and subsequent deletes causes further random write allocations. This is then further compounded as you add additional Guest systems in a virtualized environment reeking havoc across the hypervisor, host, network, and storage systems.

V-locity solves the problem by providing the intelligence to NTFS of each Windows VM using behavioral analytics to prevent these random I/Os from occurring in the first place. V-locity will help NTFS pre-allocate a consecutive range of logical clusters for the initial file write such that the requests will be larger and more sequential in nature. When the file is closed any over allocation of space is truncated to the last used cluster written by the application. Subsequent reading of the file is also benefited because all the read I/Os are also done in a more sequential manor.

In addition, V-locity reduces the read I/O traffic by intelligently and dynamically caching the frequently accessed read data blocks. Think about it… What's faster than the fastest SSD/Flash Memory/Caching Tiered storage device???? That's right, not having to send the I/O to the storage device in the first place. A memory to memory data transfer is many times faster than having to go storage. This in turn lessens the impact of I/O traffic and solves the problem of the I/O blender effect by making Windows and the NTFS I/O request as efficient as possible right from the beginning of the I/O traffic stream.

**Q2: Do you have any 3PAR customers who also use V-locity to address the I/O blender problem**?

**I am interesting in seeing how much the performance has been increased under the 3PAR architecture. I am sure V-locity can improve the performance because 3PAR does not address the I/O blender issue in a VM environment, while GridStore is able to address the I/O blender problem at the source, as shown in Figure 2.**

**Figure 2**

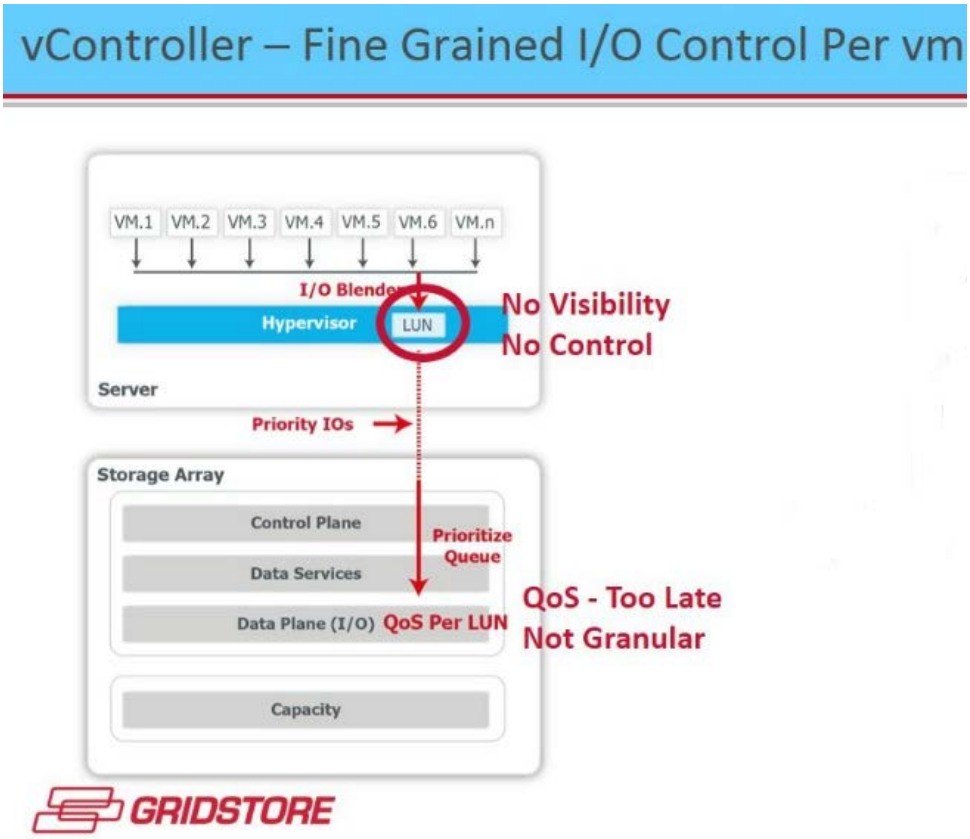


**Courtesy of GRIDSTORE**

A2: We generally don't track customer success and performance gains based on their hardware, but most customers report a 50%-300% benefit in areas such as reduced I/O traffic, response time, and workload throughput. V-locity solves the I/O blender problem not by adding additional hardware, but by preventing and eliminating the wasteful NTFS "Split I/O" and caching the Read I/O requests using "available" memory within each Guest VM such that there is fewer I/O request sent from each VM Guest to the Hypervisor->Host->Network->Storage. Every component downstream, including Gridstore receives a positive benefit from V-locity.

**Q3: Does V-locity support Windows Server 2012 R2? If so, what's an average performance increase from your customers' feedback if all Guest OSs are based on Windows Server 2012 R2 in hosts?**

**Let's assume that my environment is based on Windows Server 2012 R2 with SMB 3.0 enabled from clients. SMB 3.0 has improved the NTFS efficiency dramatically by combining multiple write requests into a single larger request within OS, thus, dramatically reducing the multiple write requests based on SMB 2.0 (Windows OS' splitting files into multiple write requests)**

A3: V-locity does fully support Windows Server 2012 R2 with SMB 3.0. Each customer environment is different and is more affected by the nature of their applications read/write activity and availability of available memory than the combination of hardware and flavor of Windows & SMB enhancements. The approach I take isn't what I have done for others, but what real improvements can I make in your real production environment. That's where the "rubber meets the road" and would be the true measurement of success in using V-locity.

**Q4: Since HP 3PAR (v3.1.2) supports Microsoft ODX that equates Windows Server 2012 + Hyper-V offloaded copy (bypassing Hypervisor or Hyper-V), I am interesting in learning what's the percentage of performance will be increased if both ODX and SMB 3.0 are turned on?**

A4: While I would never dismiss hardware claims as they do have a real purpose in benchmarking isolated and control workloads to document a given benefit of one platform over another. The challenge is with the mix of applications in the real world under production workloads. I am confident that I can give you at lease a minimum of 25% improvement (reduction) in I/O activity. This means you will be able to process the same amount of data faster and/or be able to handle more workflow with no decrease in application responsiveness. Are you ready to take the V-locity challenge?

Jeremy's comments:

I also read the article and found a significant improvement after installing the v-locity software, although your conservative estimate is at 25%. Below you will find a few excerpts from the article:

"Also, Condusiv says if V-locity doesn't boost performance of your system by 25 percent, the software is free. The new release optimizes the physical server before data is routed to NAS or SAN resources.

Like the version for VMs, the new V-locity release is aimed at Microsoft workloads such as Exchange Server and SQL Server and others that require extensive use of I/O, according to company officials."

Since the County of Los Angeles is moving to Office 365, I think V-locity might be considered to improve the performance on an QL application.

**Q5: Per one storage expert, a third party SW (e.g., V-locity) might bring a detrimental effect because it lacks an end-to-end control, meaning V-locity agent is not installed at the storage layer to communicate with the agent at the guest level. What do you think?**
**By the way, Nimble Storage with CASL Architecture delivers a sequential write with Read-Cache acceleration. However, it might be too late since V-locity takes care of the issue before I/O activities arrives at the storage layer.**

A5: The operative word here is "might". This implies that no firsthand testing or measurement has been done by this storage expert as a statement of fact. What I can say for a fact is that V-locity was designed to enhance the I/O traffic flow by reducing the unnecessary and inefficient I/Os within the NTFS file system layer. This is accomplished with a few Microsoft signed light weight and small file filter drivers combined with a system resource monitoring and behavior analytic architecture which make it all but impossible to notice any negative overhead associated to V-locity. V-locity was also purposely designed as a complete software solution and solely contained within the Windows Guest VM. It is not an Agent nor is it an Appliance and has no dependencies to acquire storage layer information and pass that data back and forth to evaluate or process. V-locity is therefore complementary to the storage layer such that if fewer or more efficient I/Os are sent then the storage systems can focus more on what they do best.

**Q6: What's the architecture difference between V-locity and Atlantis architecture?**

A6: As previously stated, V-locity is installed on the Windows Guest VM and is a software only approach. No additional hardware is required. Atlantis is installed between the hypervisor and the local server RAM. At last look, Atlantis achieves performance gains primarily by write I/O optimization using deduplication and compression. It's best used on top end systems as it quite CPU intensive and they recommend systems with 256GB of RAM (192GB is required. It is a solid approach, but Atlantis does nothing to fix the problem of inefficient I/Os bombarding the host from each of the Guest VMs.

**Q7: May you kindly tell me whether Microsoft PSS will support any OS issue if a V-locity SW is installed on any Guest OS, as shown in Figure 3?**

**Figure 3**

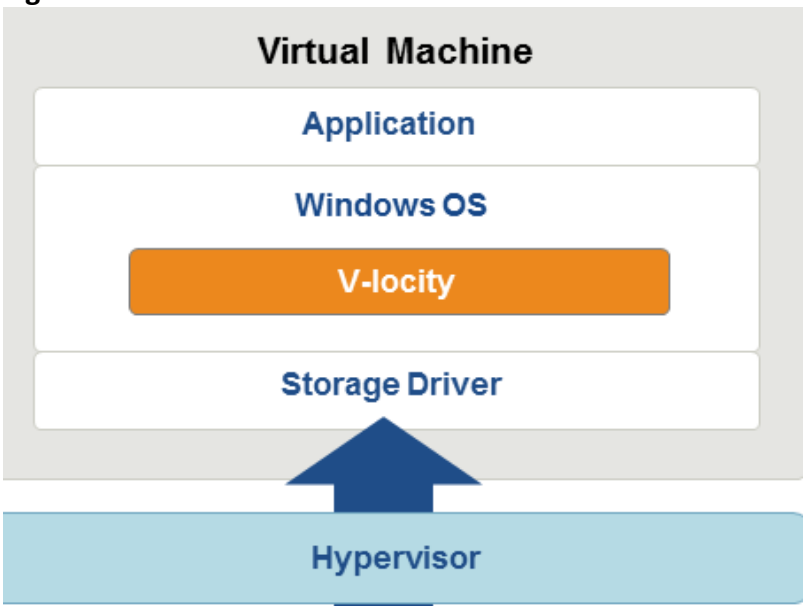


**Courtesy of Condusiv Technologies**

**According to your answer for Q5 as quote "a few Microsoft signed light weight and small file filter drivers", Microsoft should continue to support its OS, even though the V-locity SW has a few signed light weight and small file filter drivers, meaning approved by Microsoft.**

**Note: For example, Microsoft PSS will not support any OS issues arising from WAN Optimization Controller (WOC) in Office 365.**

A7: Yes, Microsoft will continue to provide support. Of course, you'll find a few individuals who are not fully up to speed on what Microsoft has approved, but the info below should help set the record straight. :-)

- All of V-Locity drivers are certified and signed by Microsoft (See Logo report listed below)

- Link below, you can see us listed in the Windows Server Catalog

http://www.windowsservercatalog.com/results.aspx?text=Condusiv&=Go&bCatID=1282&avc=10&ava=0&OR=5&chtext=&cstext=&csttext=&chbtext=

- We regularly attend the Microsoft Filter Driver Plugfest events, where we run Microsoft stress tests against our products and well as run compatibility tests against other Microsoft and 3rd party vendors, which include a wide variety of applications

- If necessary, the customer can run their own compatibility testing as part of the PoC process.

## Windows hardware certification report: Approved

| | |
|---|---|
| Submission ID: | 1632942 |
| Submission date: | 1/22/2014 |
| Hardware certification completion date: | 1/23/2014 |
| Company: | Condusiv Technologies |
| Product name: | V-locity |
| Category: | Device |
| Product type: | File SystemOther Device |
| Qualification level: | Signature Only - Microsoft Windows Server 2003 family, x64<br>Signature Only - Microsoft Windows XP family, x64<br>Signature Only - Microsoft Windows Vista family, x64<br>Signature Only - Windows Server 2008 family, x64<br>Logo - Device - Compatible with Windows 7 x64<br>Logo - Certified for Microsoft Windows Server 2008 Release 2 family, x64<br>Certified for Microsoft Windows 8 Client family, x64<br>Certified for Microsoft Windows 8.1 Client family, x64<br>Certified for Microsoft Windows Server 2012 R2, x64<br>Signature Only - Microsoft Windows XP family, x86<br>Signature Only - Microsoft Windows Server 2003 family, x86<br>Signature Only - Microsoft Windows Vista family, x86<br>Signature Only - Windows Server 2008 family, x86<br>Logo - Device - Compatible with Windows 7<br>Certified for Microsoft Windows 8 Client family, x86<br>Certified for Microsoft Windows 8.1 Client family, x86 |
| Marketing names: | N/A |

**Q8 If SSD-based products are used (e.g., VNX-F array), can those extra I/O issues be mitigated via Flash?**

**Scenario 1: Just install one PCIe card into a host, all VMs on the host might be taken care of without installing V-locity software one-by-one on each VM.**

**Scenario 2: Installing a shared SSD appliance (e.g., IBM FlashSystem 840) can also mitigate the extra I/O requirement due to virtualization, in addition to address a burst mode of I/O scenario, as well as shared all flash pool efficiently.**

A8: Adding an SSD into the mix will help mitigate some I/O performance problems, but it does not address the root cause. You will never achieve the full potential of SSD storage if the file systems of the Guest systems are performing unnecessary I/O requests. Think of our solution (V-locity) as a means of reducing the amount of I/O activity for a given workload. It

really is all about I/O avoidance rather than individual I/O speed. Even with an SSD there is a measurable time an I/O takes to complete.

Remember… The greatest performance gain is not having to do the I/O to the storage device in the first place. This reminds me of what Albert Einstein once said, "A clever person solves a problem, but a wise person avoids it." V-locity optimizes and reduces the workload being generated at the source of where the I/O traffic originates within the Guest. If you do add SSD or use SAN/NAS storage V-locity will make that storage solution more efficient and perform faster simply because we eliminate the wasted and unnecessary I/O traffic allowing the storage hardware/offload engines/appliances to focus on the important I/O tasks at hand.

**Q9: May you tell me whether Microsoft SMB 3.0 has improved the NTFS performance, meaning has addressed the NTFS inefficiency over SMB 2.0 described, as well as illustrated in Figure 4?**

**Figure 4**

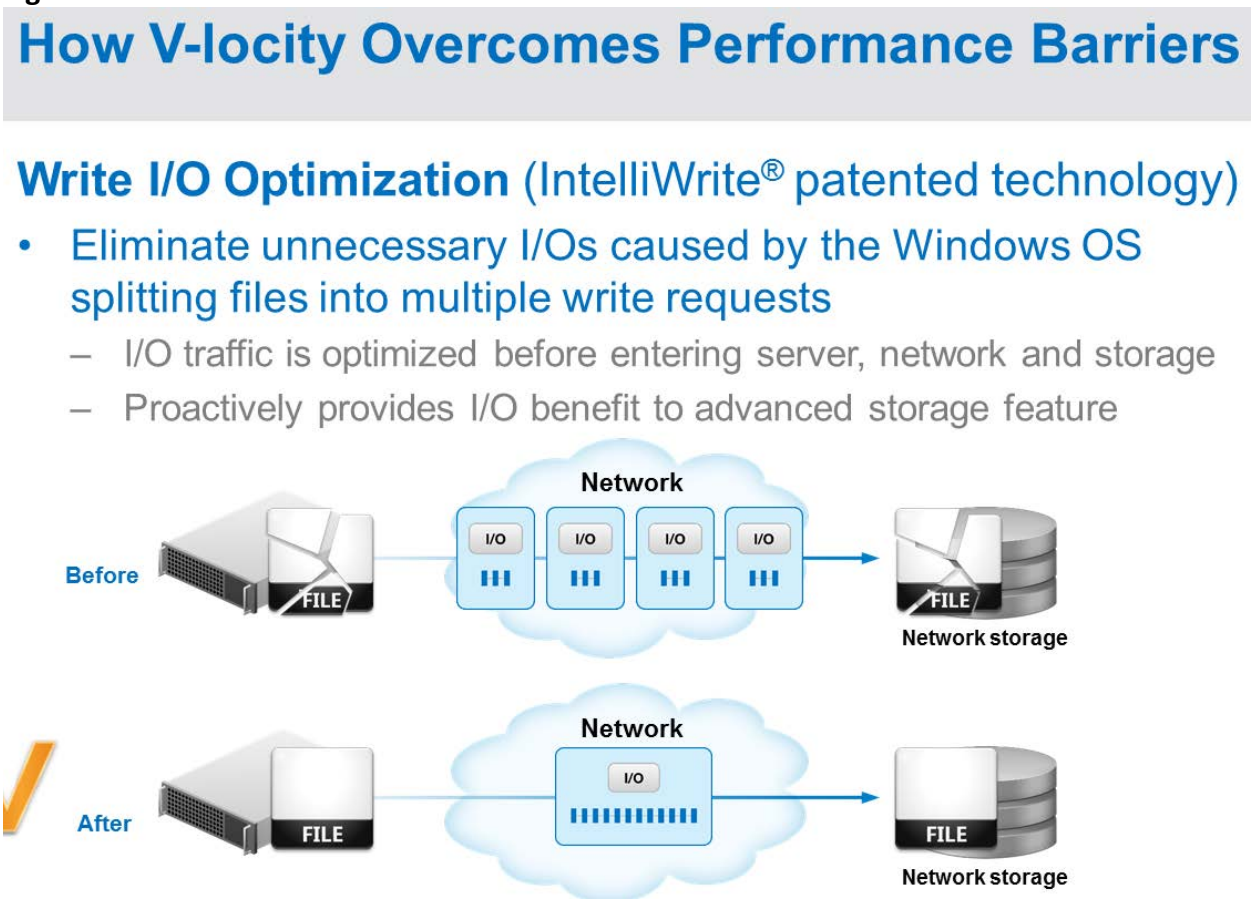


**Courtesy of Condusiv Technologies**

**Note: Jose Barreto is a program manager on Microsoft storage team, and wrote a blog summarizing the differences in many versions of SMB capabilities:**

http://blogs.technet.com/b/josebda/archive/2013/10/02/windows-server-2012-r2-which-version-of-the-smb-protocol-smb-1-0-smb-2-0-smb-2-1-smb-3-0-or-smb-3-02-you-are-using.aspx

A9: Yes, Microsoft SMB 3.0 has made some improvements to file read/write performance within NTFS, but does not solve the problem as I described in my answer to Question 1 above. V-locity is the only complete solution and comes with a Benefit Analysis mechanism to capture I/O performance measurements of your real world production environments "before" and "after" V-locity has engaged on your system.

## Recommendation

First, whenever possible, a latest version of SMB should be used in order to maximize the scalability, the best performance due to maximum transmission unit (MTU) increased (version 2.1), as well as SMB Multichannel support for higher availability, in addition to enhanced security in a Windows environment.

Keep in mind that Windows Server 2012 Hyper-V and Windows Server 2012 R2 Hyper-V only support SMB 3.0 for remote file storage. This is due mainly to the availability features (SMB Transparent Failover, SMB Witness and SMB Multichannel), which did not exist in previous versions of SMB. The additional scalability and performance is also very welcome in this virtualization scenario. The Hyper-V Best Practices Analyzer (BPA) will warn you if an older version is detected.

Second, always perform a proof of concept (POC) process in order to determine how much V-locity can improve a Windows Application performance.
Third, calculate a break-even point to warrant for installing the V-locity to improve a Windows Application performance.

## Challenges and Caution

Condusiv Technologies faces the fact that many storage vendors' newer products such as Peripheral Component Interconnect Express (PCIe) flash adapter stored on a host, as well as a hybrid storage array with flash first design (EMC VNX 2) or an all-flash-shared appliance (e.g., PureStorage), can boost IOPS dramatically, even though the help for more IOPS is too late, as illustrated in Figure 2 above.

The current V-locity advantage referenced in the conclusion section below might be significantly reduced if a future Windows OS will address the Windows split I/Os effect, as illustrated in Figure 1.

## Conclusion

The V-locity can eliminate unnecessary I/Os (Windows split I/Os effect) caused by the Windows OS splitting files into multiple write requests by putting many small writes into a single big write sequentially, similar to Nimble Storage's CASL Architecture (Cache Acceleration for read-only and Sequential Layout for writes), as illustrated in Figure 4.

It can optimize reads and writes at the source, and resolve I/O blender issue, as illustrated in Figure 1, in addition to caching hot data in available server memory.

## Recommended Reading


1. SunCoke Energy Solves Months of Oracle Performance Troubleshooting with V-locity VM
2. Gartner Cool Vendors in Storage Technologies, 2014 (**15 April 2014** ID:G00262188)
3. V-locity 4 Review: Optimize Your Virtual Guest Machines
4. Magic Quadrant for Solid-State Arrays (28 August 2014 ID:G00260420)


## Acknowledgement


Thanks Alex Edwards, MAJOR ACCOUNT EXECUTIVE – WESTERN REGION and Howard Butler, SENIOR DIRECTOR FIELD ENGINEERING for providing me a live presentation and a follow-up Q&A session via E-mail.

Thanks Condusiv Technologies for allowing me to use one graphics to illustrate How V-locity overcomes performance barriers.