



VM-Aware Software Defined Storage

Optimizing NFS for virtual server workloads

Solution Whitepaper

Version 2.0
May 2018

Table of contents:

Executive overview 3

Introduction 3

VM-aware storage for VMware vSphere 3

 VM storage provisioning 4

 Intelligent client-side caching 5

 VM movement with VMware vMotion 5

 VM snapshots and clones 5

 Datastore migration 6

 High availability with vSphere Metro Storage Cluster (vMSC) 6

 Additional features 6

Summary and conclusion 7

Additional resources 7



2350 Mission College Blvd, Suite 500
Santa Clara, CA 95054

hedvig.io

Executive overview

As Software-defined storage is transforming data centers from static, inflexible infrastructures to modern architectures that are dynamic and agile, enabling organizations to quickly adapt to rapidly changing business requirements. The Hedvig Distributed Storage Platform provides VM-aware storage that simplifies storage management for virtualization administrators. Hedvig is designed to meet the unique requirements that a dynamic virtual server environment places on storage and provides the flexibility to meet the needs of any application. Taking advantage of a distributed systems architecture with industry-standard server infrastructure and public cloud services to form a scalable private, hybrid, or multi-cloud storage system, the Hedvig solution cost-effectively supports the needs of organizations as their environment grows and business requirements change.

Introduction

Hedvig, the industry's most complete software-defined storage platform, delivers a modern solution for virtual server environments. With built-in intelligence for the management of virtual machine (VM) storage through the entire VM lifecycle, Hedvig software helps ensure performance and availability while delivering a lower total cost of ownership.

Virtualization presents a unique set of challenges for storage. This includes the need to support multiple application types, handle consolidated and random I/O, deliver seamless VM movement, and ensure availability across sites. To meet these challenges, Hedvig has designed its support for virtual server infrastructure to be VM-aware, offering intelligent provisioning and management of storage volumes to ensure customers can achieve their goals in hypervisor-based clouds.

VM-aware storage for VMware vSphere

The VMware vSphere data center virtualization platform offers a wide range of capabilities for managing virtual machine infrastructure at large scale. Beyond the hypervisor, advanced features that enable workload mobility help organizations ensure applications are available and running at maximum efficiency. To take full advantage of these VMware capabilities, the storage platform supporting a vSphere environment must be able to provide the intelligence to seamlessly enable data availability and performance no matter where in the environment a VM may be running.

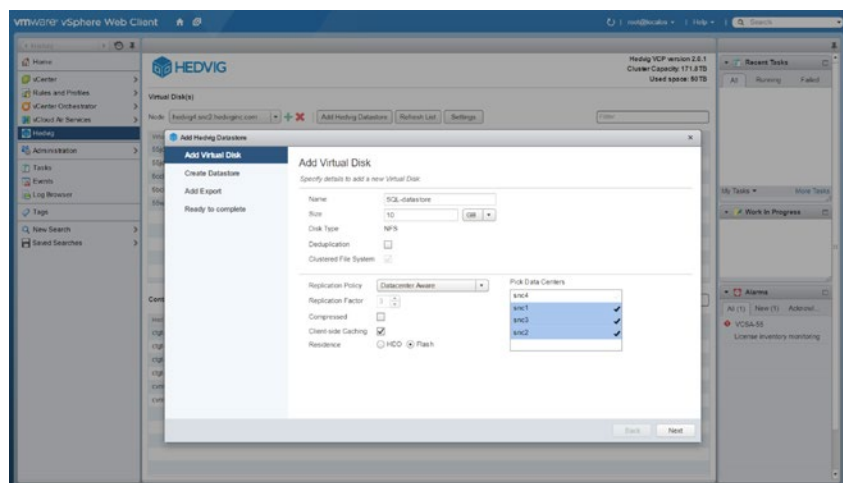
Hedvig meets these requirements through VM-awareness built-in to the Hedvig Distributed Storage Platform. Tight integration with vSphere ensures that operations throughout the VM lifecycle from datastore provisioning to VM movement and migration are simple, seamless, and automated. Specifically, the Hedvig NFS protocol implementation provides the unique ability to create and manage Virtual Disks per-VM while providing intelligent storage services tied directly to each VM.

VM storage provisioning

To simplify provisioning and management of storage for a VMware vSphere environment, Hedvig provides a vSphere Web Client Plugin that integrates Hedvig functionality with VMware vCenter and the VMware vSphere Web Client. This puts Hedvig's comprehensive suite of enterprise storage capabilities at the fingertips of a virtualization admin. With a few clicks administrators can configure and share NFS-based datastores provisioned as virtual disks, a volume abstraction created from the pooled resources available within the Hedvig storage cluster. Each datastore is customizable with a set of features to best meet the performance, protection, and disaster recovery needs of a given application.

As VMs are created and assigned to a datastore, Hedvig recognizes each virtual machine disk (VMDK) file and automatically creates a corresponding “child” virtual disk within the NFS “master” virtual disk that corresponds to the datastore. Through awareness of VMs and corresponding VMDK files, the Hedvig software is able to perform management and operations at the virtual machine level. This includes features that facilitate performance, efficiency, seamless VM movement, and data migration.

Figure 1 — Provisioning a datastore with the Hedvig vSphere Web Client Plugin



Intelligent client-side caching

For VMware datastores that have client-side caching option enabled during provisioning, the Hedvig Storage Proxy installed on the ESXi host stores a working set of data blocks to local SSD/PCIe drives to accelerate reads. This block cache returns data directly from local flash media, dramatically improving efficiency, accelerating read operations and avoiding network hops when accessing VM data. Through VM-awareness Hedvig manage caching per-virtual machine enabling intelligent cache management as VMs start, stop, and move within a vSphere cluster.


VM movement with VMware vMotion

For users who take advantage of vSphere vMotion to move virtual machines from one physical server to another, Hedvig’s VM-aware storage solution recognizes the operation and ensures consistent, continuous access to underlying virtual disks. Datastores provisioned on Hedvig can be shared on all hosts in cluster, each with access to the same NFS export IP address. A VM that moves locations connects and resumes I/O to its virtual disk through the Hedvig Storage Proxy on the new vSphere host. For datastores with caching enabled, block cache operations begin on the new host to optimize application performance. If a VM is moved back to its original vSphere host, Hedvig again recognizes the change, resumes I/O, and starts a new block cache operation to ensure consistency with any data written while the VM was running on the alternate host. Through this seamless support for VM movement with vMotion, Hedvig enables administrators to make infrastructure decisions and easily move VMs as needed.

SVM snapshots and clones

Another key advantage of Hedvig's VM-aware solution is the ability to take advantage of metadata-based, zero-impact snapshots and clones driven directly from the vSphere Web Client interface. With Hedvig you can create snapshots and clones of VMs in a matter of seconds. The Hedvig software first calls the VMware snapshot framework to flush any active I/O for the VM. Once the metadata-based Hedvig snapshot is made, the VMware snapshot is removed. This approach provides the benefit of eliminating the consumption of extra capacity overhead typically required with the VMware snapshot framework. A snapshot can be used to revert the VM to a specific point-in-time, or to create one or more clones. Clones can be used to create test VMs or new active production VMs. These VMs are space-efficient as each will be linked to the underlying source virtual disk on the Hedvig platform.

Figure 2 – Virtual disk actions in the Hedvig vSphere Web Client Plugin



Datastore migration

An important, but less frequently used operation within a vSphere environment is the migration of a virtual machine and its VMDKs from one datastore to another. The most typical reasons for using this operation is to perform maintenance or to address performance issues that arise from over-utilized VMFS datastores. To support this operation VMware provides Storage vMotion, a process that copies VM data to a new location while the virtual machine is running.

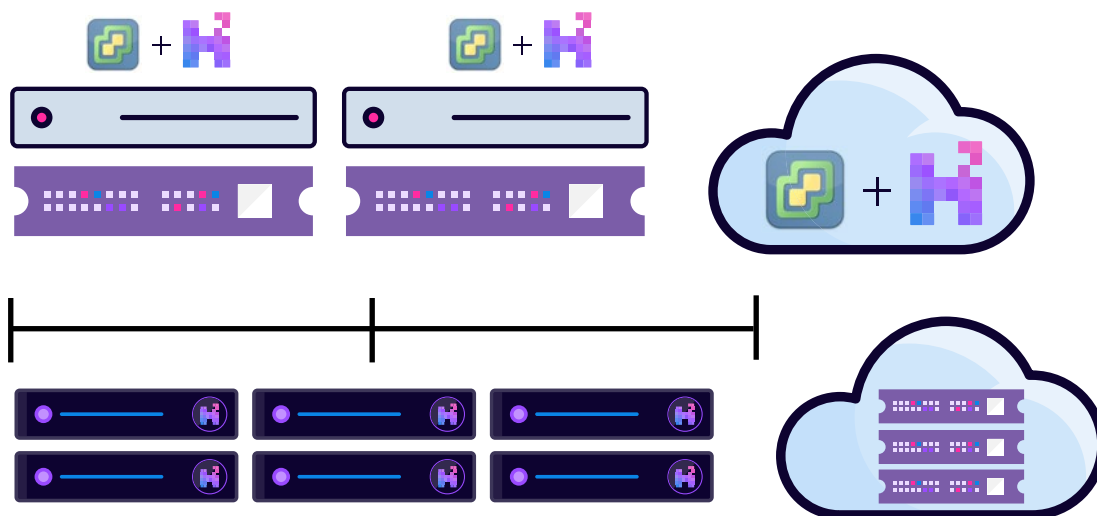
The impediment with Storage vMotion is that a physical move of an entire VM data set requires significant network I/O to transfer the data and consumes CPU resources. The operation will typically have an impact on the other VMs on the involved vSphere hosts. To alleviate this resource impact when moving VMs between datastores within a Hedvig cluster, Hedvig provides a VM-aware migration capability that quickly reassigns VM data from a current datastore to a different datastore. This feature, available through the Hedvig vSphere Web Client Plugin, simply unlinks the child virtual disk corresponding to the VM from the source master NFS virtual disk and links it to the destination virtual disk in seconds while ensuring that VMware is aware of the new addition to the datastore.

High availability with vSphere Metro Storage Cluster (vMSC)

In environments where disaster avoidance and high availability for applications is a requirement, a VMware deployment can take advantage of stretch clustering. VMware vSphere Metro Storage Cluster (vMSC) is one such solution and enables a vSphere cluster to be active across two sites. The vMSC setup requires a single storage subsystem that spans datacenters. The Hedvig solution is an ideal fit for these deployments, providing a single multi-cloud capable storage cluster that can span multiple locations to support an active/active stretched cluster.

Like the operations with vMotion, VM movement on a stretched cluster takes advantage of the Hedvig Storage Proxy on each of the vSphere hosts to deliver the most optimal access path for the VM to its data. As a VM moves across a stretched cluster, it accesses data locally regardless of its physical location. What makes this possible is that the NFS-based datastore stored on Hedvig can be defined with a replication policy that distributes data across datacenters. The local Storage Proxy is aware of where copies are stored in the appropriately directs I/O to local nodes to minimize latency. Hedvig's built-in consistency mechanisms ensure that data written in the new site is successfully replicated across all relevant nodes and sites within the defined cluster.

Figure 3 – Multi-cloud storage operations for stretched clusters



Additional features

In addition to the VMware-specific features described above, Hedvig natively provides additional enterprise storage services that provide benefit in a virtualization environment.

This includes:

- › Global deduplication—Maximizes storage efficiency with deduplication of dedupe-enabled virtual disks across the entire cluster
- › Source-side deduplication—Improves backup throughput and reduces the amount of data required to be sent across the network
- › Compression—provides additional data reduction post deduplication for maximum efficiency
- › Thin provisioning—Enables datastores to be provisioned without requiring preallocation and reservation of unused physical disk capacity.

For a full overview of the Hedvig Distributed Storage Platform features and advantages, download the Hedvig Technical Overview Whitepaper.

Summary and conclusion

Virtualization transforms enterprise IT, increasing productivity, lowering costs, and improving business continuity. VMware has led the way to an era of software-defined, service-oriented cloud architectures and on-demand IT. The Hedvig Distributed Storage Platform is the industry's most complete software-defined storage solution built to be VM-aware and deliver elastic, resilient, modern storage for VMware vSphere. Organizations with fast-growing, virtualized applications and data in a single data center or across multi-site deployments benefit from the ability to dynamically provision and scale storage in private and public clouds with unprecedented simplicity and speed.

Additional resources

Hedvig and virtualization: <http://www.hedvig.io/resources/topic/use-case?cat=servervirtualization>

Hedvig product information: <http://www.hedvig.io/product>

Hedvig technical overview whitepaper: <http://www.hedvig.io/technical-overview-whitepaper>



2350 Mission College Blvd, Suite 500
Santa Clara, CA 95054

hedvig.io

LICENSES AND COPYRIGHTS

Hedvig Inc. believes the information in this publication is accurate as of its publication date. The information is subject to change without notice. The information in this publication is provided as is. Hedvig Inc. makes no representations or warranties of any kind with respect to the information in this publication and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any Hedvig Inc. software described in this publication requires an applicable software license. All trademarks are the property of their respective owners.

©2018 Hedvig Inc. All rights reserved. | Version 2.0 | Published in the USA