

VirtuCache + VirtuStor Case Study



USE CASE

- Backup Storage using Veeam and SuperMicro

LOCATION

- New York, USA

VIRTUALIZATION AND STORAGE ENVIRONMENT

- Disaster Recovery Hosts: 2 x Second User SuperMicro servers running VMware Essentials Licenses
- Disaster Recovery Storage: 3 Server VirtuStor cluster running on repurposed SuperMicro servers with 24 TB raw capacity. Two Live Instances of data, so useable capacity was 12 TB.
- Veeam was used to backup and replicate VMs to the VirtuStor Storage cluster.

CHALLENGES

- Deploy offsite Disaster Recovery at a very low cost

SOLUTION

- VirtuStor allowed St James Hospital and other hospitals in the collective to share IT infrastructure for high performance Disaster Recovery

Building a Cost-Effective Disaster Recovery Solution from Existing Servers

Why deploy Disaster Recovery (DR) infrastructure when you already have backups?

The main reason to deploy DR infrastructure is to protect against the rare case where you lose your part or your entire datacenter. In case of a datacenter wide outage, you can flip over to your offsite DR data center. The DR datacenter is a replica of your entire production infrastructure, from applications running in VMs to servers, storage, and network configuration.

DR and backups though related are two different things involving slightly different processes and underlying technologies, resulting in varying times to restore the VM state and data.

Other than the rare datacenter wide outage, there are more frequently occurring failure situations where leveraging DR infrastructure is a better choice than restoring from backups. For instance, even if one VM were to fail, you can flip over to the replica of that VM in your DR site and have end users use the replica VM as the interim production VM till you restore your primary production VM. You can restore your production VM from your replica VM while end users are using the replica VM (now acting as interim production VM) in parallel.

With DR infrastructure, you can also search and replace corrupt or deleted files, databases, mailboxes, across all VMs without restoring the VMs.

Provided you have available rack space, and if you already have a process for doing backups, the incremental cost to deploy DR can be reasonable. You can repurpose older servers to use as both DR VMware hosts and DR SAN storage (how to repurpose older servers to build SAN storage is described in later sections), then you can use VMware Essentials licenses, and lastly you need backup/DR software. A software like Veeam that you might already be using for backups, has DR functionality in its basic license, so no additional licensing cost are incurred here.

What is CEPH storage? And it's relationship with VirtuCache and VirtuStor?

CEPH is open source storage software that runs on commodity servers. It clusters servers together and presents this cluster of servers as an iSCSI appliance. VirtuStor enhances CEPH with an iSCSI module to interface with VMware developed software for VAAI (storage offload from VMware); and built an easy to use GUI.

Servers of any make, model, size, and antiquity can be 'hot' added to an existing VirtuStor cluster to add capacity or improve performance.

From a pricing point of view VirtuStor gives smaller hospitals the enterprise storage features they need, but at much lower costs than traditional SAN storage appliances.

The VirtuCache Difference

Built a low cost, high performance Disaster Recovery solution that could be the foundation of a shared DR infrastructure across other facilities.

VirtuCache + VirtuStor Case Study



Building a Cost-Effective Disaster Recovery Solution from existing servers

Why is CEPH/VirtuStor suitable for shared DR storage?

VirtuStor storage has its origins at cloud service providers (SPs). The fact that commodity servers can be used to build SAN storage was important to cloud SPs to keep their hardware costs low. Low cost storage is a key requirement for on-premises DR storage.

Since CEPH is also used by cloud SPs, it also has features to isolate and encrypt data and storage I/O path between multiple organizations that might be using the same CEPH storage cluster, a requirement if different organizations are to share the same storage hardware.

As you scale out the storage cluster, cost per capacity reduces dramatically. Starting at \$2/GB for raw 5TB storage, the cost drops to 20 cents/GB for 300TB of storage. And so it is cost effective for smaller IT departments to pool together their DR/backup budgets to get larger amounts of storage for their DR infrastructure.

Sharing compute using VMware.

VMware lends itself well to a shared DR infrastructure as well. Since each VMware physical server can host a maximum of 512 VMs, large number of VMs can be deployed on only a 2-host VMware cluster.

DR Infrastructure at St. James Hospital.

Currently the DR infrastructure at St. James has:

- 2-host compute cluster using VMware Essentials license running on repurposed servers. The license cost is \$600 for VMware Essentials license.
- 3-server VirtuStor cluster for iSCSI storage using older servers but with new storage media. It has raw capacity of 24TB (usable of 12TB). The cost for the 12TB usable is about \$15K.
- Veeam is used to replicate data from St. James' production VMware cluster to this DR infrastructure. The cost of Veeam Enterprise Plus Essentials is ~ \$7K.
- So the cost of the entire infrastructure with one time services fee to put this together was \$30K.

By simply adding a few more hard drives and SSDs, this infrastructure had the capacity to accommodate the DR workload for two more hospitals of the same size as St. James.

The incremental cost to share St. James' DR infrastructure was \$10K per year for replicating 30 or so VMs and 10TB of data.

Furthermore, this idea of shared backup and DR infrastructure can be shared across other hospitals in the collaborative that St. James belongs to, both VMware and VirtuStor storage clusters could be scaled up by adding more servers to each cluster.

This configuration could support hundreds of organizations with a single clustered DR environment.