WHITE PAPER

# Virsto for VDI, Hyper-V Edition

## Changing the Storage Economics of VDI in Hyper-V Environments

## Introduction

Virtual desktop technology promises to streamline administrative operations for company desktops while at the same time providing a secure and better protected solution overall to keep workers productive. But storage has presented some significant challenges to successful virtual desktop infrastructure (VDI) deployments in both the performance and cost areas that have in many cases delayed or derailed projects.

Virsto offers a pure software-based virtual storage solution specifically designed for VDI environments that changes the economics of VDI deployment significantly without having to buy new storage hardware. This solution cuts desktop provisioning and deployment times by hours or days, depending on the size of your deployment, increases the desktop density that any given storage configuration can support by at least 2x, and ensures higher, more predictable storage performance from your existing storage. **With Virsto running in Hyper-V, you will be able to meet your performance requirements with 50% less storage of your choice than if you are running Hyper-V without us, and that translates directly to significant cost savings for you.**

This white paper provides a discussion of the storage challenges in VDI environments and explains how **Virsto for VDI, Hyper-V Edition**, addresses these problems in Hyper-V R2 environments more cost effectively than any other solution on the market. This is a bold claim, but once you see what we do and how we do it, you can decide for yourself.

## Why Is Storage in VDI Environments So Expensive?

VDI environments present some unique challenges, so if you are trying to size and budget your storage configurations based on past experience with physical desktops or even with virtual servers, you may have a few surprises in store.

*The Type of Storage You Use*

First, the storage technology in your physical desktops today is probably IDE. A VDI project will consolidate these desktops onto one or more Hyper-V Hosts in your data center where you will likely be running enterprise-class storage. You can buy 1TB of desktop-class IDE storage at computer retail outlets for a little more than $100, but you'll be paying at least 20x that much on a $/GB basis for even the cheapest enterprise-class storage you will be using in your centralized, data center environment.

*The Amount of Storage You Need*

Second, when you create your first pilot configurations you may notice that your storage seems to be under-performing from an IOPS point of view. You'll be more likely to notice this if you're working with a disk-bound storage configuration in your pilot, and we'd like to suggest that you work to find that disk-bound threshold so that you can more accurately gauge what will be needed to meet your IOPS requirements as you roll your VDI project out. If your pilot is small, don't make the mistake of using an over-built storage configuration that you don't take into account as you roll the deployment out to more desktops, as this will hide any storage performance problems you might run into as you move to production, leading to unexpectedly low performance. An over-built storage configuration wastes

money, and when you roll out your VDI deployment you'll probably want your storage configurations sized so that they provide the performance you need most cost-effectively.

Why does storage seem to under-perform in VDI environments? It's because of a phenomenon called the "VM I/O blender". For a physical server deployed according to the legacy client/server model, you had a dedicated application running on a dedicated server, often with its own dedicated storage. Given that there was one I/O stream coming out of that server, there were optimizations that could be performed to provide pretty good storage performance. This is very different from a VDI environment where you may have 50 – 70 virtual desktops, each running their own independent I/O workload, on a single physical host. VDI environments tend to be particularly write-intensive. The I/O stream coming from that host is significantly more random and significantly more write-intensive than it ever was on most dedicated application servers.

Spinning disks were built to accommodate both reads and writes, but it is well known that they perform at their worst on very random, very write-intensive workloads. That is because the rotational latencies and seek times begin to increase with more random writes, eventually coming to dominate the entire data transfer time for any given write. This impact is not trivial. In fact, at the macro level in virtual computing environments – either server or desktop - it results in storage that appears to under-perform by as much as 30% - 50%.

This has two impacts. First, it slows the application performance at the desktop enough so that end users notice the difference. The first time you mentioned to end users that you were thinking about virtualizing their desktops, you probably heard concerns about what would happen to their performance. In many VDI environments, this storage performance slowdown *will* result in slower perceived application performance at the desktop if you don't do something about it.

The second impact is the cost of doing something about it. The "VM I/O blender" effect really boils down to the fact that your storage configuration is not generating enough IOPS. A common response is to throw hardware at the problem – deploy more disk spindles, buy exotic storage technology like solid state disk (SSD), or maybe even invest in a higher performance, higher end disk array. Regardless of the approach you take, you are going to be spending a lot more money on storage than you originally planned.

And there's yet another performance concern, this one unique to VDI environments. With physical desktops, each of which were standalone with their own dedicated hardware, you didn't have to deal with phenomena like boot, login, and application storms. If all your employees generally get to work around 8 AM, they're all going to boot their system up and login within a short period of time, say 20-30 min. With physical desktops, an employee booted and logged in to his dedicated desktop so there were no performance concerns with it. With virtual desktops, you may have literally thousands of employees all trying to boot their desktops and login against the consolidated hardware in the data center at roughly the same time. If you've got 70 virtual desktops per Hyper-V Host, then that host will need to be able to deal with that "spike" in activity in a timely manner. If it can't, you'll get complaints about response time delays from end users all attributed to "desktop virtualization".

Boot, login, and application storms all introduce a much wider variability in peak and average IOPS requirements than exist in virtual server deployments. For this reason, administrators are forced to "over-build" their storage configurations in VDI even more to deal with them. What it boils down to is,

as bad as storage performs in virtual server environments due to the "VM I/O blender" effect, it performs even worse in VDI environments. You'll have to throw a little *more* hardware at VDI environments to build them back up to the performance that you need.

The question is whether or not you can do that either before you run out of storage budget or before your cost/desktop becomes so high that it's just cheaper to stay with physical desktops.

*Provisioning and Deployment*

You don't just roll new desktops out once. In fact, if you're like most shops, you'll probably be refreshing your virtual desktops somewhere between 8 and 16 times per year to deal with upgrades, patches, new application rollouts, etc. A snapshot technology which can provision new desktop images rapidly from a "golden master", is very space-efficient, and can provide good performance can make things much easier. The need for these three features together is what has driven many VDI projects towards the use of high end, enterprise-class storage arrays.

It's not that you can't get the provisioning and deployment performance you want, it's how much you have to spend to get it.

*Desktop Data Protection*

And finally, how will you protect this new centralized data store, and what are the costs associated with that? When storage was sitting out on your users' desktops, they were responsible for protecting their own data. In practice what that generally meant is that it wasn't being protected. Now that the data is sitting in your data center, will you be using things like RAID and backup to protect it? If so, this is going to increase your storage costs even more. RAID configurations protect against disk failures through redundancy, and that redundancy reduces the usable storage capacity you have available for data in any size RAID array. The capacity "overhead" varies with the RAID configuration chosen (RAID 1, RAID 5, etc.). Plus, if you're backing that data up, you now need to think about things like how backups may impact users, how much storage you need for backups, etc.

When a single user's disk failed, only one user was impacted. When a disk in a VDI configuration fails, that may impact tens or hundreds of users. RAID is probably not something you can do without. You may also decide that you're not going to back up the user data to avoid some of this expense. Keep in mind, however, that much of an enterprise's intellectual property is sitting in users' desktops, and VDI gives you a much more efficient way to protect that data than when it was sitting on physical desktops that were distributed throughout your organization. The decision to deploy backup for virtual desktops is a decision, however, that must be made by each enterprise that deploys VDI.

*The Bottom Line on VDI Storage*

From an administrative point of view, VDI is very attractive as it significantly decreases the time, effort, and risk associated with management and maintenance operations while at the same time providing more centralized control for protecting corporate data (security, backups, etc.). If you've decided to deploy VDI regardless of the costs, any approach which reduces the cost/desktop merits a glance. If your deployment has been delayed or derailed because you ran out of storage budget, you might be able to get those projects going again if you could cut the storage cost/desktop in half.

## Virsto: Changing the Storage Economics of VDI Deployments

When deployed with Virsto, Hyper-V will provide high performance snapshot/clone technology that you can use without incurring any performance degradation, thin provision all storage all the time, and exhibit higher storage performance – all using any heterogeneous, block-based storage on the back end. Taken together, these features will cut your storage cost/desktop by at least half.

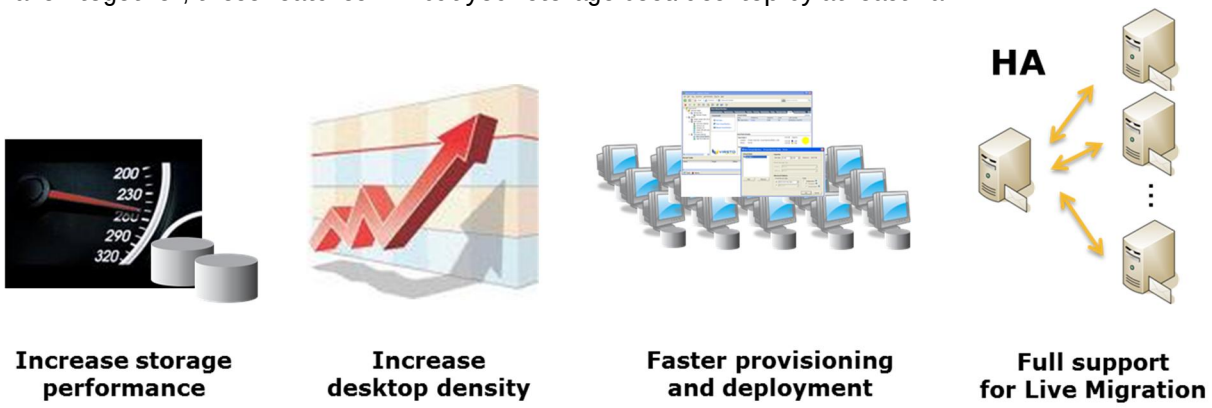


Figure 1. Virsto's combination of technical benefits results in at least a 50% *decrease* in the cost of storage/desktop.

But we do one other thing that no block-based storage alternative on the market does – we make what you have to do to manage storage as easy as NetApp *at a much lower price*. This means that the server admin, not the storage admin, can "self-provision" storage to meet their needs quickly, easily, and safely – without having to involve the storage admin. We let you operate at the "virtual machine" (VM) level – with Virsto, you don't have to deal with LUNs like all the other block-based storage alternatives, we let you deal with individual VMDKs. That means you can snapshot or Live Migrate individual VMs without having to worry about what the underlying storage configuration looks like.

When deployed with Virsto, Hyper-V offers a compelling value proposition for VDI environments:

- You will at least double the desktop density that your storage configuration can support, which means that you'll need to deploy less storage to meet your needs
- You'll be able to handle boot, login, and application storms without having to over-build your storage at significant additional expensive with exotic storage technologies like SSD
- You will cut your provisioning and deployment times by at least 35% and as much as 70%, depending on how your storage environment is configured
- You'll support critical Hyper-V features like Windows Volume Shadowcopy Services (VSS), Data Protection Manager (DPM), Windows Server Failover Clusters, and Live Migration that you may want to use to help minimize the impacts of failures and maintenance operations
- You'll get the performance of block-based storage with the ease of use of managing at the virtual hard disk (not the LUN) level

We spent the first half of this paper talking about why VDI deployments needed 30% - 50% more storage to perform at needed levels. We've just told you that a Hyper-V environment running Virsto will need at least 50% *less* storage to perform at a given level. Now let's discuss why that's true.

## Virsto for VDI, Hyper-V Edition

Virsto is a "storage hypervisor", implemented in software which deploys in Hyper-V, that provides the same basic benefits for storage that server virtualization technology brought to servers: we let you utilize your existing storage hardware at a much higher rate so that you can get more out of it in terms of performance, capacity, and management features.

Installing Virsto in a Hyper-V Host sets up a virtual storage layer that, while it looks like it is presenting standard Hyper-V "fixed" virtual hard disks (VHDs) to virtual desktops, is actually providing a storage device that is thin provisioned, performs at the limit of the underlying storage hardware (whatever that is), supports very scalable snapshot/clone technology, and is fully cluster-aware to support Hyper-V tools like Live Migration – *all at the same time*. This moves you away from native hypervisor storage options that force you to choose between performance, efficient space utilization, and rapid provisioning, allowing you to get all three without having to invest in high end, very expensive, enterprise-class array technology – it is "no compromise storage" for Hyper-V-based VDI environments.

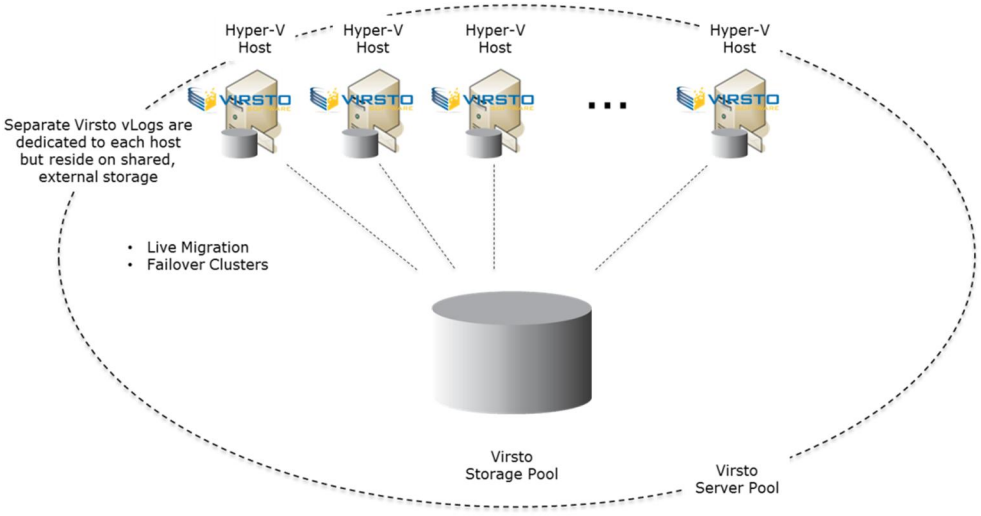


Figure 2. The Virsto Server Pool defines the ecosystem within which Live Migration is supported.

When Virsto is deployed with Hyper-V, it conforms to a pretty standard "Hyper-V cluster" configuration that we call a "Virsto Server Pool". The Virsto Server Pool is comprised of all the Hyper-V Hosts running Virsto, and the Virsto vLogs and Virsto Storage Pool running on shared, external storage. You can think of the Virsto Storage Pool as the same thing as the primary storage. During initial installation and configuration, the underlying physical storage is provisioned once. After that, Hyper-V administrators can spin VMs up or down at their whim without having to involve storage administrators to allocate new storage, create new storage devices, or reclaim storage from deleted VMs – all of that is handled automatically by Virsto. And they can do this efficiently, accurately, and safely.

Virsto installs in the parent partition of each Hyper-V Host, establishing a new virtual storage layer. Within this virtual storage layer, each Hyper-V Host gets its own dedicated vLog while at the same time sharing a pool of storage with additional Hyper-V Hosts that have Virsto installed. The storage objects presented from this shared storage look to the virtual desktops just like Hyper-V fixed disks, and they

are called Virsto vDisks. Desktops running on vDisks have no idea they are not just running on standard VHDs, but they certainly notice the performance improvements. Virsto's vDisks run at the same speed or greater as pass-through disks but they are at the same time thin provisioned and support very scalable snapshot/clone technology – all of which operate at the native performance of the underlying storage hardware. This means that, for most runtime operations, they are managed through Microsoft System Center management tools like Virtual Machine Manager, Hyper-V Manager, Data Protection Manager, and Operations Manager - just like standard VHDs. For additional integration points, there is even a Systems Center Operations Manager management pack available at no extra charge to Virsto's Hyper-V customers.

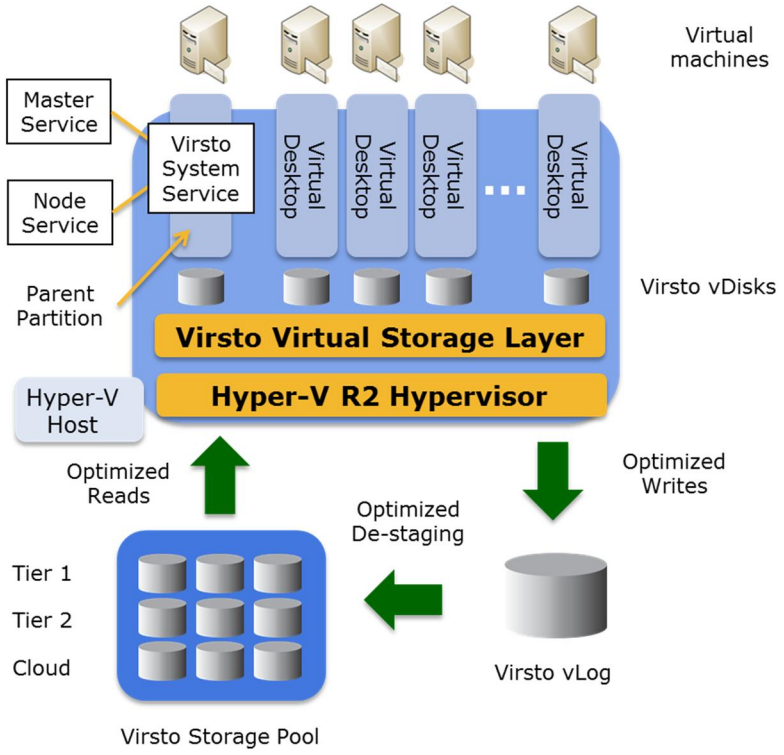


Figure 3. Virsto's logging architecture is the secret to its performance, capacity utilization, and management benefits.

The secret to Virsto's functionality is in the logging architecture that we implement. All writes from the host go to this log, which provides write acknowledgements back to the desktops. This effectively turns the very random write-intensive I/O pattern from the host into an almost 100% sequential I/O stream. The write acknowledgements occur at the *sequential* not the random write performance of whatever device is used in the log, which can speed up the perceived performance of that storage by anywhere from 30% - 10x, depending on the type of storage it is and how the storage has been configured. This speedup occurs on your *existing* storage, and doesn't require any new storage hardware purchases. A secondary process picks the writes up off the back end of the log asynchronously, and writes them through to primary storage, which Virsto calls the Virsto Storage Pool. The vLog resides on non-volatile storage that is shared, so that any writes that hit the log are immediately recoverable. If a host fails, Virsto can recover that log on another Hyper-V Host running Virsto that is attached to that same shared storage. The automatic recovery of that log on another node (i.e. failover) requires the use of additional

tools, like failover clusters, but was specifically designed this way so that it could support any high availability (HA) requirements a VDI customer might have.

This log architecture borrows from a time-proven design that has been in use in enterprise databases for decades, and is likely in use in your own environment if you are running databases like Oracle, DB2, or SQL Server. What we've done is deploy it in the "storage" instead of the "application" layer so that any application can benefit from the performance improvements it offers.

The vLog provides a very low cost, space-efficient way to handle very high peak IOPS requirements without overbuilding the storage configuration. Logs generally need to be no larger than 10GB to handle quite "bursty" write-intensive I/O workloads like login storms. If you need a larger log to accommodate your particular environment, you can just create a larger log. Boot and login storms are one of the prime drivers of SSD purchases in VDI environments, but Virsto provides what can be a much more cost-effective way to solve the same problem with spinning disk. Keep in mind, however, that Virsto can be deployed with SSD, and will increase the performance of SSD when the two are used together. You have the flexibility to decide what is best for your environment.

There is another very interesting implication to this design. From the point of view of the Hyper-V Host, the performance of the back end storage subsystem is entirely determined by the performance of the Virsto vLog. This allows operations like thin provisioning and snapshot/clone creation and usage to be performed *without any performance impacts whatsoever*. This is what allows Virsto vDisks to operate at the storage performance limit of the underlying hardware, just like a pass-through disk, while at the same time supporting thin provisioning, very rapid snapshot creation, and high performance clones. Virsto can support literally thousands of snapshots and/or clones, providing a far more scalable technology in this area than even the most expensive, high end arrays.

The ability to use very space-efficient snapshots and clones without concern for performance impacts makes your Hyper-V environment much easier to manage. Virsto's snapshots use a re-direct on write technology and are immutable, point-in-time images of any vDisk or set of vDisks you select, whereas Virsto clones are created from the snapshots when you want to read or write to them. You can create snapshots almost instantly without consuming any additional storage capacity, facilitating their use to make operations like data protection, maintenance (where they can be used to easily provide a quick rollback safety net), desktop provisioning, and VM template creation much easier. Clones can also be created very rapidly and consume no additional storage until you start to write to them.

In failover clusters, Virsto vDisks co-exist with Hyper-V clustered shared volumes (CSV), enhancing the performance of these environments. A CSV is a volume that is available to directly read from and write to by all nodes in a failover cluster, and it's the best place to put the VM configuration data in a failover cluster. Virsto vDisks can be used for the VM images and user data in failover clusters, and can add significant value here. In this environment, vDisks deliver high performance, thin-provisioned VHDs that support very scalable snapshot/clone technology that makes Hyper-V perform even better during runtime operations as well as during VSS/DPM backups. Virsto supports very granular operations at the VM level – not the LUN level – even though it is running on top of block-based storage. This means you get the flexibility you need to Live Migrate individuals VMs where you want them without having to worry about how you've laid out your underlying LUNs or the latency impacts of NAS.

*Managing Virsto in a Microsoft System Center Environment*

Because vDisks look like standard Hyper-V VHDs, they fit very nicely into the familiar System Center management framework. Because Virsto leverages block-based storage, it can provide higher performance than true NFS options but, despite the fact that it is block-based, Virsto allows all storage operations to occur at the granularity of individual VMs. This overcomes the major problem with other block-based storage alternatives, which is that they can only manage at the LUN level.

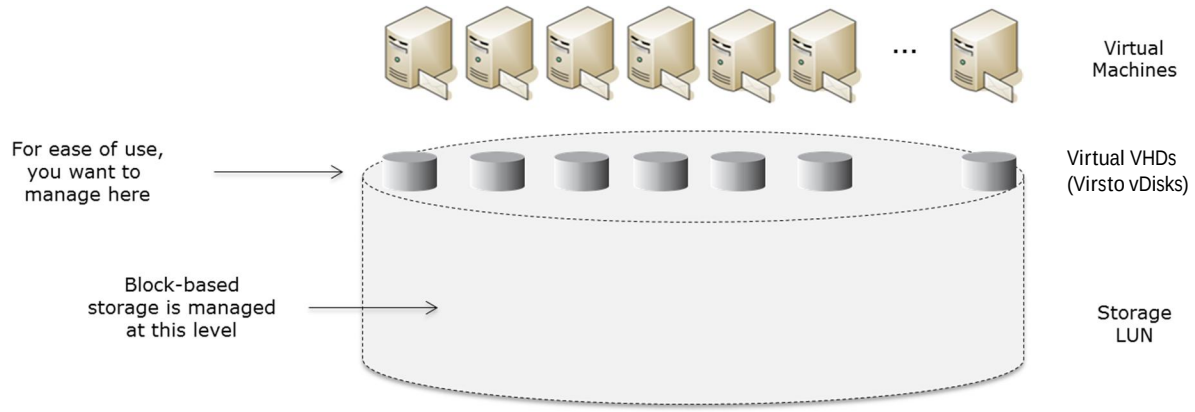


Figure 4. The difference between VM and LUN level storage management.

Understanding this distinction is critical to understanding why Virsto makes storage as easy to manage as NAS. Most block-based solutions only know about LUNs, which are the lowest level storage objects they can manage. This means when you create snapshots or want to move things around, like VMs with a Live Migration, you can only do that at the LUN level. But what you want to operate on are individual VMs, which is really the VHD level. You might want to snapshot a particular VM to provide a safety net to rollback to if necessary during maintenance operations or do a snapshot backup of just one VM so you can reduce the impact of backups. Most block-based storage doesn't know about the individual VMs, just the LUNs. So if you want to snapshot a particular VM, you actually have to snapshot the entire LUN that it's on. Since a storage administrator is going to minimize the number of LUNs they have to manage for ease of use, you will almost always have a lot of VMs on a particular LUN. So by snapshotting an entire LUN just to get one VM, you end up snapshotting a lot of VMs that you may not be interested in. If this sounds wasteful, it is. It takes extra time and can use up extra storage capacity.

Let's look at the Live Migration example. For a given LUN where 15 VMs may reside, there may only be 2 that I want to move around for whatever reason. If I can only manage at the LUN level, I have to move all of them around. That takes up more cycles and time to move a bunch of VMs that I'm not interested in.

What administrators want is the ability to select individual VMs and perform storage operations like snapshots and Live Migrations on just those. This is what NAS does for you, and one of the reasons why "file-based storage" is viewed as easier to manage than "block-based storage". This is also what Virsto does for you, while at the same time providing you the performance of block-based storage. And deploying Virsto instead of a high end NAS box can save you a lot of money. With this unique "ease of use" distinction, Virsto makes Hyper-V administrators more productive, allowing them to serve their own storage needs quickly, easily, and without risk.

For Virsto for VDI, Hyper-V Edition, Virsto's management strategy is to fit into the native System Center management environment. The Virsto management GUI is in fact a Microsoft Management Console (MMC) snap-in. All of the initial installation and configuration is done from there. Once installed, administrators have the choice of managing through the standard VMM wizards for most operations. For certain common operations in VDI environments, like VM creation or desktop provisioning, Virsto provides their own tools that, while they closely follow the look and feel of the native Hyper-V Manager and VMM workflows, provide optimizations not available within those tools today for operations like bulk creates/deletes, attach/detach, space reclamation, and rapid provisioning of thousands of desktops. When a VM is deleted, Virsto reclaims that space immediately and returns it to the storage pool – you do not have to wait until you've rebooted that Hyper-V Host.

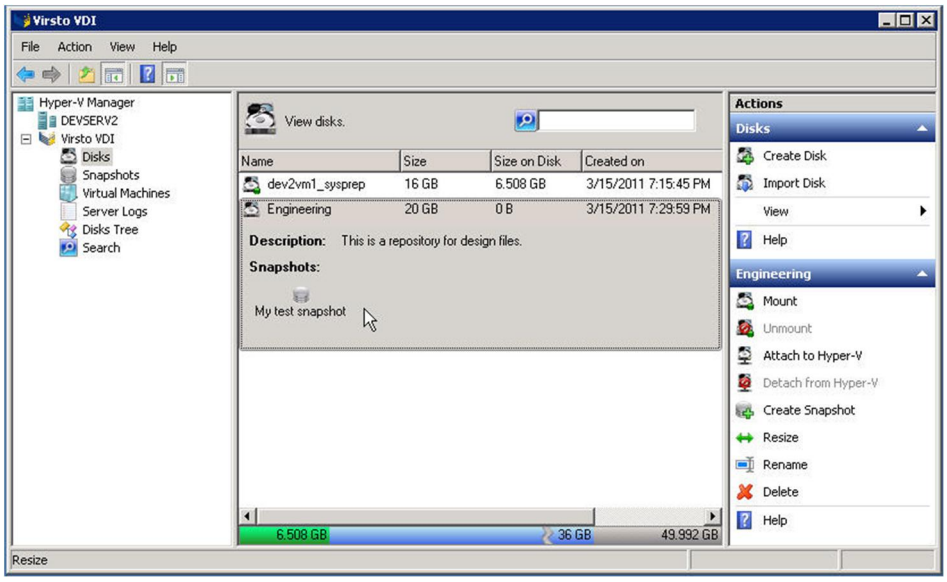


Figure 5. Virsto's MMC snap-in offers familiar management workflows to Hyper-V administrators for managing storage.

Virsto supports a storage tiering capability that provides additional options for creating blended storage configurations that cost-effectively meet performance requirements in VDI environments. Virsto supports up to 4 storage tiers, statically defined, and can house a variety of different types of block-based storage in these tiers, including SSD, FC, SAS, SATA, iSCSI, and even cloud-based storage (as long as it is a block-based target). A particularly common use of storage tiering in VDI environments can take a "golden master", which contains data like operating systems, applications, and other shared data which is commonly accessed, and places it on a very high performance, SSD-based tier. The user data for each individual desktop does not require as high a performance, and can be placed on a lower-cost tier which resides on spinning disk.

Virsto also supports a variety of other tools in common use in Hyper-V environments, including Windows Server Failover Clusters, CSVs, Live Migration, Windows VSS, DPM, Virtual Machine Manager, Hyper-V Manager, and Operations Manager. Although Hyper-V customers may not want to configure failover clusters for use with individual desktops, there are projects where it is advantageous to be able to automatically recover failed desktops deployed using persistent mode or to fail over all desktops on a given host prior to performing maintenance on it. Virsto supports these operations.

## Summary

Virsto provides a storage hypervisor, implemented at the server hypervisor level through a simple software deployment, that changes the economics of VDI deployment to enable cost-effective deployments. Virsto improves the desktop density that any given storage configuration can support by at least 2x, which means that you can either spend 50% on your storage or support twice the number of virtual desktops before you buy more storage. It works with any heterogeneous, block-based storage, yet offers the ease of management of NAS-based storage to make it fast, easy, and safe for Hyper-V administrators to provision and manage their own storage as they spin VMs up and down.

With Virsto deployed, you'll notice the following from your Hyper-V environment:

- All virtual storage appears to perform at the storage performance limit (pass-through disk speed) of the underlying hardware, despite the fact that it is thinly provisioned to save on storage capacity, which can be 30% - 10x faster than your storage is now
- You'll provide the performance needed to meet peak IOPS requirements for boot, login, and application storms without having to overbuild your storage configuration, thereby saving significant cost
- You can freely use snapshot/clone technology to facilitate data protection, maintenance, template creation, and other operations without incurring any performance degradation and without having to buy high end, enterprise-class storage arrays
- You'll be more productive on tasks where you have to provision storage for your desktops because you won't have to involve the storage administrator during VM spin up/spin down (after initial configuration), which on initial rollout and refresh tasks can save you hours or days
- You'll be able to leverage native Hyper-V features like Windows Server Failover Clusters and Live Migration as needed to improve availability or minimize the impact of maintenance operations on your desktops

Starting at a list price of only $2,800 per host, there is no other option on the market today that can provide the storage performance, cost, and management advantages really needed to make VDI environments affordable. Virsto will let you deliver high performance for your virtual desktops predictably and cost-effectively, and could be the difference between moving to a VDI environment that saves you time on administrative tasks or staying with the old physical desktops because of end users' performance concerns.

Designed for use in server-hosted virtual desktop environments, Virsto for VDI, Hyper-V Edition supports a variety of virtual desktop front ends, including Microsoft VDI and Citrix XenDesktop. Support for Quest vWorkspace is planned as well.



© Copyright 2011, Virsto Software Inc. All Rights Reserved. This document is provided for information purposes only, and the contents hereof are subject to change without notice. The information contained herein is the proprietary information of Virsto and maynot be reproduced or transmitted in any form for any purposes without Virsto's prior written permission. Virsto is a registered trademark of Virsto Software Corporation.