



# OpenStack<sup>®</sup> Technology Breaking the Enterprise Barrier

HP Helion OpenStack<sup>®</sup>



Jacek Artymiak and Lisa-Marie Namphy





# **OpenStack® Technology**

## **Breaking the Enterprise Barrier**

HP Helion OpenStack®



OpenStack Technology  
Breaking the Enterprise Barrier

© 2014 Hewlett-Packard Development Company, L.P.

Published by:

HP Press  
660 4th Street, #802  
San Francisco, CA 94107

All rights reserved. No part of this book may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or by any information storage and retrieval system, without written permission from the publisher, except for the inclusion of brief quotations in a review.

#### WARNING AND DISCLAIMER

This book is designed to provide information about OpenStack cloud computing and HP Helion OpenStack technology. Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information is provided on an “as is” basis. The author, HP Press, and Hewlett-Packard Development Company, L.P., shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or from the use of the discs or programs that may accompany it. The opinions expressed in this book belong to the author and are not necessarily those of Hewlett-Packard Development Company, L.P. Readers should be aware that Internet websites offered as citations and/or sources for further information may have changed or disappeared between the time this is written and when it is read.

#### TRADEMARK ACKNOWLEDGEMENTS

All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. HP Press or Hewlett Packard Inc. cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark. The OpenStack Word Mark and the Square O Design, together or apart, are trademarks or registered trademarks marks of OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation’s permission. HP is not affiliated with, endorsed by or sponsored by the OpenStack Foundation or the OpenStack community.

## FEEDBACK INFORMATION

At HP Press, our goal is to create in-depth technical books of the best quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the expertise of members from the professional technical community. Readers' feedback is a continuation of the process. If you have any comments regarding how we could improve the quality of this book, or otherwise alter it to better suit your needs, you can contact us through email at [feedback@hpressbooks.com](mailto:feedback@hpressbooks.com). Please make sure to include the book title in your message.

We appreciate your feedback.



HP HEADQUARTERS  
Hewlett-Packard Company  
3000 Hanover Street  
Palo Alto, CA  
94304  
USA

Phone: (+1) 650-857-1501

Fax: (+1) 650-857-5518

**HP, COMPAQ** and any other product or service name or slogan or logo contained in the HP Press publications or website are trademarks of HP and its suppliers or licensors and may not be copied, imitated, or used, in whole or in part, without the prior written permission of HP or the applicable trademark holder. Ownership of all such trademarks and the goodwill associated therewith remains with HP or the applicable trademark holder. Without limiting the generality of the foregoing: a. Microsoft, Windows and Windows Vista are either US registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries; and b. Celeron, Celeron Inside, Centrino, Centrino Inside, Core Inside, Intel, Intel Logo, Intel Atom, Intel Atom Inside, Intel Core, Intel Core Inside, Intel Inside Logo, Intel Viiv, Intel vPro, Itanium, Itanium Inside, Pentium, Pentium Inside, ViiV Inside, vPro Inside, Xeon, and Xeon Inside are trademarks of Intel Corporation in the U.S. and other countries.

# Contents

Chapter 1	<b>Introduction</b> .....	1
	Why we wrote this book .....	2
	Who this book is for .....	2
	What you will learn.....	2
Chapter 2	<b>A case for cloud computing</b> .....	5
	The new reality of IT .....	6
	We are drowning in data.....	7
	People adopt more and more mobile devices .....	9
	We expect security and accountability .....	10
	Cloud becomes the storage medium of choice.....	10
	Cloud computing to the rescue.....	11
	The world around us is getting smarter .....	11
	We expect our data centers to shrink and expand on demand .....	11
	Mobile apps are hungry for bandwidth.....	13
	We want to know more.....	13
	What do we do now? .....	15
	What the cloud needs to deliver .....	16
Chapter 3	<b>Why we need open source clouds</b> .....	17
	Hybrid delivery.....	18
	Open standards are a good thing .....	19
	An operating system for the cloud .....	20
	Why HP chose OpenStack technology .....	20
	No love for proprietary clouds .....	21
	Open = Trusted.....	21
	Maturity.....	23
	Appeal to our own people .....	23
	Dynamic growth.....	24
	Industry-wide support.....	25
	Global community .....	25
	Sound governance model .....	26
	An easy decision to make .....	29

Chapter 4

**Just how serious is HP about OpenStack technology?** ..... 31

HP's contributions to the OpenStack project..... 32

- Involved from the beginning ..... 32
- Project governance ..... 32
- Financial support..... 33
- Strategic commitment..... 33
- Commitment to the project governance and processes..... 34
- HP is always one of the top five code contributors to OpenStack..... 34
- Dedicated staff..... 35
- Successful commercial deployments of OpenStack technology ..... 35
- Transparency and accountability..... 36

Expert leadership ..... 36

Coopetition and giving back to the community ..... 36

HP Helion ..... 37

- An all-hands effort ..... 37
- What is HP Helion? ..... 38
- It's a journey ..... 38

Chapter 5

**What is HP Helion OpenStack?** ..... 39

Why we created HP Helion OpenStack..... 40

Main advantages of HP Helion OpenStack ..... 41

- Simple installation ..... 41
- Hardened code ..... 42
- Improved stability ..... 42
- Faster improvements..... 42
- Well-paced innovation ..... 43
- Simplified management ..... 43
- Automated delivery of additional functionality and content ..... 43
- Interoperability ..... 44
- Good fit for all data centers, large and small ..... 44
- Simple data sovereignty ..... 44
- Lower cost of migration ..... 45
- A truly open cloud operating system ..... 45
- Model once, deploy everywhere ..... 46
- Rapid provisioning..... 46
- Sharing resources..... 46
- Multi-cloud provider support ..... 47
- Hybrid provisioning ..... 47
- Supported by HP ..... 47

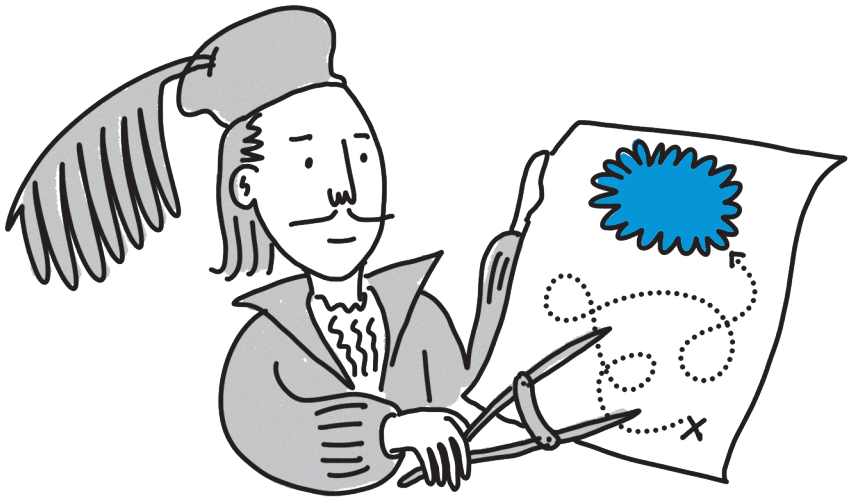
HP Helion foundational technology..... 48

Partner and ISV support ..... 48

Chapter 6	<b>Inside HP Helion OpenStack</b> .....	49
	The architecture of HP Helion OpenStack .....	50
	The HP Helion OpenStack kernel .....	52
	HP Helion OpenStack value-added services .....	58
	Administering HP Helion OpenStack .....	60
Chapter 7	<b>Use cases</b> .....	63
Chapter 8	<b>Where do I go from here?</b> .....	67
	Online resources .....	68

Chapter 1

# Introduction



---

## In this chapter

- ✓ Why we wrote this book
  - ✓ Who this book is for
  - ✓ What you will learn
-



## Why we wrote this book

Welcome!

We wrote this book to introduce HP's views on the future of data center models, cloud computing, and OpenStack technology, and to introduce our foundational platform, which is based on OpenStack software and provides a common architecture for hybrid delivery across the HP Helion portfolio.

This is the story of our cloud journey, an explanation of the choices we made, and an invitation to join us on that journey into the future. We explain why we chose OpenStack technology, and how we are going to support our clients on their own *cloud computing* journey.

## Who this book is for

If you are a CTO, data center administrator, systems architect, or an IT professional-looking for an enterprise-grade, hybrid delivery, cloud computing solution that is open, trusted, and reliable, we wrote this book for you. It is also for (and dedicated to) all you pioneers who fearlessly pushed the needle forward and moved the industry to the next major phase of its technology lifecycle. We had a great time producing this book and we hope you will enjoy reading it.


## What you will learn

This book explains how cloud computing is a solution to the problems facing data centers today and to highlight the cutting-edge technology (including OpenStack cloud computing) that HP helped bring to the current stage. Once you more fully understand this relatively new technical approach, we believe you too will see that it's the right choice for the future of cloud computing.

In the following pages you will learn about the challenges facing data centers today, why cloud computing is the best technology we have to cope with those challenges, and why OpenStack technology provides the best cloud computing operating system for your data center. You will learn how deeply involved in OpenStack cloud computing we are, what cloud computing means for the future of HP, and how we built HP Helion OpenStack—an enterprise-grade OpenStack distribution ready to help you realize your hybrid cloud delivery needs.



# A case for cloud computing



---

## In this chapter

- ✓ The new reality of IT
  - ✓ The need for a new data center model
  - ✓ A case for cloud computing
-

## The new reality of IT

Are you ready? The way you deploy, manage, and use your IT infrastructure is about to change.

Cloud computing is about to take over data centers around the world and dispel the old ways of thinking about IT infrastructure. And we are not talking about taking over the odd surviving mainframes, but the data centers designed to meet the demands of the old (yes, already old!) client-server model.

Don't believe it? Remember the mighty mainframes that ruled the data processing world and stared down at the tiny UNIX boxes? They are gone now. Client-server data centers are on the same path.

### Where are they now?

Mainframes turned out to be very similar to dinosaurs. Too big, too slow, not capable of adapting to the changing world, and way too expensive to buy and run. Buying a mainframe was like building a power plant—only governments and banks could afford that kind of investment. Thirty years ago it seemed like they would be with us forever. And yet, they are practically extinct today.

A feisty, young challenger—the client-server model based on UNIX, sealed the fate of the mainframe. It looked like a toy at first, but then Oracle, Sybase, and Informix showed the world that the client-server architecture was a much more affordable business proposition. So good, in fact, that the architecture made it possible to build numerous Internet-based businesses. It drove demand for new, bigger data centers all over the world. UNIX was in and mainframes were out.

Today, only a few universities bother to teach courses that cover the mainframe architecture, and those that do it spare mainframes some time only because they consider that knowledge essential to understanding the origins of the modern data center. Much like the dinosaur—mainframes are long gone, but studying their fossils helps us better understand other living organisms.

Such is life...

The old adage “nothing is constant except change” could not be more apt than when applied to the world of IT. Today, the world of IT must deal with one of the greatest shifts the data center business has ever seen. The highly dynamic, unpredictable, and rapidly growing nature of demand for compute, networking, and storage created a need for this dramatic change.

**Cloud computing is about to take over data centers around the world.**

The simple fact is the traditional data center architectures no longer deal effectively with the new reality of IT—a fact that is beginning to dawn on data center facilities’ providers and users worldwide. Four trends define the new style of IT: *Big Data*, *mobility*, *security*, and *cloud*.

## We are drowning in data

Pundits continue to argue about what Big Data is exactly. Meanwhile, every day we pretend to cope with this data tsunami of our own making.

The days of people performing data entry at myriads of PCs is long gone. No human army could keep up with all the data collected these days.

Smart phones, smart buildings, smart appliances, smart cities, smart cars, smart roads, combined with country-wide CCTV networks, traffic monitoring networks, environmental monitoring networks, industrial monitoring networks—these technologies barely existed five years ago. Today, all of these and more run untold numbers of data-producing and data-crunching applications all day, every day.

Each new device or app generates gigabytes—even terabytes—of data per day. All that data needs its own backend capable of sending, requesting, and processing information on a massive scale, stretching the limits of the old hardware, software, and data centers. Quite frankly, we need to be planning for the *exabyte age*, because that future is rapidly unfolding before our eyes. Check out these statistics published in the first half of 2013:

- Scientists at CERN run experiments that generate over one *petabyte* of data per second during Large Hadron Collider experiments.<sup>1</sup>

Currently, most of that data has to be discarded, because they cannot store and process it all for technical and financial reasons. Even with international funding, it is too expensive to pay for a massive, purpose-built IT infrastructure capable of capturing and processing vast amounts of data generated in a split second.

- Twitter delivers over 500 million tweets per day.<sup>2</sup>

If every tweet is roughly a 4KB JavaScript Object Notation (JSON) object, it means that Twitter has to handle over 2TB worth of data every day. All that data needs to be stored, delivered, and analyzed. It cannot be lost and it must be available for retrieval in the future.

The 2TB estimate does not include images and other forms of media Twitter stores and delivers using its infrastructure. It also does not include the CPU cycles and network resources necessary to create and resolve links using the Twitter URL shortener. And let us not forget that Twitter does it all in real time, without dropping data even when a global event causes a sudden spike in demand for its resources.

- Instagram captures, stores, and delivers over a billion likes per day.<sup>3</sup>

Instagram likes are a small, but very important subset of the data that Instagram's backend has to handle with great care and speed. Each like is a simple message represented by a JSON object, possibly under 1KB. But even if it is that tiny, Instagram is still dealing with at least 1TB of like data per day. That data carries important information about who likes what, when, and on what platform the like was created and delivered on.

- Tumblr stores 3TB of new data per day.<sup>4</sup>

Images, videos, audio files, blog posts, shares, likes, you name it. All that data has to be captured, stored, analyzed, and served, but because of its highly visual nature, the infrastructure Tumblr needs is even more sophisticated than that of Twitter.

These all represent some of the more extreme cases, yet any mobile application or online service that becomes popular faces similar problems of availability, capacity, and scalability—usually when they are least prepared for it. We see this happen when someone discovers something trendy and adoption spreads like wildfire. Or, the opposite extreme also happens when the next big thing means that old favorites get abandoned and lots of expensive servers sit idle.

Each new device or app generates gigabytes—even terabytes—of data per day.

## People adopt more and more mobile devices

With feature phones being replaced by smart phones, and personal computers being replaced by tablets, the use of always-on, always-connected devices is growing exponentially worldwide. People often carry a phone or two plus a tablet and maybe even a smartwatch.

These billions of devices need to be connected 24/7 making capacity planning a nightmare at best. No one can predict how many devices will be turned on and, once activated, when and which apps users will install.

The majority of smart devices are designed to communicate with the home base; they check for and download software updates, send usage data, or simply perform their job taking temperature, recording movement, or uploading holiday snaps.

It is impossible to predict the amount of data they will generate or the Internet traffic ebbs and flows they may produce. Still, when it comes to those who buy our devices and use our apps, this unpredictability is not an acceptable excuse for poor service.

Despite all attempts at prognostication, nobody knows which device is going to get popular, when, and what will be its rate of adoption. Therefore, backend capacity planning is an impossible task. Still, no excuse—your data center must cope with those demands, whether short-lived or ongoing, or your customers will dump your products and services for the competition.



## We expect security and accountability

You have seen the headlines telling tales of security breaches and the risk of identity theft. Every business and person knows that security, trust, transparency, and accountability are crucial in every context, but especially with regards to digital data. Everyone wants to know who has access to data and how we can preserve its integrity, and sovereignty. Both security and accountability encompass legal and political issues of strategic importance to businesses, governments, and ordinary people.

## Cloud becomes the storage medium of choice

Cloud storage has actually been around for more than half a decade and is being used by hundreds of millions of people and companies to store their documents, pictures, videos, life streams, metadata, etc. Smart devices come with a choice of two or more cloud storage accounts. The masses and the businesses have voted with their dollars and attention, preferring not to

Despite all attempts at prognostication, nobody knows which device is going to get popular, when, and what will be its rate of adoption. Therefore, backend capacity planning is an impossible task.

manage different file formats and storage devices: instead, their computers, smart phones, and tablets do it for them.

The demand for cloud storage is only going to continue to grow as we

connect every device capable of running a TCP/IP communication stack to the Internet. Cloud storage only paints the first part of the picture, however. Using the cloud for compute functions means the continual adoption and expansion of cloud platforms.

Equally as important as cloud storage, are cloud-based applications: from office suites to image processing to other kinds of applications not yet envisioned. Apps are pushing the adoption of cloud technology. With the mammoth amounts of data and proliferation of devices using apps in multiple formats and on multiple platforms such as binaries, HTML5, or emulated apps, cloud emerges as the only smart choice to deliver what businesses and individuals expect.

## Cloud computing to the rescue

While Big Data, mobility, security, and cloud comprise current IT trends, it is obvious that cloud also encompasses a way to deal with the other three. The challenges posed by these trends have one thing in common—they are, by nature, highly dynamic and unpredictable. It is hard to predict how much capacity needs to be provisioned when some obscure mobile app suddenly goes viral gaining tens of thousands of users overnight.

Let us take a closer look at some of these challenges.

## The world around us is getting smarter

Although not all of us are early adopters picking up the latest and greatest smart device, worldwide more and more people use an increasingly large number of devices equipped with sensors that gather and relay data about their own state, users, location, environment, the apps they are running, and the usage patterns.

Many of these smart sensors are in reality tiny computers capable of running the TCP/IP communication stack suite as well as being capable of communicating using HTTP and Secure Sockets Layer (SSL). By their function, then, each sensor automatically becomes a web client in need of servers to talk to. Large networks of smart sensors enable the creation of smart cities, smart cars, smart roads, smart buildings, smart appliances, smart quantified self-monitors, and a plethora of smart environmental sensors to form what is broadly called the *Internet of Things*.

## We expect our data centers to shrink and expand on demand

We continue to process more and more information and increasingly rely on computers and networks in all aspects of our lives, thus increasing the demand for data processing power, storage, and networking—an incalculable and inconstant demand.

From small startups to big organizations like banks or retailers, the need for data processing capability varies dynamically in ways that a traditional data center cannot deal with in a cost-efficient manner. The only approach

available to these administrators is to design and build their data centers to handle the worst-case scenarios, a very costly option for a traditional brick-and-mortar data center model.

It may help to think of capacity planning in a different way.

### **What the Dutch can teach us about capacity planning**

If you know the variables influencing the demand for processing power, bandwidth, and storage capacity, you can handle capacity planning with confidence. When that demand is unknown, the best you can do is plan for the worst-case scenario.

Consider how the Dutch must plan for the ever-present threat of flooding from the North Sea. Years of protecting their reclaimed land have shown that the water level can vary from zero to ten meters. To deal with the possible variance, it makes no sense to just take the average and build a dam five meters tall—it must be at least ten meters whether water levels rise or not.

The same principle applies to data centers. If you need 1,000 servers to handle a spike in the demand for processing power on just one day in a month, you still need to keep those servers humming in their racks 365 days a year, even if that capacity is utilized only on payday.

Building and maintaining a data center capable of handling spikes in demand—even predictable ones—that happen on fewer than 20 days per year is not only expensive, but it is also wasteful; it requires an enormous amount of capital to design, build, and keep the data centers running.

The new data center has to adapt to the changes in the demand for processing power and has to be able to do it within minutes. A sudden spike in the amount of data sent by billions of devices may occur instantly, without warning as well as be over just as quickly—delivering and configuring servers to deal with a sudden spike makes no sense and inevitably comes too late. The old data center model offers no answer to these questions, making businesses less competitive and less nimble by tying up a lot of capital resources.

## Mobile apps are hungry for bandwidth

One of the reasons these spikes happen: mobile devices run a variety of applications that are expected to always work. Earlier we discussed the growing adoption of all sorts of mobile devices. Now let us consider the bandwidth capacity this growth requires.

Many applications require daily update. The infrastructure that handles those updates must be able to cope with the challenges of distributing the right payloads to the right clients over unstable connections that offer high latency and very limited bandwidth.

To better understand the magnitude of the problem, realize that an update to an app with a small 100KB binary and one million users generates a bandwidth bill for over 90GB. When Apple ships a new version of iOS, the bandwidth requirement is north of 600MB.

Email, Twitter, Facebook, conferencing apps, and calendars are just a few applications that have to be updated, sometimes multiple times per day. And you never know and cannot plan for a sudden spike in demand for bandwidth, storage, and computing power when millions of people want to find out or comment on the latest news or entertainment phenomenon.

## We want to know more

The billions of connected devices and apps cause a massive headache for anyone charged with capacity planning—all those always-on, always-syncing sources of data generating output in a variety of often unstructured formats, all expecting real-time communication with the home base.

Those communications no longer center around uploads, downloads, and access control. Smart devices today “speak” with the home base and with each other. They stream audio and video, and process more data than desktop PCs used to plough through five years ago.

And it is not just business apps, games, or social media services that need IT infrastructures that can facilitate communications in real time. Even something as apparently simple as a shopping list app requires a reliable, secure backend capable of delivering and handling real-time updates.

Beyond storing and retrieving data, we now search for meaningful patterns, trends, sentiments, and early warning signals inside vast quantities of data. We used to want to store and retrieve digital still images. Today, we want to run them through image recognition software looking for faces, brand logos, and more.

We used to be happy to have the ability to chat using plain text. Today, we wrap what we are saying online in JSON or XML structures and mix it with images, videos, and audio files to search for actionable data points. And now, we look for clues that let us find out more about the people behind the messages. We want to know more about the location of the event captured on a digital still. We want to monitor sentiment, approximate gender, annual

income, and other data points. And we want to do it in under a second.

Beyond storing and retrieving data, we now search for meaningful patterns, trends, sentiments, and early warning signals inside vast quantities of data.

As we get better at devising more accurate data analysis algorithms, we often want to reassess old

information along with the new, creating a temporary demand for additional compute, network, and storage resources.

The myriads of types of data formats and analytical algorithms that dig through the mountains of data present a huge challenge to the application architects who have to use new tools, new algorithms, and vast amounts of central processing unit (CPU) power, storage space, and network bandwidth to keep up with the ebbs and flows of the tidal waves of data and the demand for immediate answers.

With this explosion in the amount and the variety of data, gone is our ability to reliably predict changes in demand for compute power, storage capacity, or network bandwidth.

Data centers of today have to deal with a deluge of information coming from an ever-increasing number of new sources capable of generating gigabytes, often terabytes of data per day as we cover our world with one network after another, each smarter than the previous one, each capable of gathering more information than ever before.

## What do we do now?

The traditional data center model does not have enough flexibility to deliver the resources necessary to handle such unpredictable spikes in demand. Even the best logistics will not help if you need a thousand servers delivered, configured, tested and put into production overnight. If somehow you could manage to stand up a new batch of servers overnight, once the spike ends and the demand ceases, their supplier will not take them back from you; the time and money is gone.

A better solution: affordable cloud computing infrastructure. Thanks to the hard work of the people involved in hardware virtualization and cloud computing technologies, we are now at a point in time when it is possible to launch hundreds or thousands of servers with a single script or an API call. Another call can decommission them in a few minutes, all done without having to re-wire physical switches, servers, and disk arrays.

If you are ready to address the data and usage problems facing all of us now, consider the following advantages cloud computing offers:

- Virtual hardware can be provisioned and decommissioned within minutes even on a weekend (when the server suppliers are closed and the data centers operate with minimal staff). This feat is impossible to achieve with bare metal servers. Nobody delivers and deploys real hardware servers as fast as you can deploy virtual servers, with the click of a mouse button or a call to a RESTful API.
- When you no longer need the cloud resources you ordered, all you have to do is delete them and you are no longer paying for them. Cloud computing replaces capital expenditure with operational expenditure, making operating your own data center much more affordable.
- You can create large-scale backends that exist and operate for short periods of time. Coming back to the Large Hadron Collider experiments mentioned earlier, the problem of capturing, storing, and processing data from scientific experiments can be solved using temporary, on-demand compute resources. Or consider how banks benefit from on-demand storage and compute resources to lower the cost of operating their own data centers when they need to deploy extra capacity on paydays or bond settlement days.

## What the cloud needs to deliver

While having the ability to add compute, networking, or storage on demand is very appealing, cloud computing platforms must change to provide openness, interoperability accountability, and the greatest promise of cloud computing—hybrid delivery.

Even the best cloud computing platforms today do not interoperate very well. They are warehouses of black boxes of goods with varied quality. Their

Cloud computing platforms must change to provide openness, interoperability accountability, and the greatest promise of cloud computing—hybrid delivery.

proprietary nature makes it very costly to move infrastructure from one public cloud to another or from a private cloud to a public cloud. This has to change if we want to reap the


most attractive benefits of the cloud. And it has changed, as you will learn on the following pages. Read on to find out where the cloud innovators are taking us.

## REFERENCES

- 1 CERN.  
[home.web.cern.ch/about/updates/2013/04/animation-shows-lhc-data-processing](http://home.web.cern.ch/about/updates/2013/04/animation-shows-lhc-data-processing)
- 2 Securities and Exchange Commission.  
[www.sec.gov/Archives/edgar/data/1418091/000119312513390321/d564001ds1.htm](http://www.sec.gov/Archives/edgar/data/1418091/000119312513390321/d564001ds1.htm)
- 3 Digital Market Ramblings.  
[expandedramblings.com/index.php/important-instagram-stats/](http://expandedramblings.com/index.php/important-instagram-stats/)
- 4 High Scalability.  
[highscalability.com/blog/2012/2/13/tumblr-architecture-15-billion-page-views-a-month-and-harder.html](http://highscalability.com/blog/2012/2/13/tumblr-architecture-15-billion-page-views-a-month-and-harder.html)

## Chapter 3

# Why we need open source clouds



---

### In this chapter

- ✓ Delivering on the greatest promise of cloud computing—hybrid delivery
  - ✓ The need for open standards in the cloud
  - ✓ The need for an operating system for the cloud
  - ✓ The question of trust
  - ✓ What the industry thinks of the OpenStack project
  - ✓ A vibrant community
  - ✓ A sound governance model
  - ✓ The many reason for HP's involvement in the OpenStack community
-



Cloud computing is a very compelling concept, but the everyday reality is unapologetically practical and leaves little or no time to ponder upon the fascinating concepts of compute, networking, or storage abstraction.

After working on our own in-house cloud computing projects for a while, HP saw great promise in the OpenStack cloud computing project, and in 2011 decided that it was time for a change in thinking about the cloud—from it being just one of many products and services in our portfolio to treating it

as *the* platform for delivery of future products and services.

We will go into more details later, but first it may be helpful for you to know what OpenStack

Cloud computing will never deliver what it promises unless we make hybrid delivery a reality, and there is an acute need for an operating system for the cloud.

technology is. From [openstack.org](http://openstack.org), “OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter, all managed through a dashboard that gives administrators control while empowering their users to provision resources through a web interface.”

What we needed was a technology platform that would help us realize the full potential of cloud computing in the enterprise. Two of the most important things we learned from talking to our customers were the realizations that cloud computing will never deliver what it promises unless we make *hybrid delivery* a reality, and that there is an acute need for an operating system for the cloud.

## Hybrid delivery

Nobody can build and maintain enough data centers to adapt to the changes in demand for data processing resources. As our customers began to investigate the cloud as a possible solution the following models of cloud delivery emerged:

- **Private cloud**—Owned and managed by the customer.
- **Managed cloud**—Owned by the customer, but managed by a third party.
- **Public cloud**—Cloud as a service, so the customer only pays for the resources they used.

Most customers who use the cloud to adapt to the new reality of IT use a mixture of private, public, and managed clouds from different providers.

One of the greatest impediments to a wider adoption of cloud computing in the enterprise is the complexity of redeploying the cloud-based infrastructure on a different cloud platform from a different provider.

What we need is a solution that enables hybrid delivery so moving data or applications between private, managed, and public clouds becomes relatively simple. These truly dynamic clouds would make it possible to provision extra resources on a public cloud when the private cloud cannot cope with a spike in demand. The way things are at the moment, bursting into another cloud is only possible within the realm of the same provider, which does not satisfy the requirement for hybrid delivery.

## Open standards are a good thing

When the world of IT adopts a standard, formal or informal, everybody moves forward faster. We have seen numerous examples of this approach working in practice, such as the standard C library, POSIX, TCP/IP, or HTTP.

We need the same thing to happen in the world of cloud computing, because open standards in the cloud will:

- Accelerate the process of migration from the old data centers to the cloud.
- Allow customers to compare different clouds and choose one that meets their needs.
- Allow providers to compete in a more transparent environment.
- Create a rich ecosystem of vendors providing value-add enhancements to the core platform, without breaking compatibility with the core and with other add-ons.
- Enable interoperability between public and private clouds.

## An operating system for the cloud

When we begin to think of data centers as pools of compute, networking, and storage resources, they start to resemble the basic design of a computer, albeit a huge one. And any computer, whether a tiny embedded device, a powerful server, or a smartphone in our pocket, needs an operating system.

The more open and sophisticated the operating system, the better the hardware becomes; the more services it provides, the better the applications built on top of it become. A good operating system is like a coral reef supporting a complex ecosystem. A lot of the advances in the history of computing would not have happened were it not for the free and open BSD distributions, Linux, Android, and the like.

What we need now is the equivalent of Linux for the cloud. It must support a variety of hardware used in data centers and expose the data centers' capabilities via simple, standard APIs. Only then will it become interesting for developers to build apps, services, and additional content that enhance the platform.

Examples of such add-ons for a cloud operating system include drivers for new storage hardware, application monitoring tools, energy efficiency monitoring, pre-built virtual images, pre-built infrastructure templates for ready-made mobile app backends, and more.

## Why HP chose OpenStack technology

In our search for a cloud operating system, HP considered all options available in 2011, proprietary and open. As we went through the list of our requirements, it became increasingly apparent that the most important question asked by our customers was the one of trust. It may have been disguised as a question about reliability, security, scalability, compliance, openness, transparency, or data sovereignty, but our customers ultimately wanted to know if they could trust their cloud platform.

## No love for proprietary clouds

While proprietary public clouds have been very successful, they cannot positively answer a number of questions asked by their customers. For example, none of the existing proprietary public clouds are available for private cloud deployments.

Proprietary cloud vendors are not keen to give their code to customers and let them run it on their own infrastructure. That rules out proprietary public clouds in the hybrid cloud delivery model. The complexity, the costs, and sometimes the plain incompatibility of the public cloud with the customer's private cloud make bursting impossible or very difficult.

While proprietary public clouds have been very successful, they cannot positively answer a number of questions asked by their customers.

Another problem with proprietary public clouds is the fact that they are black boxes: you do not know what software they are running, and you do not know what hardware they are

using. If you want to burst into another proprietary cloud, you need to re-design your infrastructure again, which may be just too expensive.

As we looked at various proprietary cloud platforms, we quickly realized that they were not the right fit for HP and its customers.

## Open = Trusted

When HP started looking for a product we could use as a basis for our cloud operating system and as a cloud delivery platform, we knew that we needed a product that was developed using an open source model that lets everyone who is interested check the quality of the code and its inner workings, and suggest improvements, or submit fixes.

Problems with open source software typically get fixed fast, because software developers in the open source community want to make the code better. Open source software shortens the development process by removing the need to negotiate new license agreements or commit engineering time and resources to implement patches and changes. This open source development process makes open source software trustworthy in the eyes of customers, who no longer trust proprietary black boxes.

On the flip side, we were not interested in projects that had loose governance and were subject to the whims of developers. Our customers want to know that the open source projects they rely on for their business are:

- **Well-organized**—There is a clearly defined process for code submissions, reviews, removal, and maintenance.
- **Stable**—There are roadmaps, features get delivered on time, and there is a long-term vision and commitment.
- **Business-friendly**—Successful open source projects are often those that do not oppose commercial entities adding value to the ecosystem. (See Apache, Linux, WordPress, and others.)
- **Open to feedback from users**—Projects that listen to their users are more likely to make their users happy.

The OpenStack project was not the only open source project of its kind in 2011, but it certainly did exhibit these traits. Here is how the OpenStack project defines “open” on their own blog (source: [openstack.org](http://openstack.org)):

- “Open source (not open core) with appropriate license (in our case, Apache 2.0 which is OSI<sup>1</sup> approved, GPLv3 compatible, and DFSG<sup>2</sup> compatible)
- Open Design
- Open Development
- Open Community”

## Data sovereignty

The project’s commitment to openness helps tremendously in this day and age when customers no longer trust vendors, governments, or each other. As more and more data is stored in the cloud, the following concerns begin to be voiced in conversations related to the issue of trust:

- Data integrity
- Data sovereignty
- Industrial espionage
- Sabotage

The increasingly complex political and legal landscape that businesses and governments deal with daily only serves to exacerbate these concerns.

Proprietary cloud providers cannot give satisfactory answers to the questions related to these hot issues. Going forward we believe the obvious choice for a cloud platform must be open. The hard reality of doing business in this day and age is that trust is at a low ebb and black boxes only increase wariness.

## Maturity

When HP started to look for a candidate for a cloud operating system, we had a choice of building it from scratch, buying one, or joining a mature open source project. Already very mature in 2010, when the OpenStack project was officially announced and HP officially joined in 2011, it had three major releases on its CV.

OpenStack technology is designed to be massively scalable, with no hardware dependencies.

Even at three years old, OpenStack technology is ready for enterprise-grade applications. Large organizations like banks or telecoms are already using it in

production, furnishing the best testament for a software project.

The maturity of the OpenStack platform shows in its design: it has a modular architecture with each component communicating with others and with the outside world via a set of simple, well-defined RESTful APIs.

All OpenStack projects are designed to be easily extended to support new hardware, new hypervisors, new protocols, and new identity verification mechanisms. OpenStack technology is designed to be massively scalable, with no hardware dependencies.

## Appeal to our own people

HP did not have to force our staff to join the OpenStack project. People working in different divisions at HP have been involved in the development of OpenStack technology before the company itself got involved.

It happened because the OpenStack platform was and still is seen by HP's own engineers as the cloud technology solution they were looking for and something they wanted to get involved in. We did not have to force our engineers to work with the OpenStack project. They started submitting code before the company even had a chance to form an official strategy.

This widespread internal adoption of the project proved to be a strong predictor of success that could only benefit our customers. When engineers work on something they would work on even if nobody asked them to, the result often produces high quality code along with a sense of ownership and responsibility that cannot be enforced via an official corporate policy. Bluntly speaking, there is no better indicator of a project's future success than the engineers' willingness to work on it voluntarily.

## Dynamic growth

For a project started in July 2010, the OpenStack project has had a phenomenal run. So far, it has maintained a steady stream of releases approximately every six months, which is an amazing pace of development considering the size and the complexity of the source code. Here are some facts: <sup>3</sup>

- In 2010, the OpenStack project was a collection of just two projects (Nova and Swift).
- In 2013, there were already nine projects under the OpenStack umbrella.
- There are now over 1,500 developers contributing code to the OpenStack code base.
- The number of integration tests has grown from just 70 to over 700.
- The total number of all tests, unit and integration has passed 15,000.
- The Grizzly release alone added 230 new features and 7,620 patches from 517 contributors.
- The number of developers working on the OpenStack project is greater than the number of developers working on any other open source cloud operating system project.

These measures establish the OpenStack project as the most dynamically growing open source project in history.

## Industry-wide support

The OpenStack project has wide industry support from many major vendors. The list of Platinum Members of the OpenStack Foundation includes some of the largest companies in the business:

- AT&T
- Canonical (Ubuntu)
- HP
- IBM
- Nebula
- Rackspace
- RedHat
- SUSE

Other large companies that support the OpenStack project include Cisco, PayPal, and Yahoo!.

## Global community

A strong, enthusiastic community is crucial to the long-term survival of any project, but it is especially important for the survival of open source projects. A large, dynamic community of developers, users, and supporters serves to make the project grow and thrive. The global OpenStack community ticks all those boxes.

The OpenStack project has a global community of more than 10,000 supporters, developers, and users in nearly 100 countries actively participating in code development, online discussions, deployment, and support. More than 200 companies are involved in the OpenStack ecosystem.<sup>4</sup>



## Sound governance model

The OpenStack community is a meritocracy. The project's legal status is taken care of by the OpenStack Foundation. The project is governed by the following bodies:

- The OpenStack Foundation and its board of directors.
- The Technical Committee.

Each governing body is responsible for a different area of the project's activities and has its roles and responsibilities defined in the Foundation's bylaws. Along with these two official governing bodies, OpenStack benefits from an active user community.

### The OpenStack Foundation

According to its bylaws, the OpenStack Foundation's purpose is to “develop, support, protect, and promote the open source cloud computing project which is known as the OpenStack Project.”<sup>5</sup>

### The board of directors

The board of directors of the OpenStack Foundation is responsible for the strategic and financial oversight of the Foundation resource and its staff. It cannot exceed 24 members.

Since no organization can have more than two members on the board of directors, it is impossible for any one entity to dominate the direction of the project and cause tensions that may lead to a breakup of the project.

### Members

The OpenStack Foundation recognizes three types of members, with different voting powers and numbers of seats on the Foundation's board of directors:

- **Platinum**—No more than eight members. These members are companies who make significant strategic commitments to the OpenStack project. They are expected to commit funding and resources to the project.

Platinum members become involved at that level because they have aligned their corporate strategy with the mission of the OpenStack project. To achieve their goals they also assign their own staff (developers, legal, documentation, etc.) to work on the OpenStack project. In exchange for their contributions to the project, Platinum members appoint their representatives to the board of directors. Each member can appoint a maximum of one representative.

- **Gold**—No more than 24 members. This level of membership requires lower financial and resource commitments. Members elect their own representatives to the board of directors. This process allows for another representative affiliated with a Platinum member to be appointed to the board of directors. However, if this happens, it is resolved by the fact that any organization may have at most two representatives on the board of directors.
- **Individual**—Unlimited. These members may elect representatives to the board of directors.

The seats on the board of directors are split in three equal parts among Platinum, Gold, and Individual members.

### Technical Committee

The Technical Committee of the OpenStack Foundation is responsible for defining and delivering the technical goals of the OpenStack software project. The Technical Committee has absolute and final say on all technical matters related to the project. The members of the Technical Committee are elected by the active technical contributors on a staggered basis, half of the committee being up for election every six months.

### User Committee

The User Committee was created to give the users of OpenStack technology a way to relay their requirements and concerns in an organized, formal fashion. Its goals are as follows:

- Consolidate user requirements. This gives users a voice on all matters related to the project and helps the Technical Committee and the board of directors define future goals of the project.

- Present user requirements to the OpenStack Foundation board of directors and to the Technical Committee.
- Create guidelines that instruct developers when and how to ask for user feedback.
- Track OpenStack deployments and use. The project maintains an extensive analytics website that gives deep insights into the project's progress.
- Disseminate OpenStack user stories. These are essential to help accelerate the adoption of OpenStack technology.
- Work with the global community of OpenStack user groups. With over 10,000 members, the global OpenStack community must be managed in a way that helps each member participate and profit from being a member.

### Legal Affairs Committee

This is a legal advisory and oversight committee responsible for compliance with legal requirements. One of its most important tasks is oversight of the project's intellectual property landscape.

### New projects

The OpenStack Foundation defines an official process for the inclusion of new projects under the OpenStack umbrella. Each new project that wants to join the OpenStack project must go through the following stages:

- **External**—The first stage in the life of a new project begins with an idea and an initial implementation, and ends with a finished design and a working implementation that all interested parties can agree upon. Usually this work is done on StackForge, which is a portion of the project infrastructure provided for the gestation of as-of-yet unofficial projects.
- **Incubation**—This is the second stage in the life of a project, once it has been accepted as a project aligning to the direction that OpenStack technology would like to head. At this stage, the project gets to use official OpenStack resources and works towards being ready for

integration with the rest of the projects, but it is not yet in a position to be depended upon by the rest of OpenStack. Incubation ends with a graduation review from the Technical Committee resulting in a vote on whether to integrate or not.

- **Integrated**—The project gets released as a part of the official OpenStack integrated release. All integrated projects are expected to work with each other and are all tested for interdependency.
- **Core**—The project is nominated by the board of directors of the OpenStack Foundation to join the OpenStack Core. Once a project is core, it can use the OpenStack trademark. Discussions continue regarding what ongoing core definitions mean for consumers.

### Unified licensing

The license for the OpenStack project and all of its member projects is the Apache 2.0 license. Having a single license streamlines the legal side of project adoption. The Apache 2.0 license is well known and understood and chances are that by now it has already been vetted and accepted by the customers who might benefit from OpenStack technology.

### Funding

Platinum Members provide the majority of funding for the OpenStack Foundation. The remainder is provided by Gold Members, Corporate Sponsors, and Startup Sponsors.

## An easy decision to make


With so many arguments in favor of adopting the OpenStack platform, it seemed unwise for HP to not consider it, or to instead follow the proprietary route. The only other option, developing our own cloud operating system, just did not make business sense.

Going at it together—rather than alone—proved the smartest route, and that is why on July 27, 2011, HP officially announced its intention to join and support the OpenStack project.

## REFERENCES

- 1 Open Source Initiative. [opensource.org](https://opensource.org)
- 2 Debian Free Software Guidelines, a part of the Debian Social Contract. [www.debian.org/social\\_contract](https://www.debian.org/social_contract)
- 3 CloudScaling. [www.cloudscaling.com/blog/cloud-computing/openstack-at-3-this-is-what-winning-looks-like/](https://www.cloudscaling.com/blog/cloud-computing/openstack-at-3-this-is-what-winning-looks-like/)
- 4 Stackalytics. [www.stackalytics.com/](https://www.stackalytics.com/)
- 5 OpenStack. [www.openstack.org/legal/bylaws-of-the-openstack-foundation/](https://www.openstack.org/legal/bylaws-of-the-openstack-foundation/)

# Just how serious is HP about OpenStack technology?



---

## In this chapter

- ✓ The history of HP's involvement with the OpenStack project
  - ✓ How HP participates in the OpenStack ecosystem
  - ✓ HP's commitment to OpenStack technology
  - ✓ HP Helion
-

HP has been involved in the OpenStack community from the early days and was instrumental from the beginning of the project. Our first involvement was a form of unofficial, grass-roots participation by HP's engineers, who simply wanted to get involved in a cool project, or felt that OpenStack technology could potentially help them solve the problems they were facing at work. HP engineers working on our own public cloud found out about OpenStack technology and started contributing their time and code.

From this humble and unofficial beginning, HP moved to the point of considering OpenStack project for its cloud computing platform. As noted in the last chapter, we evaluated many proprietary and open platforms and came to the conclusion that OpenStack cloud computing met all of HP's requirements. Once this decision was made, HP jumped in with determination.

## HP's contributions to the OpenStack project

HP is one of the major contributors to the OpenStack project in terms of funding, resource allocation, testing, community participation, code contributions, training, and commercial deployments.

### Involved from the beginning

HP employees have been involved in the OpenStack project from its early days in 2010, and formally since July 27, 2011 when HP announced its intention to join and support OpenStack.

HP is a Platinum Founding member of the OpenStack Foundation.

### Project governance

HP takes its involvement in the project seriously and we take responsibility for all facets of the project, including its governance, where HP is involved at all levels with:

- Two members on the Board of Directors

- Members on the:
  - Technical Committee (3 elected members)
  - Legal Committee
  - User Committee
  - Project Incubation
  - Training
  - DefCore Committee
- Project Technical Leads:
  - Horizon
  - Ironic
  - TripleO
  - Trove

Our involvement in all aspects of the project governance gives us direct access and a chance to influence project strategy, governance, roadmap, blueprints, and development. This ensures that our customers know that they are working with people who have intimate knowledge of all aspects of OpenStack technology, and that their feedback will be listened to and relayed to other members of the OpenStack project.

We cannot force anyone to implement or accept the changes we are proposing, but we can make sure that they are at least heard and considered.

## Financial support

As a Platinum Member, HP is committed to provide its share of funding for the project. This is on top of our voluntary commitment of infrastructure, time, power, staff, and code.

## Strategic commitment

As a Platinum Member of the OpenStack Foundation, HP has pledged to align its corporate strategy with the OpenStack mission.



## Commitment to the project governance and processes

HP is fully committed to the OpenStack governance model and the checks and balances mechanisms implemented in the project's bylaws, which specifically prevent a takeover of the project by any one member, tame ambitions, and force everyone to work together on a common vision.

Wherever we find the process too slow from our point of view, we create add-ons compatible with the open source core of the OpenStack platform. That means we have to bear the burden of maintaining their compatibility with the core.

These decisions are not predatory, but rational—we cannot expect anybody else to write code that supports our own products or services or implement our ideas for the future. We have to take care of that ourselves, and quite often we open source parts of our extensions and add-ons back to the OpenStack code base.

The board structure specifically prevents any party from hijacking the project. We accept and encourage that model, because we believe it is essential to the project's growth and long-term survival. By doing this, we protect our clients' best interests and investments.

## HP is always one of the top five code contributors to OpenStack

When the OpenStack project began its life as an open source project, it had two major code contributors and two projects. Today, there are over ten projects and hundreds of source code contributors, yet HP is always in the top five employee contributors, including being the number one contributor for the Havana and number two for Icehouse. These contributions represent a massive investment in developer time on HP's part, and we plan to continue that commitment in the future.

## Dedicated staff

While developers are a significant part of our staff assigned to work on the OpenStack project, they are by no means the only ones participating in the project. Other employees of HP include operations personnel, documentation teams, training teams, and QA staff.

Our employees act as Technical Leaders for some projects and our legal staff help watch over the project's intellectual property and its legal status.

Running our own OpenStack-based public cloud gave us unique and invaluable practical experience.

HP is the only organization other than the OpenStack Foundation itself that provided a dedicated group of employees to the OpenStack Developer Infrastructure and Continuous Integration projects.

and Continuous Integration projects.

We also contribute OpenStack cloud accounts required to keep the OpenStack Developer Infrastructure and Continuous Integration projects running.

## Successful commercial deployments of OpenStack technology

When we committed ourselves to the OpenStack project, we quickly started using the software to run our public cloud, because that was the only way we could find out how well the OpenStack platform would cope with customer demands as well as determine what it needed in order to become a viable solution for the enterprise.

Running our own OpenStack-based public cloud gave us unique and invaluable practical experience deploying thousands of OpenStack compute nodes and multiple petabytes of storage managed by OpenStack technology.

The hands-on experience we gained helped us understand how OpenStack technology needs to be hardened, deployed, and enhanced for production use in real-life applications. We had to develop our own tools, extensions, plug-ins, and management and monitoring add-ons.

## Transparency and accountability

HP is fully aware of the fact that our involvement and our fulfillment of the promises made to customers and to the OpenStack community are open to public scrutiny. By virtue of being involved in an open source project, HP has to be committed to the vision and the rules of engagement of the project. Everybody can see our code, proposals, decisions, votes, blueprints, and comments.

Nobody, not even the largest players, can simply take from the project. They also have to contribute. The history of our involvement with the OpenStack project confirms our commitment to its principles.

## Expert leadership

The cloud team at HP includes industry veterans with extensive experience building, managing, governing, and delivering large-scale projects and enterprise-grade products based on Linux, Apache, and other major open source systems. Many of these team members used to work on open source projects at Canonical, IBM, Sun, or MySQL, and have extensive knowledge of the enterprise market.

## Cooperation and giving back to the community

HP firmly believes in sharing and giving back to the community. These values have been deeply ingrained in our culture by our founders, therefore we do not believe in a zero sum game in the case of the OpenStack community.

Even though it may seem contradictory to our goals, we sometimes work with our competitors, helping them integrate their solutions with the OpenStack platform, because we believe that we will all benefit from OpenStack technology, if we help our customers integrate it with the solutions they are already using or choose to use in the future.

We want to work with our competitors to help create a rich ecosystem around OpenStack cloud solutions where all kinds of participants, large and small, can find opportunities for success.

## HP Helion

The IT industry often spews nothing but static in an attempt to address the issues of the day. Thus, it is fair to ask, is HP truly committed to OpenStack technology? Well, how about betting everything on the cloud? That is essentially what HP is doing with the OpenStack platform and our HP Helion portfolio.

- The analysis of the trends discussed in Chapter 2, “A case for cloud computing,” has convinced us that the cloud is the future.
- The cloud is already driving many of the internal efforts inside HP and you can expect it to become even more important in the coming months and years. Once we made our choices, we became engaged with the project on all levels.

## An all-hands effort

- The cloud is having a profound effect on HP as a company. Having a clear strategy for the future can work wonders for a business as large as HP. Once we defined our goals and chose the platform, we could see ourselves reaping the benefits of those choices even before we started offering our own distribution of OpenStack software to customers.
- Just like Linux in the past, the OpenStack project is one of those projects that introduces major changes to the way a company thinks, operates, and sees its future. It has caused HP to revise our goals to align them with the HP Helion distribution.
- HP is fully committed to making cloud technology a significant distribution channel for delivery of our future products and services.

## What is HP Helion?

HP Helion is a comprehensive portfolio of products and services that make it easy for organizations to build, manage and consume workloads in a hybrid IT environment. HP Helion moves beyond cloud to become the very fabric of your enterprise. It brings together all the benefits and agility of cloud computing, all the possibilities and interoperability of open source, and all the security and reliability that enterprises need to move forward with confidence.

## It's a journey

Cloud is a journey for HP just like it is for our customers. We started with a good base, and we have already built tools and services that enhance the experience. We have been learning by doing: submitting code, helping the code get better, letting people try OpenStack software, and running our own HP Public Cloud, which is the second-largest public cloud based on OpenStack technology. We continuously test and deploy the latest code from the OpenStack core.

We are learning from all of these efforts, and what we have learned helps us better serve our customers and the OpenStack community. We are packaging our experiences as technology such as HP Helion OpenStack or as extensions to OpenStack software. We are also enhancing all of our products and services that can benefit from HP Helion OpenStack.

It is an exciting time to be involved in the OpenStack project and we are as committed to it as we can possibly be. Our commitment has allowed us to become experienced, knowledgeable guides helping our customers on their own journey into the cloud.

## Chapter 5

# What is HP Helion OpenStack?



---

### In this chapter

- ✓ Why we created HP Helion OpenStack
  - ✓ Main advantages of HP Helion OpenStack
  - ✓ Where does OpenStack and HP Helion OpenStack fit within HP's strategy
  - ✓ Who HP Helion OpenStack is for
  - ✓ Partner and ISV support
-

HP Helion OpenStack is an open and extensible cloud platform based on OpenStack technology. It is a curated, tested, enterprise-grade distribution of OpenStack cloud software, designed to deliver the best open source cloud computing technology in a stable, maintainable, and easy to install and manage package. Essentially, HP Helion OpenStack is enterprise-grade OpenStack technology, specifically designed to cater to the needs of large organizations looking for a platform that enables a hybrid cloud delivery model.

### Why package free software?

We decided to package a curated, tested, enterprise-grade OpenStack to cater to enterprise and service provider customers. To achieve that, the software must be delivered in a way that makes it easy for the customer to use. This approach has helped turn Linux into the dominant server operating system and we want to help repeat that success story with OpenStack technology.

## Why we created HP Helion OpenStack

Software projects, including those releasing the fruits of their labor as open source, often produce source code not fit for immediate consumption. The more complex the code, the more time and effort it takes to get it running. This explains the popularity of various Linux distributions, software package installation tools, and setup tutorials.

At the end of the day, customers want a working product that does not require them to spend time experimenting with various configuration options. OpenStack software faces issues similar to those faced by Linux—with over two dozen projects under the OpenStack umbrella, it takes a lot of learning and experimentation to get them all configured and running the way the user wants.

Even when open source software is packaged in a way that makes it possible to deploy by developers with a deep understanding of its internal architecture, it is not usually fit for use in the enterprise, if only because of the sheer number of possible configuration options and various dependencies that typically come with unique configuration processes.

OpenStack software is, in this regard, quite similar to the Linux kernel, which by itself is not of much use to most users. The source code of the kernel and the additional software that turns the kernel into an operating system are not very useful in source code form. Linux becomes useful for production applications only after it has been tested and packaged in a way that makes all those millions of lines of code easy to install and manage.

Just as Linux needs to be packaged and offered as a curated distribution, so does the OpenStack platform. The challenges of installing OpenStack software are even bigger, because you need to create a network of nodes, install an operating system on some of them, install and configure various components of OpenStack on top of that system, and hope it will all work together.

Our customers want to run OpenStack software, but they are not comfortable playing with various configuration settings, certainly not when they are trying to deploy the OpenStack platform on hundreds or thousands of servers.

HP Helion OpenStack was developed to address the complexities of installing and running an open source project in an enterprise environment.

## Main advantages of HP Helion OpenStack

HP Helion OpenStack adds functionality to the OpenStack platform, bringing much needed improvements for our enterprise customers. This is what we mean by “enterprise-grade OpenStack technology.”

### Simple installation

The OpenStack platform is an open source product that is very powerful and very flexible, but requires extensive knowledge of its inner workings to properly set up each component, using configuration files to make them all work together. Customers usually do not have the resources necessary to gain all that knowledge and would rather use an installer that hides all that complexity behind a simple interface.



To make OpenStack software more friendly to the operations staff, our engineers spent a lot of time enhancing the current OpenStack deployment framework. We also added a visual cloud installation and configuration tool that helps the user install a pre-configured cloud.

The tool provides a wizard-like experience to deploy complex services with easy point-and-click configuration options. The users decide which nodes they wish to deploy, and which services, and the tool handles the installation and configuration.

By making it easier to install, HP makes HP Helion OpenStack an attractive proposition for those who are discouraged from trying OpenStack software because of its “some assembly required” nature.

## Hardened code

Years of experience running the second biggest cloud powered by OpenStack technology have taught us a lot about hardening OpenStack code for enterprise deployments. This expertise helps us curate the code for our customers.

## Improved stability

By packaging OpenStack software within HP Helion OpenStack, HP removes uncertainty related to testing, validation, upgrades, and software certification. Because HP tests all patches and updates before sending them to our customers, there is less chance of disruption. HP manages the process of implementing the latest developments in OpenStack technology in an easy to manage, deploy, and test manner.

## Faster improvements

The advantages of developing HP Helion OpenStack in-house is that we can focus on innovations and improvements that the community is not focusing on today, even though these are very important requirements for enterprise and service provider customers. We can then contribute those innovations back to the community.

## Well-paced innovation

By providing a stable mechanism for applying and rolling back patches and upgrades, HP Helion OpenStack helps enterprise customers ride the wave of innovation while using a trusted, tested, and deployable platform. HP Helion OpenStack lets customers choose which new OpenStack software features they want to enable and when.

## Simplified management

HP Helion OpenStack helps enterprise customers keep up with the speed of development of OpenStack technology, in an organized, predictable manner. It takes the pain out of OpenStack upgrades and patches by automating that process.

HP Helion OpenStack plugs into a content distribution network which distributes patches and upgrades straight from HP to the customer's cloud environment. All patches and updates to the HP Helion OpenStack environment are thoroughly reviewed and tested by HP which only releases

HP Helion OpenStack helps enterprise customers ride the wave of innovation while using a trusted, tested, and deployable platform.

patches that can be safely applied to the customer's cloud without disrupting operations.

The administrators see those patches and updates in their HP Helion

OpenStack operational dashboard and can decide which ones they want to apply. Applying patches and updates requires zero cloud downtime. If necessary, updates and upgrades can be rolled back using the same visual interface found in the HP Helion OpenStack operational dashboard.

## Automated delivery of additional functionality and content

Patches and updates are not the only types of content that can be delivered to the customer's cloud. The content distribution network can also be used to deliver additional functionality in the form of plugins, extensions, new modules, and more.

The content distribution network can deliver additional content, such as virtual machine images, infrastructure templates, scripts, or training materials. HP is constantly evolving the network. It is quite possible that in the near future, customers will be able to download and install infrastructure templates for building backends for web apps, mobile apps, online shops, development environments and more.

## Interoperability

A common, open source cloud platform makes it easier to achieve interoperability between clouds offered by different providers or deployed at different data centers owned by the same organization. Workload transfer becomes as easy as executing an infrastructure template.

## Good fit for all data centers, large and small

With the advent of the HP Moonshot hardware platform, your data center need not live in a bunker, but may fit in the office next door. HP Helion OpenStack is designed to be scalable and work with OpenStack-compatible hardware infrastructures of any size, from a few old servers to an air-conditioned underground data center housing thousands of machines. This makes HP Helion OpenStack a good fit for both departmental clouds and for whole enterprises. In addition to our HP Helion OpenStack Community distribution, we also offer a an HP Helion OpenStack edition that has been optimized to scale for the very large needs of enterprises and service providers.

## Simple data sovereignty

By using open source software and technologies, and deploying their workloads on private clouds, customers can gain the convenience of large-scale cloud management and the maximum possible levels of data sovereignty.

## Lower cost of migration

Clouds based on proprietary solutions increase the cost of migration, and might even require reengineering of the whole application. This process may have to be repeated for each public cloud, which is unfeasible and ultimately limits workload portability.

Also, with HP Helion any third-party solutions based on OpenStack technology can be easily migrated along with the customer's base infrastructure. This is not possible with a proprietary cloud.

Standardizing on the OpenStack platform gives cloud users a lot more freedom of choice and avoids their being held hostage because of an investment in just one vendor's ecosystem. It also provides an opportunity for third parties to offer solutions based on OpenStack technology—that can be migrated along with the whole cloud.

The concept of not being able to get your code or data out of a third-party provider's data center is not acceptable.

## A truly open cloud operating system

HP's focus is on making HP hardware and software work well with the HP Helion portfolio and OpenStack technology. But, thanks to the open source nature of the project, our customers will be able to use HP Helion OpenStack with hardware and software provided by other vendors, too. Open source provides the ability for vendors to provide drivers that expose their solutions to the OpenStack platform. The distribution is free to license, so independent software vendors (ISVs) can easily use it to create their own applications.

HP Helion OpenStack follows the architectural principles behind OpenStack technology. Our enhancements to OpenStack software are compatible with the existing extension frameworks, or implemented as separate modules accessible via RESTful APIs.

## Model once, deploy everywhere

HP Helion OpenStack introduces a powerful set of tools for modeling and provisioning infrastructure on different OpenStack-based clouds in the form of the infrastructure templates used to describe infrastructure in terms of required resources and the relationships between them. These templates follow the industry OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) standard.

- Our implementation of infrastructure templates includes hooks for executing actual deployment scripts.
- HP Helion OpenStack includes template versioning tools to enable rollbacks and change tracking.
- The content distribution network can be used to distribute infrastructure templates for those who would rather reuse a tested infrastructure template than spend time building and testing their own.

## Rapid provisioning

Once templates have been designed and tested, administrators can reuse them and quickly recreate the infrastructures they represent.

Infrastructure templates let you model your infrastructure, connect your model to deployment scripts and use those highly abstract recipes to provision complex backends with only a few clicks.

## Sharing resources

The HP Helion OpenStack resource pool registry and capability tagging service allows customers to publish and browse information about available cloud resources. This facilitates cloud resource selection optimized for specific workloads, streamlining utilization of their clouds.

## Multi-cloud provider support

The HP Helion OpenStack resource pool registry can store information about both public and private clouds. Right now, when there is a need for extra capacity, the cloud administrator has to provision it using infrastructure templates manually. In future releases of HP Helion OpenStack, administrators will have an option to dynamically provision additional capacity by setting rules and triggers for infrastructure templates and for the matching resource pools. This will enable automated bursting from private into multiple public clouds.

## Hybrid provisioning

HP Helion OpenStack helps its users achieve the goal of hybrid delivery with the help of the infrastructure templates and the resource pool registry. HP Helion OpenStack templates let cloud administrators describe what resources they need and then choose the cloud they want to deploy that infrastructure on by selecting the resource pools that match the template's

[HP Helion OpenStack helps users achieve the goal of hybrid delivery with the infrastructure templates and the resource pool registry.](#)

requirements from a list returned by the registry.

Future releases of HP Helion OpenStack will include template

functionality that allows you to split workloads across different clouds and deploy databases on a private cloud while the web application server runs on the public cloud.

## Supported by HP

A lot of enterprise customers are facing the same issues with OpenStack technology that they faced a few years ago with Linux—there was nobody who would take responsibility and support enterprise users; there was nobody who would try to help introduce changes or improvements to the Linux source code in the form of drivers, file systems, or protocol support. HP takes responsibility. We provide paid technical support options to ensure HP Helion OpenStack stays up and running at peak performance.

## HP Helion foundational technology


HP designed the HP Helion OpenStack distribution to be the common foundation for HP Helion products and services. HP Helion OpenStack harmonizes internal efforts at HP. It is a common platform for all departments, hardware, software, and services. HP Helion OpenStack will be supported by all future HP products and services.

## Partner and ISV support

The rich HP Helion OpenStack ecosystem enables developers, software and hardware vendors, and worldwide partners to create new solutions and tap new markets. HP makes HP Helion OpenStack available for testing and integration purposes. Hardware and software vendors can certify their products with it through our comprehensive certification programs, which include consultation with our technical experts and infrastructure resources. Certified partners are able to leverage HP's marketing engines and other resources to accelerate business momentum, open new markets, and drive additional revenue.

## Chapter 6

# Inside HP Helion OpenStack



---

### In this chapter

- ✓ How HP Helion OpenStack enhances and augments OpenStack technology
  - ✓ How much OpenStack code is inside HP Helion OpenStack
  - ✓ The future of HP Helion OpenStack enhancements to the OpenStack platform
-



HP Helion OpenStack is a hardened, curated distribution based on OpenStack software designed to meet the expectations of the enterprise customer. The OpenStack project is, in many ways, similar to the Linux kernel project—the value of its source code increases greatly through the addition of the software that lets it interface with the open and proprietary products and services developed. HP Helion OpenStack packages OpenStack software and value-add services, written to extend the core OpenStack code base for hybrid delivery, into a product that enterprise customers can easily install, use, and manage.

HP Helion OpenStack offers two distributions to meet users' different needs:

- **HP Helion OpenStack**—A hardened and secured enterprise-grade product that speeds cloud deployment and cloud application development while simplifying management for large scale private, public and hybrid clouds.
- **HP Helion OpenStack Community Edition**—A pure and free-to-license distribution that speeds deployment and simplifies the management of small scale, open cloud environments and infrastructure services.

### Staying true to the community

It is our intention to build value-add extensions to the OpenStack project in order to support our customers. We take responsibility for making sure that our contributions work with OpenStack technology, and we are evaluating ways we could open source our own technology and give it to the community.

## The architecture of HP Helion OpenStack

HP Helion OpenStack retains the OpenStack platform's highly modular design and follows the conventions of communicating between components via RESTful APIs. This design makes the OpenStack platform and HP Helion OpenStack more resilient, scalable, and maintainable by separating the implementation details of various functionalities from each other.

The architecture of HP Helion OpenStack can be broken into two environments:

- **Execution**—Managing above the cloud. This is where you create projects, users, and provision infrastructure.
- **Administration**—Managing below the cloud. This is where you define and manage the configuration of your cloud, where you manage and maintain your cloud, and where you extend/expand it.

The Execution environment offers functionality that makes the cloud run on top of the customer's infrastructure while the Administration environment offers functionality required to install, configure, manage, and upgrade an HP Helion OpenStack-based cloud.

Furthermore, the Execution environment is broken into five layers:

- **Kernel**—Implements functionality similar to the kernel of an operating system, abstracting the complexity of the underlying infrastructure behind a unified, application interface implemented as a set of RESTful APIs. These APIs serve a very similar purpose to the standard libraries found in Linux or Windows operating systems.
- **Base**—Where you will find functionality that supports the work of the kernel.
- **UI**—Provides a unified user experience for managing tasks in the under cloud and the over-cloud.
- **Sub-system**—Expands supportability functionality for managing enterprise-grade cloud environments.
- **Linux platform**—Offers an optimized Linux operating system running the managed cloud controllers and nodes.

## The HP Helion OpenStack kernel

The HP Helion OpenStack Kernel is comprised of the following OpenStack components:

- Keystone (identity management service)
- Glance (virtual image management service)
- Ceilometer (telemetry)
- Heat (orchestration)
- Nova (hypervisor and virtual machine management service)
- Cinder (volume management service)
- Swift (object storage service)
- Neutron (virtual networking service)

Each of the Kernel components in HP Helion OpenStack ships with additional plugins that enhance the overall interoperability of the system.

### Identity service (Keystone)

The Keystone component provides Identity, Token, Catalog, and Policy services for other components of HP Helion OpenStack. It is a centralized service framework with an extensible plugin architecture designed to support multiple identity verification methods.

Each of the Kernel components in HP Helion OpenStack ships with additional plugins that enhance the overall interoperability of the system.

Being implementation-agnostic allows Keystone to be quickly and easily integrated with existing and future identity technologies.

Developers only need to write an extension compatible with the Keystone framework. The extension passes user credentials on to the identification service. Once a positive verification is confirmed, Keystone issues its own internal User token.

Supported user credentials include:

- Token—a secret identifier that represents the user.
- Username and an API key
- Username and password

Upon successful authentication, users registered with HP Helion OpenStack receive a token that identifies them within that cloud. Once they obtain the token, they can use it to obtain access to Nova, Glance, Swift, and other HP Helion OpenStack components. Keystone provides a Catalog service that lists endpoints accessible to a particular user.

How much access users have and to which components depends on their roles within the access policy defined by the cloud administrator.

Since Keystone exposes its functionality via a RESTful API, adding new components that require a robust authentication and authorization mechanism is very easy. Keystone is as implementation-agnostic on the inside as it is on the outside. For example, it supports a variety of database back-ends for storing user information, making it possible to choose the database your staff has the most experience with, or to connect to the organization's Lightweight Directory Access Protocol (LDAP) backend.

Among other improvements to the OpenStack platform, HP has contributed HP-IDM Admin Extensions for access to infrastructure templates provided by HP Helion OpenStack. Other HP contributions to Keystone include global identity support, account management, billing, and Customer Relationship Management (CRM) integration.

### **Image service (Glance)**

Glance is used to discover, register, retrieve, and manage the catalog of the virtual machine images within an HP Helion OpenStack cloud. Unlike database or file system backups, virtual machine images store the complete state of a machine and its file system, letting you stop and restart them at a later date. Provisioning a new server, complete with the software and often

the data it needs to operate is now a matter of loading the virtual image onto a virtual machine. This enables a number of scenarios:

- At the most basic level, OpenStack cloud administrators may use Glance to create pre-configured application servers, development environments, staging environments, etc.
- You do not have to limit yourself to single servers. It is possible to store whole backends containing load balancers, application servers, database servers, queues, and more saved as virtual machines, ready for deployment at a moment's notice.
- If you require an even higher level of sophistication, you can turn parts of your cloud off during the day and restart them at night when the prices of energy are lower. Or, if you require continuous availability, you can start and stop parts of your infrastructure in different time zones so you are always paying less for energy.

For all those scenarios, you need a robust image management service with a simple management interface and an API. That is what Glance was created for. It provides two APIs:

- **Metadata API**—For managing virtual image metadata. It is used to create, modify, download, and delete virtual images and their associated metadata records.
- **Binary Data API**—For managing actual virtual images. You can use it to store and download the images.

Glance is independent of the storage medium—you can use it with a local disk array, network storage, block storage, or object storage.

HP Helion OpenStack introduces enhanced metadata support in OpenStack via the infrastructure templates handled by the Eve and Focus services.

## Telemetry (Ceilometer)

Ceilometer provides an infrastructure to collect measurements within OpenStack. Its primary targets are monitoring and metering, but the framework is easily expandable to collect data for other needs as well. Ceilometer today is comprised of two main components: meters and agents.

Meters provide the measurement to be captured using cumulative (increasing over time), gauge (discrete items), or delta (changing over time) types. Agents are the deployment mechanism for the meters.

### **Orchestration (Heat)**

Heat is the template-driven engine in OpenStack that allows application developers to describe and automate the deployment of infrastructure. Heat uses a flexible template language that can specify compute, storage, and networking configurations as well as detailed post-deployment activity to automate the full provisioning of infrastructure as well as services and applications. The orchestration engine is also capable of performing auto-scaling of certain infrastructure elements.

### **Compute service (Nova)**

Nova implements the OpenStack cloud computing system and binds all of its components using a shared-nothing, modular, messaging-based architecture via RESTful APIs. Nova is responsible for managing hypervisors and virtual machines.

Nova is not tied to any particular hardware virtualization technology, data store, identity service, or networking topology. All of its components communicate with each other via RESTful APIs and are themselves implementation-agnostic. Such an approach results in a highly flexible framework that can be adapted to new technologies with as little effort as necessary.

HP's contributions to Nova include adding scalability across Availability Zones and Regions. HP Helion OpenStack ships with a plugin for the Kernel-based Virtual Machine (KVM) hypervisor.

### **Block storage service (Cinder)**

Cinder is an interface between cloud servers and block storage resources. It can be used to connect cloud servers to high-performance storage resources for I/O-intensive applications like databases. Cinder volumes can be used to boot virtual machines or they can be mounted under the servers' filesystems.

One particularly useful feature of the Cinder block storage component is the ability to take snapshots of the data for backup and restore purposes. Within HP Helion OpenStack, block storage is also used to store backups of virtual machine volumes.

HP Helion OpenStack supports HP 3PAR StoreServ Storage, HP StoreVirtual Storage, and other products in the HP Converged Storage portfolio. Both Internet Small Computer System Interface (iSCSI) and Fibre Channel are supported. These enhancements let enterprise customers use HP Helion OpenStack with advanced storage products and technologies, such as self-managing storage, software-defined storage, or storage federation.

The HP Helion OpenStack community edition ships with the Logical Volume Manager (LVM) Cinder plugin for raw storage and logical volume support.

### Object storage service (Swift)

Swift is a high-availability, exabyte-scale, long-term storage solution that uses commodity hardware to satisfy demands of enterprise, financial, government, military, or scientific customers who need to safely store huge amounts of data, but do not have to retrieve it often.

Easily scalable replication of data is a built-in feature of Swift. You can configure it to keep in sync multiple copies of files, disks, RAID arrays, or even whole data centers.

Typical uses include scanning paper land records, checks, or bills. Swift can also be used as a backend for a content delivery network (CDN). Major customers using Swift today include banks, telecoms, and governments.

Just like other components and smaller projects that form various parts of the OpenStack platform, Swift too uses a RESTful API to expose its functionality to the outside world. Unlike physical or block storage, Swift object stores cannot be mounted as volumes under the virtual servers' filesystems—objects get stored, retrieved, or deleted via API calls.

You do not have to run Nova to use Swift; it is a standalone system that happens to easily integrate with the compute service.

## Networking service (Neutron)

Neutron provides virtual networking services for the devices managed by the HP Helion OpenStack Nova compute service.

Neutron implements networking as a service and provides an API for defining network connectivity and addressing for virtual machines, block storage, and other parts of your cloud managed by the OpenStack platform. Neutron manages different networks that can be defined within an HP Helion OpenStack cloud:

- **Management network**—Used for internal communication between OpenStack components for the purposes of monitoring and management.
- **Data network**—Internal network used by virtual machines managed by OpenStack.
- **External network**—Provides access to the Internet.
- **API network**—Provides access from the Internet to the OpenStack API.

At its most basic level, Neutron manages networks, subnets, and ports. Extensions to Neutron provide the following additional functionality:

- **Provider network**—Handles mapping between HP Helion OpenStack network objects and the underlying networking infrastructure.
- **Layer-3 networking**—Enables packet routing between internal and external networks through floating IPs.
- **Quotas**—Manages network quotas that limit the number of networks an HP Helion OpenStack tenant can create.
- **Security groups and rules**—Simplifies security groups and rule management on a tenant level.
- **Agent management**—Manages network agents (Dynamic Host Configuration Protocol [DHCP], Neutron OVS, Neutron L3). You can use it to inform the HP Helion OpenStack networking schedulers which network agents they need to provision.
- **ExtraRoute**—Configures extra routes on the router.



- **Load Balancer as a Service (LBaaS)**—Balances network traffic for virtual machines, on a per-machine, per-network, or per-protocol basis. Implements session persistence and application service monitoring.
- **Agent Scheduler**—Used to schedule resources among network agents. Works with the Agent Management extension.
- **Virtual Private Network as a Service (VPNaaS)**—Sets up and manages VPNs for extending your cloud's private networks to other networks. Tenants can create multiple VPN connections, connect two private networks, and use IKEv1 and IPSec policies and strong encryption. Dead Peer Detection is also available as a standard.

HP's contributions to Neutron include support for user-defined networks, software-defined networks, intrusion detection, and load balancing. HP Helion OpenStack ships with the Open vSwitch OpenFlow plugin for managing and plugging into VPNs.

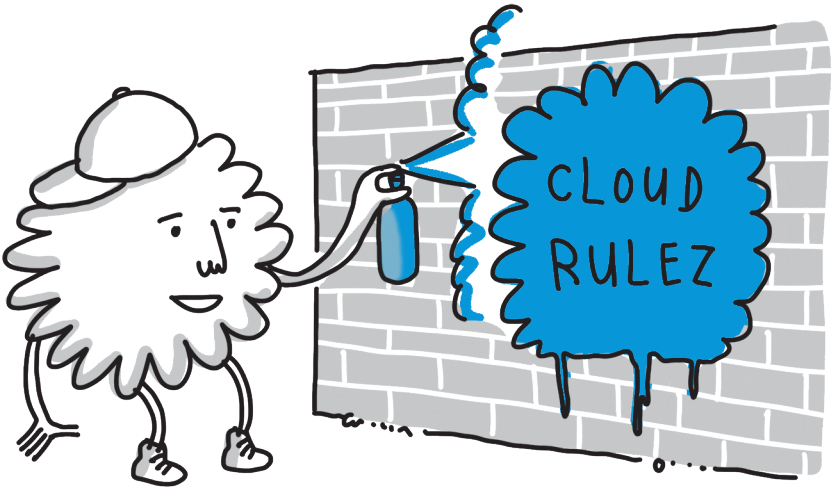
## HP Helion OpenStack value-added services

HP Helion OpenStack value-added services include:

- **Graffiti**—Registry of resource pools and a capability tagging service. Allows you to register private and public resource pools letting you find the best match for the infrastructure you want to deploy. Resource matching is greatly improved by tagging.
- **Eve**—Infrastructure Topology Provisioning Service. HP Helion OpenStack infrastructure template provisioning service.
- **Focus**—Template management, versioning, and relationship management.

### Graffiti

Graffiti is the HP Helion OpenStack resource pool registry service and capability tagging service. It lets you register private and public resource pools that can be used to deploy your cloud. Its purpose is to let you find the best match between the resources available from different providers and the resource requirements for the infrastructure you want to deploy.



Any OpenStack cloud, public or private, can register with Graffiti, making hybrid delivery a matter of choosing the best resource pool for the infrastructure you want to deploy.

## Eve

Enterprise customers expect to be able to operate at scale and any tool that helps them work in a more efficient way is a welcome addition. Mass deployments of virtual servers are handled today by various command-line deployment tools and scripts, which are lacking an extra layer of abstraction.

This problem is solved in HP Helion OpenStack with the topology design template (Eve) service. It is used to design XML-based templates that allow administrators to describe required infrastructure resources and layouts.

Eve infrastructure templates are based on the TOSCA standard. These topology designs define standard, reusable infrastructure models.

HP Helion OpenStack ships with a template Designer tool, which can be used to create XML-based templates that allow administrators to define required infrastructure resources and layouts.

You can use Eve templates to define the functionality of the infrastructure you want to deploy (for example, a test cluster of seven servers, each with

64GB of RAM and 2TB of storage space) and the tools and steps necessary to build and deploy that infrastructure (run Chef scripts to install OS, build software, and test the cluster).

Thanks to Eve, you do not have to run provisioning scripts yourself, but can execute your pre-defined template and Eve will run them for you. Eve can be used to deploy templates in any cloud registered with Graffiti.

## Focus

Focus is the HP Helion OpenStack internal service for template and document management, versioning, and for managing relationships between templates and other documents.

## Administering HP Helion OpenStack

While you can run a personal cloud on a single server for testing and experimentation, real-life enterprise clouds are made up of hundreds or thousands of servers that need to be managed in an orderly fashion, preferably using a graphical user interface (GUI).

The administration environment of HP Helion OpenStack is made up of three distinct parts:

- The Administration and user dashboard used to manage and provision infrastructure—above the cloud.
- The Operations dashboard for the cloud administrators.
- The content distribution network.

## Administration dashboard

The Administration dashboard is the portal that allows you to manage and provision infrastructure—above the cloud. It is a Horizon dashboard with added panels for our value-added services, including additional components and workflows that implement a visual interface to the add-on services found in HP Helion OpenStack.

### **Operations dashboard**

The Operations dashboard is the Horizon portal that allows you to manage and maintain the cloud infrastructure—below the cloud. HP Helion OpenStack cloud administrators can use it to: apply/roll back patches; install, enable, or disable plugins; and to scale the cloud in and out.

### **Content distribution network**

The HP content distribution network is designed to make it easy to manage, maintain, and update your HP Helion OpenStack environment. The content distribution network supports patching, patch rollback, content, and other lifecycle management activities that enterprise customers expect.



# Use cases




---

## In this chapter

- ✓ Who benefits from HP Helion OpenStack?
  - ✓ Possible use cases
-

Our customers are always building amazing things with technology, so it is difficult for us to predict all the possible use cases for HP Helion OpenStack. However, based on the initial feedback and HP's experience with OpenStack technology and our own public cloud, we have identified some very good matches for HP Helion OpenStack and the HP Helion product portfolio:

- Enterprise IT:
  - Developer environments needing to leverage the OpenStack APIs for software development, testing, integration, and staging.
  - IT when it needs to regain control, and deliver the same quality services you can get on the public cloud, just as fast (to avoid shadow IT and reduce cost).
  - Enterprise customers who wants to in-source public cloud workloads because of data privacy conflicts (e.g., Banking and Government customers).
  - Deployment testing environments for DevOps for Continuous Integration/Continuous Delivery.
- Enterprise and Service Providers:
  - Deployment of hyperscale workloads
  - Hybrid delivery
  - Workload mobility
  - Bursting
  - Cloud interoperability




- Service Providers:
  - Service provider who wants to deliver IaaS to the end customer or to ISVs.
  - Providers who want to build and resell their own public cloud.
  - ISVs who want to commercialize a SaaS model and use an API based on OpenStack technology.
  - Those who want to consume IaaS at the same price point as public cloud providers, but under local regulatory requirements.
- Software vendors and consultants:
  - Developers of value-add extensions, for example, using TOSCA standard templates in order to deliver Infrastructure to Applications as a Service.
  - Those delivering cloud-ready applications that are able to scale out on demand and can interface with the infrastructure layer to request more (or less) resources. (For example, large-scale websites that require OpenStack APIs to be exposed.)
  - Providers of ready-made infrastructure templates and template designers.
  - Training suppliers.



HP Helion OpenStack is an especially attractive proposition for enterprise customers who want to maintain compatibility with the OpenStack platform and also get enterprise-grade functionality and support.

Chapter 8

# Where do I go from here?



---

**In this chapter**

- ✓ Online resources
-

For more information on HP Helion, HP Helion OpenStack, the OpenStack Foundation, and OpenStack Technology please visit these online resources.

## Online resources

- To learn more about HP Helion OpenStack, visit [hp.com/cloud/helionopenstack](http://hp.com/cloud/helionopenstack).
- To access the content distribution network, visit [cloudos.hpwsportal.com](http://cloudos.hpwsportal.com).
- To learn more about HP Helion, visit [hp.com/helion](http://hp.com/helion).
- To learn more about OpenStack at HP, visit [hp.com/cloud/openstack](http://hp.com/cloud/openstack) and [docs.hpcloud.com/cloudos/prepare/videos](http://docs.hpcloud.com/cloudos/prepare/videos).
- To learn more about the OpenStack community, visit [openstack.org](http://openstack.org).
- To learn more about OpenStack technology read, “OpenStack Cloud Computing: Architecture,” available at [www.amazon.com/OpenStack-Cloud-Computing-Architecture-Guide/dp/0956355684/ref=cm\\_sw\\_em\\_r\\_dp\\_VSkvtb051JVZCVME\\_tt](http://www.amazon.com/OpenStack-Cloud-Computing-Architecture-Guide/dp/0956355684/ref=cm_sw_em_r_dp_VSkvtb051JVZCVME_tt).

## Acknowledgments

At HP Press, our goal is to create in-depth technical books of the best quality and value. Each book is crafted with care and precision, undergoing rigorous development that involves the expertise of members from the professional technical community. We would like to acknowledge the team of experts who helped bring this book to market.

**Authors:** Lisa-Marie Namphy and Jacek Artymiak

**Illustrator:** Craighton Berman

**HP Press Program Managers:** Michael Bishop and Deena Patel (EPAC)

**HP Contributors:**

Bill Hilf	J.R. Horton
Jerome Labat	William L. Franklin
Monty Taylor	Roger Levy
Deborah Martin	Mark Perreira

**HP Editors:**

Marla Van Baren	Denise Walters
-----------------	----------------

**Publisher:** HP Press

We want to hear from you. Send email to [feedback@hpressbooks.com](mailto:feedback@hpressbooks.com).

---

HP HEADQUARTERS  
Hewlett-Packard Company  
3000 Hanover Street  
Palo Alto, CA 94304



## About this book

This book is meant to explain how cloud computing is a solution to the problems facing data centers today and to highlight the cutting edge technology (including OpenStack cloud computing) that HP helped bring to the stage that it is today. If you are a CTO, data center administrator, systems architect, or an IT professional looking for an enterprise-grade, hybrid delivery cloud computing solution that's open, trusted, and reliable, we wrote this book for you.

This book is also for (and dedicated to) all you pioneers who fearlessly pushed the needle forward and moved the industry to the next major phase of its technology lifecycle. We had a great time producing this book and we hope you will enjoy reading it.

## About HP

HP creates new possibilities for technology to have a meaningful impact on people, businesses, governments, and society. As the world's largest technology company, HP brings together a portfolio that spans printing, personal computing, software, services, and IT infrastructure to solve customer problems. More information about HP (NYSE: HPQ) is available at [www.hp.com](http://www.hp.com).



ISBN-13: 978-1-9378-2670-3  
ISBN-10: 1-9378-2670-8



**HP Press** | [www.hppress.com](http://www.hppress.com)

Cover design: Stoere Binken Design