

Hedvig Limited Edition

# Hyperscale Storage FOR. DU

#### Learn to:

- Store data like web-scale companies
- Scale storage performance or capacity as needed
- Make the business case for hyperscale storage

Lawrence C. Miller, CISSP



A Wiley Brand

These materials are @ 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.



### by Lawrence C. Miller, CISSP



These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

#### Hyperscale Storage For Dummies®, Hedvig Limited Edition

Published by John Wiley & Sons, Inc. 111 River St. Hoboken, NJ 07030-5774 www.wiley.com

Copyright © 2016 by John Wiley & Sons, Inc., Hoboken, New Jersey

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without the prior written permission of the Publisher. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

**Trademarks:** Wiley, For Dummies, the Dummies Man logo, The Dummies Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. Hedvig and the Hedvig logo are trademarks of Hedvig, Inc. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc., is not associated with any product or vendor mentioned in this book.

LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.

For general information on our other products and services, or how to create a custom *For Dummies* book for your business or organization, please contact our Business Development Department in the U.S. at 877-409-4177, contact info@ dummies.biz, or visit www.wiley.com/go/custompub. For information about licensing the *For Dummies* brand for products or services, contact BrandedRights& Licenses@Wiley.com.

ISBN 978-1-119-29272-2 (pbk); ISBN 978-1-119-29275-3 (ebk)

Manufactured in the United States of America

 $10\ 9\ 8\ 7\ 6\ 5\ 4\ 3\ 2\ 1$ 

#### **Publisher's Acknowledgments**

Some of the people who helped bring this book to market include the following:

Project Editor: Jennifer Bingham

Acquisitions Editor: Katie Mohr

**Business Development Representative:** Karen Hattan

**Special Help from Hedvig:** 

#### Editorial Manager: Rev Mengle

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

## Table of Contents

. ..

Introduction	1
About This Book	1
Foolish Assumptions	1
Icons Used in This Book	ء 2
Beyond the Book	، 2
Chapter 1: Understanding the Hype about	
Hyperscale II	····· ·
What Is Hyperscale IT?	ې ••••••
Who Pioneered Hyperscale IT and Why?	4
Why Is Hyperscale IT Useful in the Modern Enterprise?	2
<b>Chapter 2: Assimilating Hyperscale Storage</b>	e (
Defining Hyperscale Storage	
Traditional Storage versus Hyperscale	11
Heading Off Exponential Data Growth	11
Chapter 3: Comparing Hyperconvergence	
and Hyperscale	17
Examining the Common Components	17
Recognizing the Benefits and Challenges of Both Approaches	

Chapter 4: Exploring Hyperscale	
Storage Use Cases	21
Virtualization	21
OpenStack	23
Docker	23
Backup and Archive	24
Chapter 5: Introducing the Hedvig Approach to Hyperscale	29
Hedvig Architecture Overview	
Hedvig Storage Service	
Hedvig Storage Proxy	
Hedvig Virtual Disk	35
Chapter 6: Ten Reasons Hyperscale Sto	orage for 37
Glossarv	Д1
G103501 y	

## Introduction

Current approaches to storage are broken and can't keep pace with the rate of change in today's modern business. Businesses generate, analyze, use, and store more data than ever before. As a result, IT can't even accurately predict storage requirements three months out, never mind the three to five years that typify storage refreshes.

A new approach to data center infrastructure is needed to help companies dynamically adapt to rapidly changing business requirements

## About This Book

This book explains hyperscale and hyperconverged storage technologies and architectures to help you understand the similarities and differences of both approaches, so that you can deploy the best solution whether it's hyperscale, hyperconverged, or a hybrid mix of both — to support your organization's particular storage requirements.

## Foolish Assumptions

I assume you are some combination of a cloud, data center, infrastructure, storage, or systems administrator, architect, engineer, or manager. As such, this book is written primarily for technical readers. However, it's not too technical — there's even a glossary in the back I also assume that you're looking for an innovative solution to help address your organization's storage challenges, and you're somewhat familiar with cloud and virtualization technologies. If these assumptions describe you, then this book is for you!

## Icons Used in This Book

Throughout this book, I occasionally use icons to call out important information. Here's what to expect.



2

This icon points out information you should commit to your nonvolatile memory.



This icon explains the jargon beneath the jargon!



This icon points out helpful suggestions and useful information.



This icon points out the stuff your mother warned you about. Okay, probably not. But you should take heed nonetheless!

## Beyond the Book

Although this book is chock full of information, there's only so much I can cover in 48 short pages! So, if you find yourself at the end of this book thinking "Gosh, this was an amazing book; where can I learn more about hyperscale storage?" just go to www.hedviginc.com.

## **Chapter 1**

## Understanding the Hype about Hyperscale IT

### In This Chapter

- Managing growth at scale
- Recognizing the pioneers of hyperscale IT
- Examining islands of storage in the modern business

n this chapter, you learn about the rise of hyperscale IT, who is using it, and why.

## What Is Hyperscale IT?

IT organizations have traditionally managed growth with a scale-up architecture: When additional performance or capacity is needed, resources or components (such as processors, memory, network interfaces, or hard drives) are added within a server or storage device.

*Hyperscale IT* (also known as web-scale IT) is scale-out architecture that enables a more flexible and economic approach for managing growth, by adding more servers and storage in a distributed computing environment to

# Who Pioneered Hyperscale IT and Why?

4

The hyperscale computing approach was pioneered by companies like Amazon, Facebook, and Google to support cloud environments at massive scale. These environments are typically comprised of multiple petabytes of storage and tens of thousands of servers, supporting a relatively limited number of applications used by hundreds of thousands (even millions) of concurrent customers.

A hyperscale IT architecture enables these organizations to provide value-added services for their customers in an efficient and cost-effective manner. The good news is that organizations don't have to be at the scale of web giants to take advantage of hyperscale IT. The same approach can be deployed at small scale and expand indefinitely to handle exponential growth in data, ease operation provisioning burdens, and meet ever-changing business requirements.

# Why Is Hyperscale IT Useful in the Modern Enterprise?

Today's competitive business landscape requires business innovation, agility, and flexibility. Business executives are under constant pressure to innovate, application developers must accelerate time to market, and IT infrastructure and DevOps teams must support a flexible infrastructure in order to deploy new application offerings rapidly (see Figure 1-1).



According to Forrester Research, enterprise data is growing at a rate ten times faster than storage budgets, and 58 percent of organizations take days, weeks, or months to provision storage. Only 14 percent of organizations have "cloudlike" rapid provisioning capabilities.

As a result, modern enterprises are caught between two worlds: the proverbial rock and a hard drive, er, place! (See Figure 1-2.)



These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

### 6

On the one hand, IT can't stop supporting the legacy environments that run the core of the business. The systems of record — such as enterprise resource planning (ERP), customer relationship management (CRM), and email — still need to function. These systems are key for employee productivity, but they still run on traditional, siloed compute and storage infrastructure.

At the same time, IT must find ways to support innovative new workloads, such as mobile applications, software as a service (SaaS) offerings, big data, and the Internet of Things (IoT).

As a result of this, IT has deployed islands of storage consisting of a hybrid mix of direct attached storage (DAS), network-attached storage (NAS), and storage area networks (SAN) throughout the data center. Data is arguably the most important business asset today and IT has to keep finding the right storage technologies to store the data for various traditional and new application workloads and environments (see Figure 1-3).



These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

Traditional storage arrays are sufficient for traditional workloads, such as server virtualization, virtual desktop infrastructure (VDI), backup and archive, and business continuity/disaster recovery.

Hyperconverged storage improves on traditional storage by adding simplicity and scale. For innovative new workloads such as production clouds, test and development clouds, and big data, hyperscale storage is the solution.



Hedvig addresses modern storage requirements for both traditional and new workloads with a single, software-defined storage platform. 8

## **Chapter 2**

## Assimilating Hyperscale Storage

### In This Chapter

- Learning about hyperscale storage
- Comparing traditional and hyperscale storage
- Recognizing the benefits of hyperscale storage

n this chapter, you learn about hyperscale storage, how it differs from traditional storage, and how it can help you support rapid data growth in your organization.

## Defining Hyperscale Storage

Hyperscale storage solutions decouple application compute and storage resources, enabling each to scale independently depending on your business needs performance or capacity, respectively (see Figure 2-1).

Hyperscale leverages commodity servers and a software-defined approach, scaling the resources needed for applications and storage separately. As storage needs grow, companies can add servers running software-defined storage (SDS) to the storage tier to expand capacity independently of the application tier. Data is automatically distributed across the entire cluster of storage servers as new nodes are added to the system.





Conversely, as performance needs grow, companies can add servers to increase compute power, independent of the storage tier.

With hyperscale, scale-out storage software installs on servers — referred to as *cluster nodes* — that network together to form a storage resource pool. The application tier, which may include both bare-metal and hypervisor-based systems, runs on its own server and accesses the storage pool via a software client or proxy — or directly via an application programming interface (API). Each tier scales independently.



Hyperscale is the architecture of choice for companies such as Facebook and Google that want the flexibility to meet fluctuating data demands in compute or storage capacity independently, at will.

### Traditional Storage versus Hyperscale

Traditional enterprise storage systems, such as storage area networks (SANs), are among the most expensive hardware in modern data centers, not only in terms of capital investment, but also because of rack space, power and cooling, and management overhead. Scalability in these systems is typically measured in terabytes and limited by the performance of the storage controllers.

By comparison, hyperscale storage uses low-cost, readily available hardware — standardized servers with disk and solid-state drives — and scales performance and capacity independently, based on the needs of the organization.

### Heading Off Exponential Data Growth

A hyperscale deployment is suitable for generalpurpose server virtualization, but because of its unique decoupled nature, it also supports nonvirtualized applications effectively. Hyperscale excels in supporting large and growing datasets. The architecture is ideal for big data applications and for underpinning cloud environments that use technologies like OpenStack and Docker containers.

A hyperscale approach to data center and cloud infrastructure offers a high level of elasticity by helping organizations rapidly respond to changing application and data storage needs. Hyperscale lowers total cost of ownership by taking advantage of low-cost hardware, enabling the flexibility to scale compute or storage resources as needed, and driving automation to reduce the amount of human interaction required to operate a data center — even at massive scale.

Hyperscale storage provides a number of benefits over traditional storage, including:

Increased performance: Due to the scale-out nature of both the storage nodes and the storage proxy, storage capacity and storage performance can scale independently. Additionally, there is no logical controller ceiling as with a traditional storage solution, so if you start with a three-node cluster, for example, you can seamlessly expand to a 1,000-node cluster with no negative performance implications. In fact, performance actually increases as the system scales. Intelligent use of accelerated storage in the form of solid-state drives (SSD), nonvolatile memory Express (NVMe), and random-access memory (RAM) allows you to accelerate workloads where they should be accelerated, without the need for expensive all-flash arrays when only 10 percent of your data typically needs to be accelerated.

- High-availability and data protection: Hyperscale storage is active/active/(\*n)/active across multiple sites. Each storage node is an access layer into the storage fabric. Because hyperscale storage is software-based, storage nodes can even be run in a public cloud, creating hybrid-cloud high availability. With a minimum recommended replication factor for production data of three, even in the event of a failure, you're still not at risk of data loss or corruption from a second failure.
- Lower costs: It may come as a surprise, but software-defined storage is partly defined by hardware! With a software-defined, hyperscale storage solution, you can choose the most cost-effective and strategically appropriate hardware for your needs. That could save you up to 70 percent of the acquisition cost of a traditional SAN or storage appliance.
- Scalability: This is a fundamental element of hyperscale storage. You can start with three nodes and 12 terabytes (TB) of storage, and seamlessly grow to hundreds of nodes and several petabytes of storage. And if you need to scale storage performance separate from storage capacity well, that's exactly what hyperscale storage is designed to do!
- Increased use of virtualization: Whether you want to support and power a VMware cluster, Microsoft Hyper-V, KVM, or some other virtualization platform, or you want to delve into Docker and containers, hyperscale is the way to go.

Flexibility: By abstracting the hardware layer with a software-defined solution, hyperscale storage provides maximum flexibility to provision storage that is tailored per-application, letting you choose the protocol, performance, features, location, and data protection characteristics that best serve each workload.

### Making the business case for a hyperscale approach

Refresh and capacity planning in a traditional storage environment typically require organizations to prepurchase capacity for a three- to five-year time horizon. The architecture uses a scale-up approach (add drive trays and/or additional frames) to support growth, and organizations must regularly refresh hardware and software. Perhaps most challenging (and risky), data must be migrated from the old system to the new system when storage hardware is refreshed.

A hyperscale approach enables organizations to buy capacity as needed, scale-out with additional nodes to specifically support performance or capacity needs, retire old storage nodes without business disruptions, and avoid risky data migrations.

Consider the following scenario comparing the cost of a traditional SAN to the cost of a hyperscale storage solution:

SAN

 10 petabytes of storage growing at a compound rate of 20 percent year over year (YoY) over 10 years.



These materials are @ 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

## Chapter 3

## Comparing Hyperconvergence and Hyperscale

### In This Chapter

- Getting to know hyperconverged and hyperscale
- Recognizing the differences between each approach

n this chapter, I compare hyperconverged and hyperscale storage approaches and their benefits.

### Examining the Common Components

Hyperscale and hyperconverged storage have a few things in common (not just their first five letters). For example, both approaches:

- ✓ Leverage software-defined storage (SDS).
- Abstract underlying hardware resources to drive utilization efficiency and reduce operating costs.
- ✓ Use *commodity-based* hardware platforms to lower capital costs.

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

### Recognizing the Benefits and Challenges of Both Approaches

18

A *hyperconverged* system integrates application compute and storage in a single unit designed to scale out by adding more units.

Hyperconverged storage typically leverages commodity "whitebox" x86 servers and integrates softwaredefined compute along with SDS. Compute and storage scale linearly as IT incorporates additional hyperconverged nodes. Data and workloads are auto-balanced to new nodes when they're added to the storage system.

Hyperconverged systems can be architected using software and a "bring your own server" approach, or companies can buy solutions that preassemble hardware and software.

Hyperconverged systems leverage a virtualization hypervisor and scale-out storage software. The storage intelligence — sometimes referred to as the *controller* and disk capacity are co-resident with the hypervisor on the same server or appliance. Hyperconverged systems link together into a cluster to form a single resource pool.

Hyperconverged systems are suited to a number of workloads, including general purpose server virtualization, remote office/branch office where IT staff support may be minimal (or nonexistent), virtual desktop infrastructure (VDI) where predictable linear compute and storage scaling are critical, and test/dev environments.

On the other hand, a *hyperscale* system decouples application compute and storage, so that each component can be scaled independently as required by the environment.

Hyperscale storage also leverages "whitebox" commodity servers and a software-defined approach, but scales resources needed for application compute and storage separately. You can add more storage proxies for application access and advanced features at the client-layer without adding more backend storage cluster nodes and vice versa. Over time, you also have the flexibility to decommission older storage nodes and add newer configurations with faster CPUs and higher capacity storage without disrupting business applications. No more forklift upgrades!

Like hyperconvergence, a hyperscale deployment is also suitable for general-purpose server virtualization, but because of its decoupled nature, it also supports nonvirtualized applications effectively. Hyperscale excels in supporting large and growing datasets. The architecture is ideal for big data applications, OpenStack clouds, and Docker container environments.

## Determining the right storage strategy for your enterprise: A hyperscale scorecard

Answer the following questions, add up your points, then turn the page upside down to see your results (not really, the results are at the end of the questionnaire!).

- 1. For every 250 employees in your organization, add one point.
- 2. For every 100 VMs in your data center, add two points.
- 3. For every 100 enterprise applications that your IT organization supports, add two points.
- 4. For every 100 terabytes of data in your organization, add one point.

#### Results

Tally up your points to see which storage strategy makes the most sense for your organization:

- 0 to 20 points: Stick with traditional storage (for now).
- 21 to 60 points: Consider hyperconverged storage.
- 61 or more points: Hyperscale is the way to go!

### **Chapter 4**

## Exploring Hyperscale Storage Use Cases

### In This Chapter

- Streamlining server virtualization
- Increasing agility with OpenStack
- Enabling DevOps with Docker
- Improving economics for backup and archive

n this chapter, you learn about some common hyperscale storage use cases in the enterprise, including virtualization, OpenStack, Docker, and backup and archive.

### Virtualization

Virtualization technology has transformed enterprise IT, increasing productivity, lowering costs, and improving business continuity. In a January 2015 survey conducted by Forrester Consulting, respondents listed support for virtualization and support for data-related applications as their top use cases for software-defined storage.

Server and desktop virtualization use cases each have particular storage requirements.

Server virtualization is a fundamental enabler of the modern, software-defined data center, ushering in an era of service-oriented cloud architectures and ondemand IT. With virtualization technology, enterprises can dynamically provision and manage thousands of application workloads dispersed across business units and geographies. But the scale of these deployments and the rate of data creation in virtualized data centers pushes the limits of traditional array-based storage infrastructure.

Additionally, organizations commonly run multihypervisor environments, for example, with a combination of VMware vSphere, Microsoft Hyper-V, KVM, and Citrix XenServer. To successfully manage the volume of data and applications, keep pace with the speed of business, and support multi-hypervisor environments, a new approach to storage is needed.



Additionally, storage is one of the biggest bottlenecks for virtual desktop infrastructure (VDI) environments. I/O operations per second (IOPS) during boot storms, cost per desktop, and a myriad of factors can doom a desktop virtualization project before it ever gets off the ground.

Hyperscale software-defined storage provides the capability to customize your storage deployment to fit the particular demands of virtualized data centers and hosted virtual desktops by taking advantage of highperformance flash storage options in modern servers, the cost economics of commodity infrastructure, and integrated data efficiency features.

## OpenStack

Today's enterprises need greater agility and flexibility to provision applications and infrastructure in order to capitalize on new business opportunities. OpenStack is an open source, scale-out cloud management platform. Originally conceived as a way for companies to build private "cloud storage" as an on-premises alternative to Amazon Web Services, OpenStack has evolved to provide a comprehensive set of software tools for building and orchestrating private and public cloud environments, helping organizations achieve faster time-to-market with infrastructure-as-a-service (IaaS) automation.

Organizations rely on a wide range of applications and need access to a full set of storage capabilities within OpenStack. Modern web and cloud applications add additional pressure requiring scale-out architectures to meet user demands. In the face of these challenges, traditional storage solutions quickly become a bottleneck.



Most enterprises that deploy OpenStack do so with a hyperscale storage solution. The software-defined hyperscale approach provides simplicity and flexibility for OpenStack initiatives. Optimal solutions enable capabilities needed for modern infrastructure including block and file storage protocols via Cinder, object storage via Swift, and support for private and public cloud infrastructures.

### Docker

Docker is a container-based, open source solution that enables organizations to rapidly develop and deploy applications in test/dev, production, and DevOps environments to deliver innovation and value to their customers.

One of the most common requirements for enterprises looking to use Docker containers in production environments is the capability to run stateful services like databases in containers.

Docker is designed to run distributed applications and *microservices* and requires a storage infrastructure that shares the same qualities. A hyperscale, software-defined storage solution is a perfect fit for applications running in containerized Docker environments. Hyperscale storage with Docker delivers:

- The capability to incrementally scale performance and capacity by adding nodes to an existing storage cluster
- Multiprotocol access to storage for maximum flexibility of support for container workloads
- Data services that can be assigned to individual containers to meet the performance, efficiency, and availability needs of the microservice
- Native representational state transfer (REST) application programming interfaces (APIs) that give DevOps engineers complete control

## **Backup and Archive**

The explosive growth of primary business data means there is more data to protect — and therefore even more data to store — taxing already stretched IT budgets. Enterprises have deployed purpose-built backup appliances that combine proprietary hardware and software to address the problem, but at the price of flexibility and at significant cost. These solutions create yet another island of storage and come with the same challenges and management overhead of SANs.

For backup, hyperscale software-defined storage scales out dynamically with commodity servers both on- and offsite, providing an efficient platform for secondary data for backup, archiving, business continuity, and disaster recovery. For these workloads, you can deploy hyperscale software-defined storage with a commodity hardware profile appropriate for high ingest-rate and large scale, long-term storage needs.

The flexibility of hyperscale storage lets you assign specific storage policies to storage provisioned for backup out of the same storage cluster serving other types of applications workloads.

Hyperscale storage doesn't change your backup process. The solution is plug-and-play with any backup application designed to backup to disk. You can keep all your existing policies and procedures while benefiting from the advantages of a commodity, scale-out approach to storage.

### LKAB adopts Hedvig hyperscale storage for easy scalability and predictable cost

Luossavaara-Kiirunavaara AB (LKAB) is a world-leading mining company and producer of processed iron ore for steelmaking based in Sweden. LKAB invests in expertise, facilities, R&D, and technology.

(continued)

#### (continued)

#### Challenges

LKAB generates data across a wide range of applications, including mining apps, transportation data, and general corporate data. LKAB is already experiencing a data deluge from its early Internet of Things (IoT) investments. Data storage is an area of great importance to the company — and a source of significant cost. Because of the value of its data and a long-term retention requirement, LKAB protects the information by backing it up on a frequent basis — requiring even more storage to provision and manage. Every few years as existing storage solutions reached end-of-life, the company was forced to move large amounts of data to a new storage solution and retire the older solution — including software and hardware.

The company also found it a big challenge when disk space ran out in traditional systems. It took a long time to order, receive, and install additional storage.

#### Solution

The company turned to a hyperscale, software-defined approach to storage using Hedvig to deliver a more flexible and cost-effective infrastructure.

LKAB now backs up data to the Hedvig Distributed Storage Platform that runs on commodity servers hosting solid-state flash drives and spinning hard drives. The solution offers them a dynamically scalable storage system that runs as a seamless, deduplicated backup-to-disk target. The company also expects to expand the solution to serve as primary storage to host seismic data recorded in the field by geophones, and to serve as the backend of its Microsoft Exchange email server. The hyperscale storage cluster spans servers in two data centers separated by a few kilometers. LKAB has configured the Hedvig software to replicate data between racks in a single data center as well as to a secondary site for disaster recovery protection.

To extend capacity and update the system in the future, the company will add or replace servers. The hyperscale approach with Hedvig is more predictable and has improved current operations while better equipping the company to manage future growth.

#### **Benefits**

- Lower and more predictable costs for data storage hardware, software, and support
- The simplicity of provisioning efficient, enterprise-class storage with a few mouse clicks
- The capability to scale storage as data grows without forklift upgrades or long delivery delays
- ✓ 60% data reduction, storing more data in less space
- Built-replication for multi-datacenter and cloud disaster recovery protection

These materials are  $\ensuremath{\mathbb{C}}$  2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

### **Chapter 5**

## Introducing the Hedvig Approach to Hyperscale

### In This Chapter

- Understanding the Hedvig architecture
- Driving efficiency and flexibility
- Provisioning storage
- Automating provisioning and policy creation

Virtualization, automation, and self-service are now cornerstones of the modern software-defined data center. Traditional approaches to storage don't fit this new paradigm. A new approach to storage is required to keep pace with the exponential growth of data in an efficient, flexible, scalable, and cost-effective enterprise storage infrastructure.

In this chapter, you learn about Hedvig's highly scalable software-defined storage (SDS) solution — the Hedvig Distributed Storage Platform.

## Hedvig Architecture Overview

Key technical differentiators of the Hedvig Distributed Storage Platform include:

- A distributed systems engine with intelligence at each layer and node
- Support for any compute at any scale
- Support for hyperscale and hyperconverged deployments
- Flexible features and provisioning with Virtual Disks for a "hand in glove" fit to applications
- Hybrid support for multisite/cloud environments

✓ Highly programmable

The Hedvig Distributed Storage Platform consists of three primary components (see Figure 5-1):

- Hedvig Storage Service: Installs on commodity x86 and ARM processor-based servers or in cloud environments to form an elastic storage cluster
- Hedvig Storage Proxy: Installs on application hosts via a virtual machine (VM), a Docker container, or as a standalone proxy server to present block, file, or object storage Virtual Disks provisioned on the storage cluster
- Hedvig Virtual Disk: The scalable abstraction layer for granular provisioning of enterprise storage resources and functions



Figure 5-1: The Hedvig Distributed Storage Platform architecture.

## Hedvig Storage Service

The primary component in the Hedvig architecture the Hedvig Storage Service — is Hedvig's patented distributed systems engine that drives simplicity, elasticity, and flexibility for storage provisioning, policies, and management.

The Hedvig Storage Service software transforms existing server and storage assets — including solid-state drive (SSD) and flash media and hard disk — into a fullfeatured elastic storage cluster. The software deploys on commodity servers in a private data center and in public clouds to create a single storage cluster that is inherently hybrid.

Granular selection of features empowers administrators to avoid the challenges and compromises of a "one size fits all" approach to storage.



The Hedvig Storage Service is Hedvig's distributed systems engine that operates as an optimized key value store and is responsible for writing data directly to the storage media. It

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

supports provisioning of block (iSCSI), file (NFS), and object (S3 and Swift) virtual disks and captures all random writes into the system, sequentially ordering them into a log structured format that flushes sequential writes to disk. This provides the capability to ingest data at a high rate, as well as optimize the disk utilization.

The Hedvig Storage Service consist of two primary processes:

- Data process: The data process is responsible for the layout of data on raw disk. Hedvig storage nodes create two units of partition: *storage pools* (a logical grouping of three disks in the node) and *containers* (contiguous chunks of data). Virtual Disks are divided into containers, each residing in a unique storage pool in a given storage node. Containers replicate based on the replication policy associated with the Virtual Disk.
- Metadata process: The metadata process is responsible for how and where data is written. Metadata also tracks all reads and guarantees all writes in the system, noting the container, storage pool, replica, and replica locations of all data. Metadata is a key component of the underlying storage cluster, but it is also cached by the Hedvig Storage Proxy, enabling metadata queries from the application host tier without traversing the network.

The Hedvig Storage Service uses a combination of synchronous and asynchronous replication to distribute and protect data across the cluster and supports unlimited active data centers in a single cluster with up to six copies of data. Writes are committed when a majority of the storage nodes respond. For example, for a Virtual Disk configured with a replication factor of three, two nodes must acknowledge (ACK) back as having successfully received data for a write to be considered successful and acknowledged to the Storage Proxy.

The Hedvig Storage Service gracefully manages disk failures by accessing data from other replicas across the cluster. The process is handled automatically upon disk failures performing a "wide-stripe" rebuild — a process that leverages all relevant nodes and disks for a fast rebuild. The rebuild process is a background process that happens without impact to primary input/ output (I/O). During node failures, reads and writes continue as usual from the remaining replicas.

## Hedvig Storage Proxy

The Hedvig Storage Proxy is a lightweight abstraction hosted in either a VM or a container at the application tier, or as software on a separate proxy host, providing storage access to each physical host. The Storage Proxy ensures the Hedvig platform operates in existing environments without requiring changes to hypervisors, guest VMs, operating systems, or applications, thereby preserving virtualization and storage administrator workflows and eliminating the need to adopt new processes and procedures to operate a distributed, elastic storage cluster.

The Storage Proxy acts as a protocol converter and broker. It presents Virtual Disks as locally mounted block storage or file shares and traps local I/O, converting traffic to the Hedvig remote procedure call (RPC) protocol for communication to the underlying storage cluster.



The Storage Proxy deploys as an open virtual appliance (OVA) package or Docker container for easy provisioning, and it can be managed by any existing virtualization tools and orchestration frameworks.

The Storage Proxy runs in user space and doesn't require modification of the kernel. As a result, it's an extremely flexible process that can run in any compute environment — bare metal, hypervisor, container, or a cloud — including VMware vSphere, Microsoft Hyper-V, KVM, OpenStack (via Cinder), Docker, and Citrix XenServer.

For block storage, the Hedvig Distributed Storage Platform presents a Virtual Disk via a Storage Proxy as a logical unit number (LUN). Access to the LUN as an Internet Small Computer System Interface (iSCSI) target is then given to a host on the same system. The Hedvig Virtual Disk is then a locally mounted disk with all the properties applied during Virtual Disk provisioning, such as compression, deduplication, and replication. After the Virtual Disk is in use, the Storage Proxy intercepts disk activity and relays it to the underlying cluster.

For file storage, a Virtual Disk is presented via the Storage Proxy as a network file system (NFS) export. The administrator exports the Virtual Disk (the equivalent of presenting the LUN) and grants access to specific Storage Proxies. The Storage Proxy acts as an NFS server that guest VMs can then utilize and browse.

For object storage, deployment of a Storage Proxy isn't necessary. Virtual Disks are presented as object buckets or containers and can be accessed directly via Amazon Simple Storage Service (S3), Swift, and representational state transfer (REST) application programming interface (API) calls.

The Storage Proxy also provides a client-side caching ability that delivers read caching by taking advantage of local SSDs or Peripheral Component Interconnect Express (PCIe) devices, and also reduces I/O across the cluster by deduplicating all data before transmitting over network links.

To protect against a single point of failure, the Storage Proxy installs as a high availability active/passive pair. This pair uses multipath I/O (MPIO) for seamless block (iSCSI) failover and IP failover for file (NFS). If one Storage Proxy instance is lost or interrupted, operations seamlessly fail over to the passive instance to maintain availability. This occurs with no intervention necessary by applications, administrators, or users.



During provisioning, the administrator can indicate that a host will use a clustered file system, such as VMware virtual machine file system (VMFS). This automatically sets internal configuration parameters to accommodate multiple hosts using the volume and ensures seamless failover when using VM migration to a secondary physical host running its own Hedvig Storage Proxy. During live VM migration, any necessary block and file storage "follows" guest VMs to any host.

## Hedvig Virtual Disk

The Virtual Disk is the fundamental abstraction unit of the Hedvig Distributed Storage Platform. Organizations can spin up any number of concurrent Virtual Disks each thinly provisioned and instantly available. The creation, provisioning, and removal of Virtual Disks can be accomplished in just a few seconds. Every storage feature such as compression, deduplication, and replication can be switched on or off to fit the specific needs of any given workload.

Virtual Disks are provisioned from the Hedvig graphical user interface (GUI), command line interface (CLI), or via API calls directly to the cluster. Virtual Disk provisioning can also be scripted and automated.



In addition to the provisioning with granular policy selection, users can perform data management tasks, including snapshots and clones on a per-Virtual Disk basis. Snapshots happen instantly due to the unique metadata architecture of the Hedvig platform and there is no limit to the number of snapshots or clones that can be created. Users can set new policies for Virtual Disk clones — for example, selecting a different replication factor and residence.

The wide range of storage features that can be set on a per-application basis with a Hedvig Virtual Disk include:

Protocol: block, file, or object

🖊 Size

- Residence (hard disk or exclusively flash)
- Replication factor (number of replicas)
- Replication policy (agnostic, across racks, across data centers)
- Client-side caching
- Inline global deduplication
- ✓ Inline compression

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

### **Chapter 6**

## Ten Reasons Hyperscale Storage for the Enterprise Is Logical

#### In This Chapter

- Deploying on any hardware for any environment
- Reducing complexity and cost
- Improving performance, availability, and resilience

Our pointy chinned friend (the "Dummies" man) has taken a cue from our pointy eared friend (Spock) to point out a few things to remember about hyperscale storage!

- Greater flexibility: Hyperscale storage is software-defined storage (SDS) that abstracts the hardware layer, lets you define storage as you need it, and enables you to use your choice of commodity infrastructure for storage.
- More virtualization and cloud options: Deploy storage for any type of environment, including different operating systems, hypervisors, containers, or clouds.

These materials are © 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

- Intelligent application-aware storage: Applicationaware storage ensures that individual virtual machines (VMs) and applications get exactly the storage resources they need, whether read/write I/O-intensive or low latency, for example.
- Simpler capacity planning: With a scale-out architecture, there are no complicated algorithms or formulas to determine how many shelves or controllers you need to add to your storage configuration. Simply add compute or storage nodes independently, based on your needs (performance or capacity).
- Massive CAPEX reductions: Hardware represents about 60 to 70 percent of the cost of a traditional storage area network (SAN). Hyperscale storage leverages commodity servers so that you can reduce those costs and eliminate the need for separate add-on solutions like replication.
- Significant OPEX savings: With its automation, self-healing architecture, and resiliency under different node/drive failure conditions, hyperscale storage drives significant OPEX savings.
- Shorter provisioning cycles: The simplicity of storage provisioning, with an intuitive, Amazonlike (AWS) interface rather than more complex provisioning constructs, reduces storage management overhead and accelerates automated, selfservice storage provisioning to support rapidly changing business requirements. When you need storage, you get it with a few clicks (of a mouse, not your heels) — not in a few days!

- Improved performance: As you add nodes, the overall performance of the entire storage system increases. With an independent, scale-out architecture, you can add the exact performance that your unique needs dictate to address any particular performance issue.
- Better availability: Survive disk, node, rack, and data center failures. Your data center applications are automatically directed to data copies, alternate data centers, or public clouds — your critical business applications don't need to fail just because the primary copy of data fails.
- Better data protection: Hyperscale storage leverages modern technologies and techniques like self-healing, and multidata center and hybrid cloud replication for more reliable, hands-free, automated data protection.

### 40

These materials are @ 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

## Glossary

**Amazon Simple Storage Service (S3):** An online file storage web service designed for data backup and archiving.

**ARM processor:** A family of CPUs based on the RISC (reduced instruction set computer) architecture developed by Advanced RISC Machines.

**block storage:** Data storage that is organized as a sequence of bits or bytes of a fixed length or size.

**Cinder:** The block-based storage component of OpenStack.

**cluster:** A group of server nodes that form a single scale-out resource pool.

**container:** Operating-system-level virtualization where the kernel of an operating system allows for multiple isolated user-space instances, instead of just one.

**Docker:** An open platform for developers and sysadmins to build, ship, and run distributed computer applications inside software containers.

**file storage:** Data storage that is organized in a hierarchical file system and uses nodes to manage access control and file permissions.

**hyperconverged:** A software-defined storage deployment that scales storage nodes in lockstep with application hosts. **hyperscale:** A software-defined storage deployment that scales storage nodes independent from application hosts.

**Internet Small Computer System Interface (iSCSI):** A storage networking protocol that works on top of the Transport Control Protocol (TCP) and allows SCSI commands to be sent over local-area networks (LANs), wide-area networks (WANs), or the Internet.

**key value store:** A data storage paradigm designed for storing, retrieving, and managing records using a key that uniquely identifies the record, and is used to quickly find the data within the database.

**Network File System (NFS):** A distributed file system protocol that allows a client computer to access files over a computer network like local storage.

**nonvolatile memory Express (NVMe):** A logical device interface specification for accessing nonvolative storage media attached via PCI Express (PCIe) bus.

NVMe: See nonvolatile memory Express.

**object storage:** Data storage that uses variable-sized data containers organized into a flat address space.

**OpenStack:** An open source cloud computing project.

**Proxy:** An entity (a computer system or an application) that acts as an intermediary for requests from clients seeking resources from other servers.

**S3:** See Amazon Simple Storage Service.

SCSI: See Small Computer System Interface (SCSI).

SDS: See software-defined storage.

**Small Computer System Interface (SCSI):** A set of standards for physically connecting and transferring data between computers and peripheral devices.

**software-defined storage (SDS):** A software-based implementation of storage services that can run on any commodity hardware.

**solid-state drive (SSD):** A nonvolatile storage device that stores persistent data on solid-state flash memory.

storage cluster: See cluster.

is strictly prohibited.

**Swift:** The object-based storage component of OpenStack.

### 44

These materials are @ 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

These materials are @ 2016 John Wiley & Sons, Inc. Any dissemination, distribution, or unauthorized use is strictly prohibited.

### **Build enterprise** storage with modern distributed systems

Virtualization, automation, and self-service are cornerstones of the software-defined data center. Hyperscale storage fits this new paradigm with an agile infrastructure to stay ahead of exponential data growth.

- Connect to any compute, run in any cloud — build flexible elastic clusters using any commodity server
- Provide block, file, or object eliminate islands of storage with a unified, all-software storage platform
- Scale to petabytes of data start with a few nodes and scale to thousands, adding performance and capacity as needed

Lawrence C. Miller has worked in information technology in various industries for over 25 years. He has written more than 75 other For Dummies books on numerous technology topics.



A Wiley Brand

### **Open the book** and find:

- How hyperscale IT evolved and why it's needed in the modern enterprise
- When to deploy hyperconverged or hyperscale storage
- Why hyperscale is the best storage approach for a wide variety of workloads, including virtualization, OpenStack, Docker, and backup

Go to Dummies.com<sup>®</sup> for videos, step-by-step examples, how-to articles, or to shop!

> ISBN: 978-1-119-29272-2 **Book not for resale**

### WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.