

Solution Sheet: V-locity and VDI

Increasing VDI Performance and Efficiency with V-locity®

Introduction

Virtual Desktop Infrastructure (VDI) is compelling. A virtual desktop provided to a user just like any application promises an easier way to manage endpoints, to reduce management costs per endpoint, and to reduce infrastructure costs overall. In the real world, VDI is still in exploration mode: organizations are assessing it, some have adopted it, but many are wondering about the real costs associated with it.

VDI: The Myth and the Reality

There's been no shortage of buzz around VDI, including use cases and predictions that turn out to be the opposite from theory to reality. One is the focus on persistent VDI, with a common use case being remote software developer workstations. But in reality, persistent and non-persistent (stateless), or mixed VDI environments, are dominating the landscape: consider call centers, hospitals, airports—these are all cases where users do not require a customized workstation, but log in to thin clients that share the same applications, services, and functions as other virtual desktops in the environment.

Another myth has been that VDI implementations would be mostly read-intensive, probably based on testing in controlled environments with fixed test scripts. But in the real world of mixed VDI environments, this is not the case. Most VDI deployments exhibit largely random I/O requests that easily reach read/write ratios closer to 50/50 to 40/60 and even as much as 20/80% during normal working operations. Because write I/O is so much heavier on storage at the backend, it's critical to understand read/write ratio and how it impacts IOPS on the supporting storage subsystem.

Another myth is that VDI has to be expensive. Of course it may be, but does it have to be? One thing we know for sure, IOPS can be the “silent-killer” of a successful VDI implementation. So if you're undertaking a Proof of Concept (POC), look for a good online calculator and run some actual numbers.

First, get a sample of a range of users and determine their aggregate peak I/O. Use that number to determine the amount of

I/O per virtual desktop to calculate for future growth. Let's say you get 25 IOPS per virtual desktop. Whether those 25 IOPS are mostly read or mostly write will determine what you need from the storage device. In any deployment—small, midsize, or large—that difference has a huge financial impact. I/Os are not created equal, and failing to understand IOPS in a VDI environment will significantly impact storage and infrastructure costs.

Daunting? This is why VDI is still primarily in exploration mode. Organizations wonder, are we ready for this? How expensive is VDI in the long run? Is the ease of management worth the cost of an infrastructure that can handle up to 80% write activity and still meet SLAs with business units?

When you've compiled enough data for a dissertation and you've done exhaustive hypothetical calculating, consider a different approach—one that has nothing to do with spindles, Flash, SSDs, and more hardware. It's about reducing the amount of I/O required for any given file in the first place.

V-locity: Debunking the Myths and Changing the Economics of VDI

In virtual environments, I/O streams from multiple VMs (virtual machines) is funneled down to storage. In VDI, this blender effect of random I/O is even more intense, creating a chaos of write I/O from VMs and clients that overwhelms the host, network, and storage. The departure from the traditional single application on a single physical server has had a profound impact on the underlying layers that service application demand. This impact is becoming better understood as it relates to the surplus of

V-locity at a glance:

- Non-intrusive software-only solution to improve VDI performance by 50% or more
- I/O optimization intelligence reduces amount of I/O required for any given workload by 50% or more
- Eliminates sluggish performance without adding new hardware
- Provides an ideal user experience regardless of application or thin or fat clients

unnecessary I/O—perhaps the biggest performance barrier in most environments, particularly in VDI.

V-locity[®] is I/O optimization software that accelerates application performance by 50% or more with no additional hardware. Ideal for read and write I/O-intensive VDI deployments that consist of persistent and non-persistent machines, and zero, thin, or fat clients, V-locity is installed at the VM or guest OS layer, as a thin, light-weight filter driver to optimize reads and writes as close to the application as possible, where I/O operations originate.

For write I/O, V-locity's IntelliWrite[®] eliminates the performance penalty associated with random I/O patterns by writing files sequentially as a single contiguous I/O, which also speeds up subsequent reads. For read I/O, V-locity's IntelliMemory[™] acts as a behavior analytics engine that determines what blocks should be cached in available server resources and how long they should be there, based on the application itself, frequency, and time. IntelliMemory reduces read I/O because it doesn't predict which blocks of data you need—it knows exactly which blocks of data you need.

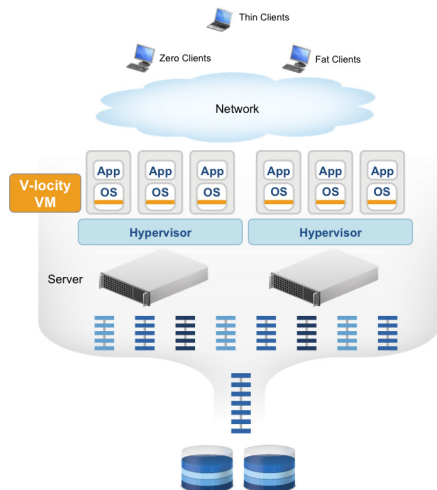


Figure 1: With V-locity, only minimal I/O remains, and I/O behavior is optimized to enable 50% faster applications

IntelliWrite (Write I/O Optimization Technology)

Whereas the Windows file system typically breaks a file into many pieces before write, with each piece requiring its own input/output operation, V-locity reorganizes this random pattern of split I/Os and aggregates writes to behave sequentially as a single, contiguous I/O, allowing organizations to achieve far greater throughput with fewer I/O operations required for any given workload.

In virtual and VDI environments, the problem of heavy write I/O is compounded as multiple VMs and pools of virtual desktops share the same storage resources, resulting in highly random I/O

behavior. Without IntelliWrite, the generation of unnecessary I/Os not only slows VMs on the same host, it also impacts all virtual desktops and applications shared by users.

IntelliMemory (Write I/O Optimization Technology)

V-locity's IntelliMemory is a server-side caching engine unique in where it resides in the architecture. By living at the VM OS layer, alongside the applications that generate read requests, IntelliMemory leverages a behavior analytics engine to learn application behavior (providing non-persistent systems run for at least one hour) and stores the metadata in logs.

Understanding application behavior ensures the most effective caching method possible—reducing I/O traffic since frequently accessed data can now be cached within available server memory. This increases IOPS performance, reduces latency, and lowers overhead on shared storage by reducing physical read I/O requests on storage devices. IntelliMemory's read cache is rebuilt for each session; so, the longer the session, the better the performance. And IntelliMemory does not cause resource contention on VMs because it throttles usage dynamically if an application needs more memory.

Dramatic Reduction in I/O = Dramatic Reduction in VDI Costs

The read/write I/O ratio in a typical VDI deployment can be as high as 20/80%, which causes bottlenecks, latency, and all-around poor user experience. The answer to the problem has traditionally been hardware: more storage, faster HDDs (or more HDDs to distribute I/O), adding Flash, adding SSDs, or maybe SSDs with caching software.

But this approach is expensive and increases the footprint of an already complex data center. It's an issue that makes organizations ponder the economics and practicality of implementing VDI.

V-locity is a non-intrusive software approach, which significantly reduces the total cost of VDI. With V-locity, organizations get maximum VDI efficiency and performance without the CapEx of adding new hardware, or the associated OpEx to deploy and manage it.

By intelligently reducing the amount of I/O traffic in a VDI environment, organizations can lower the cost of a VDI project so it is accessible and practical for all organizations, regardless of the type of VMs or applications users are accessing.

V-locity is available for free evaluation. With V-locity's Benefit Analyzer, you can measure and manage I/O performance using real-world workloads—not predictions—and plan for less hardware.

Add V-locity to your plan for a high-performing and much more efficient VDI deployment—no additional hardware required.

Learn more at www.Conduktiv.com