



- **Reduce storage provisioning times by up to 75%**
- **Improve vSphere storage performance by 3x – 9x**
- **Reduce vSphere storage capacity consumption by up to 90%**
- **Double the VM density any given storage configuration can support**
- **Enable safe “self-service” provisioning to improve IT responsiveness**

“We saw a 9.7x improvement in desktop performance using Virsto, which allowed us to reduce our storage cost per desktop by 70%.”

Kyle Murley

Server, Desktop, and Web Services Team Lead

San Diego State University

Server virtualization technology has achieved mainstream use, improving flexibility in IT environments as well as significantly lowering the cost/virtual machine (VM). Server hypervisors achieved this by at the same time virtualizing and improving the utilization of server resources. But because legacy storage architectures perform quite poorly in virtualized environments, storage costs have actually increased.

The concept of a “storage hypervisor” is very attractive. By

Increasing the utilization of storage hardware to speed provisioning, improve performance, and reduce capacity consumption, it will fundamentally change the way vSphere handles I/O to significantly reduce costs. And by virtualizing storage, it will improve flexibility, making management easier. The storage hypervisor is what is needed for virtual computing to comprehensively deliver on its promise.

WHY IS STORAGE SO EXPENSIVE IN VIRTUAL ENVIRONMENTS?

Storage costs are unexpectedly high in VM environments because of a phenomenon known as the “VM I/O blender”. For a server deployed according to the legacy client/server model, there was a dedicated application running on a dedicated server, often with its own dedicated storage. Given that there was only one I/O stream coming out of that server, there were optimizations that could be performed to deliver good storage performance.

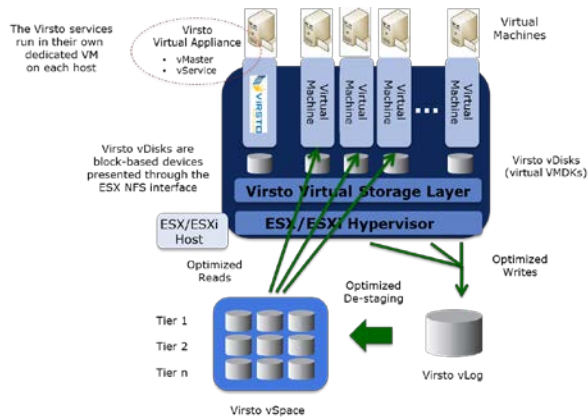
This is very different from a virtual computing environment where you may have anywhere from 5 to 100 VMs, each running their own independent workload, on a single physical host. That creates a very random I/O stream that slows down spinning disks, resulting in increased latencies and fewer IOPS. This impact is not trivial – it can cause storage to underperform by as much as 30% - 50%.

Because of these issues, administrators generally find themselves having to invest in additional spindles, expensive options like solid state disk (SSD), or enterprise-class disk arrays to obtain needed performance. If you want rapid provisioning and efficient storage capacity consumption, you’ll need to buy even more storage since conventional thin provisioning lowers performance as well. Either way, storage ends up costing a lot more.

VIRSTO CHANGES THE STORAGE ECONOMICS OF vSPHERE

When deployed with Virsto, vSphere will reduce VM provisioning times by up to 75%, increase overall storage performance by 3x – 9x, reduce storage capacity consumption by up to 90%, and enable safe, reliable self-service provisioning that does not require storage administrator involvement – all using any block-based, heterogeneous storage (spinning disk or SSD). Taken together, these capabilities shorten desktop deployment times and test cycles, cut storage cost per VM by at least half, and increase the VM density that any given storage configuration can support, without having to invest in any additional storage hardware.

Virsto fundamentally yet transparently changes the way vSphere handles storage I/O to achieve these gains. Applications are unaware of Virsto, and familiar vCenter management workflows are still used to manage the storage, leveraging existing skill sets.



Virsto works transparently in ESX environments while providing significant provisioning, performance, and management advantages.


THE VIRSTO ARCHITECTURE

Virsto installs in each ESX host as a virtual storage appliance, establishing a virtual storage layer that presents storage objects that look to ESX just like eager zeroed thick VMDKs. Called Virsto vDisks, they are compatible with native vSphere functionality including VMware HA, VMotion, and VMware Tools.

Virsto transparently implements a log architecture that has far reaching implications for storage operations in vSphere environments. All writes hit the vLog all the time, effectively sequentializing them to provide predictable storage performance with lower average latencies. Because of this, the vLog for each host always appears to the VMs to operate in streaming write mode, delivering significant performance gains. The vLog provides a very cost-effective way to accommodate high I/O spikes like those generated by boot and login storms without impacting overall performance and throughput and without having to purchase any additional storage hardware. Either spinning disk or SSDs can be used in the vLogs – Virsto will speed up both.

Writes are asynchronously de-staged from each vLog and written to vSpace, a datastore that is shared across all ESX hosts running Virsto. The layout of data in vSpace is optimized to prevent fragmentation and deliver very high sequential read performance. All vDisks in vSpace are thin provisioned, yet deliver better performance than thick VMDKs because of the vLogs. Additional performance gains come from the fact that the read and write workloads are segregated onto different spindles.

The Virsto datastore looks to ESX like an NFS datastore, even though it is running on heterogeneous, block-based storage. Virsto supports up to 4 tiers in the datastore, allowing the creation of blended storage configurations. For example, an SSD-based tier 0 can be used to store “golden masters” in VDI environments, while more cost-effective SATA can be used for private desktop data. The Virsto datastore can be up to 1PB in size without any performance degradation.



Virsto’s thin provisioned storage outperforms eager zeroed thick VMDKs, even when snapshots and clones are in use.

Virsto’s design delivers almost instant storage provisioning, and can support literally tens of thousands of snapshots (vSnaps) and writable snapshots (vClones) without incurring any performance degradation at all. This rapid provisioning can cut hours or days off of the storage provisioning times associated with initial desktop deployments or refresh, and can shortcut VM spinup/spindown times to help get products out the door faster in test/dev environments. Virsto also safely and reliably manages all storage allocation and reclamation, enabling self-service provisioning and de-provisioning of VMs for end users and administrators that may not have storage expertise.

The storage objects Virsto presents are individual VMDKs, allowing storage operations like snapshots, live migration, fail-over, and replication to occur on individual or selected VMs without regard to how the underlying storage LUNs are configured. This makes Virsto storage much easier to manage, not to mention more efficient in its use of storage capacity, time, and network bandwidth.

RAPID DESKTOP PROVISIONING

Virsto also provides Rapid Provisioning Wizards for use with the two most popular VDI front ends: VMware View and Citrix XenDesktop. These provisioning wizards preserve the familiar View and XenDesktop workflows while leveraging Virsto capabilities to rapidly provision very high performance storage. The wizards support initial deployment and desktop refresh, transparently integrating with the relevant desktop front ends through connection brokers. This allows desktops created with Virsto to take advantage of Virsto’s capabilities while still being managed with familiar tools.

VIRSTO: THE STORAGE HYPERVISOR FOR VMWARE

Virsto will improve the utilization and flexibility of existing storage to make vSphere deliver on the comprehensive promise of virtual computing. If this sounds interesting, let’s talk.

