



General Purpose Reference Architecture: HPE Helion Eucalyptus

Detailed solution example for
Eucalyptus private clouds

Contents

Introduction.....	2
About HPE Helion Eucalyptus.....	2
Eucalyptus architecture.....	3
Core Eucalyptus services.....	3
Management Console.....	3
Object Storage Gateway (OSG).....	4
Cloud Controller.....	4
Object Storage Provider (OSP) and Scalable Object Storage (SOS).....	4
Cluster Controller.....	4
Storage Controller.....	4
Node Controller.....	4
Software availability.....	4
Eucalyptus configuration.....	5
Networking.....	5
Eucalyptus using HPE DL360p Gen9 ProLiant rack servers.....	7
Capacity planning for Eucalyptus.....	9
S3-compatible object storage capacity.....	9

HPE Helion Eucalyptus

The open source AWS-compatible private cloud platform for Enterprise

Introduction

About HPE Helion Eucalyptus

HPE Helion Eucalyptus, hereafter “Eucalyptus,” is an open source platform that allows you to build an Amazon Web Services (AWS)-compatible, on-premise cloud. It is designed to run on commodity hardware and provide an implementation of popular AWS-compatible services, such as EC2 (Elastic Compute Cloud) and Auto Scaling. The platform is highly integrated, which results in a robust architecture and a user-friendly installation and configuration experience.

Eucalyptus can dynamically scale up or down depending on application workloads and is uniquely suited for enterprise clouds, delivering production-ready software that supports the industry-standard AWS APIs, including EC2, S3, EBS, IAM, Auto Scaling, Elastic Load Balancing, Cloud Watch, and CloudFormation. The benefits of using Eucalyptus for a private cloud are highly efficient scalability, organizational agility, and increased trust and control for IT.

Purpose and use case

The purpose of this document is to provide guidance on how to deploy the Eucalyptus 4 platform in a reference configuration on top of Hewlett Packard Enterprise hardware.

The architecture of the platform is designed specifically for pilot or proof-of-concept deployments that aim to test the functionality of Eucalyptus. It is not designed to accommodate heavy EBS or S3 usage, and other reference architectures should be consulted if performance of the deployment is a priority. This document also describes the Eucalyptus software components and their role in the platform, together with further information on capacity planning and scaling considerations.

This reference architecture uses HPE ProLiant DL360p Gen9 servers as an example of system configuration, although other HPE server models of similar spec can work equally as well.

Intended audience

This document is intended for solution architects or technical implementers who are building their own on-premise Eucalyptus cloud platform. You should:

- Be familiar with infrastructure configuration and management.
- Have an understanding of network configuration in such an environment.
- Be familiar with the key concepts and the basic Eucalyptus architecture, as discussed in this document.

Additional information

Official Eucalyptus product documentation includes full material on the installation and configuration procedure. This reference architecture does not implement high availability or full redundancy for Eucalyptus software components, and the reader should rely on disaster recovery procedures outlined in the official documentation.

Eucalyptus architecture

Eucalyptus is a scalable and distributed cloud platform made up of a number of core software components that provide the functionality for the cloud services. Use of the VMware broker is not in scope for this document, and so this architectural component is not discussed.

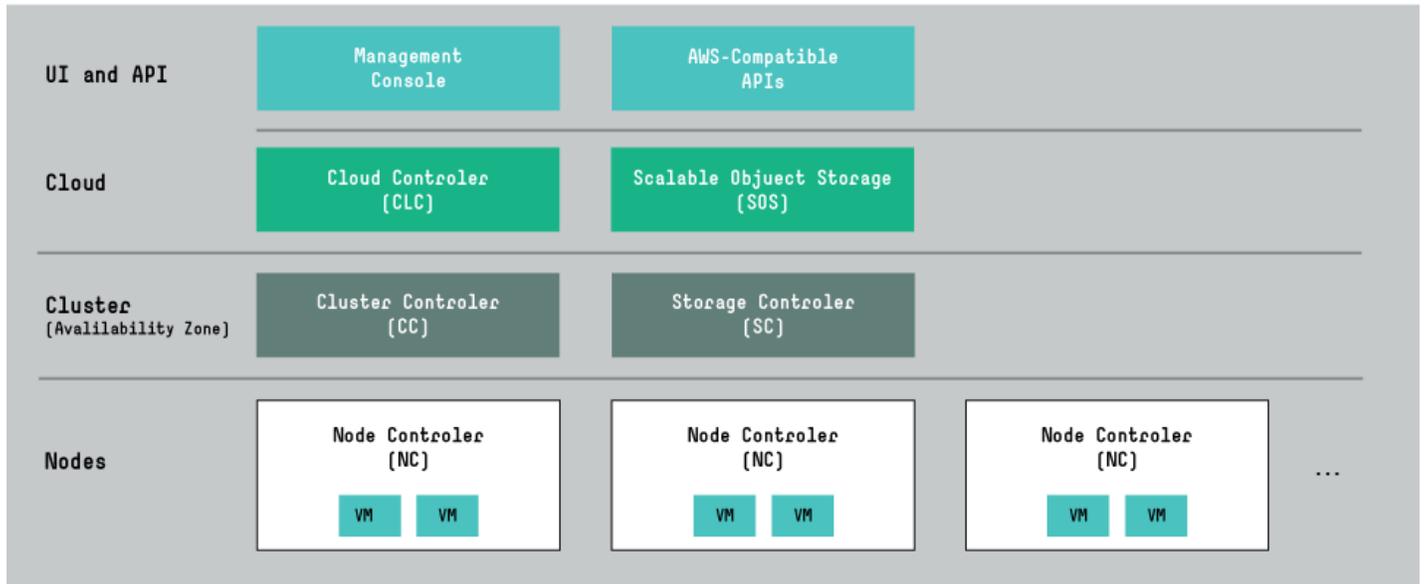


Figure 1. Eucalyptus conceptual architecture

Core Eucalyptus services

This section briefly describes the core Eucalyptus services and provides an overview of the role of each one in turn. These services will be immediately familiar to those who have already explored or implemented a Eucalyptus environment.

User-Facing Services

The majority of the AWS APIs are handled by a component referred to as the User-Facing Services (UFS). This component implements web services for the AWS-compatible service implementations and APIs. This component serves endpoints for clients and users interacting with the cloud platform. This component has been placed on the same host as the Management Console and should be accessible from the client-facing "public" network.

Management Console

The Eucalyptus Management Console is an easy-to-use, web-based interface that allows cloud users to provision and manage resources; it also gives cloud account administrators powerful tools to manage users, groups, and policies.

In this reference architecture, the Management Console resides on the same host as the API services. It is possible to load balance multiple consoles (perhaps running as instances) to provide redundancy, but this is optional.

Object Storage Gateway (OSG)

The User-Facing Services also include the Object Storage Gateway (OSG). The OSG is the Eucalyptus service equivalent to AWS Simple Storage Service (S3). The OSG is a pluggable service that allows infrastructure administrators the flexibility to implement scale-out storage on top of commodity resources using open source and commercial solutions that implement the S3 interface. Eucalyptus provides a basic storage implementation, known as Walrus, which may suit evaluation and smaller cloud deployments. For large-scale and increased performance, users are encouraged to connect the OSG to dedicated storage solutions such as RiakCS. In this reference architecture, the Object Storage Provider (OSP) is RiakCS, which provides enhanced S3 capability through peak performance and horizontal scale out.

Cloud Controller

The Cloud Controller (CLC) is a Java program that hosts the database for resource tracking in the cloud. Only one CLC can exist per cloud. In Eucalyptus 4, the Cloud Controller also handles DNS for services and resources within the cloud, along with the CloudFormation-compatible service implementation.

Object Storage Provider (OSP) and Scalable Object Storage (SOS)

The Object Storage Provider (OSP) is the term used to describe the storage platform that works in conjunction with the OSG to provide the S3-compatible object storage service. Walrus is one such provider and is a basic single-host storage provider for the OSG. It uses any POSIX-compliant file system for object storage. In this reference, we use Basho's RiakCS as the OSP for the cloud deployment.

Cluster Controller

A cluster is equivalent to an AWS availability zone, and a single Eucalyptus cloud can have multiple clusters. The Cluster Controller (CC) is written in C, acts as the network ingress for a cluster within a Eucalyptus cloud, and communicates with the Storage Controller (SC) and Node Controller (NC). The CC manages instance (i.e., virtual machine) execution. In this reference architecture, there is only a single availability zone or Eucalyptus cluster defined. Further clusters could be added through the addition of another Cluster Controller, Storage Controller, and Node Controller. Refer to the official Eucalyptus documentation for further information.

Storage Controller

The Storage Controller (SC) is written in Java and provides the equivalent of AWS Elastic Block Store (EBS) functionality. The SC communicates with the Cluster Controller (CC) and Node Controller (NC) within the distributed cloud architecture and manages Eucalyptus block volumes and snapshots. If an instance needs to write persistent data to disk, it would need to use an EBS volume served by the Storage Controller. The SC interfaces with storage systems, including local, NFS, iSCSI, and SAN. In this reference architecture, the Storage Controller resides on the same host as the Cluster Controller and is intended to use the 3PAR SAN for EBS volumes. For use cases that require extensive EBS storage, a supported SAN is typically recommended. Refer to the official documentation for more information on the supported vendors and models.

Node Controller

The Node Controller (NC) is a service component written in C, hosts the virtual machine instances, and manages the virtual network endpoints. The NC downloads and caches images from the OSG, as well as creates and caches instances on its local disk. In our reference architecture, the Node Controllers each reside on physical hosts.

Software availability

Eucalyptus is distributed as binaries in RPM format, which enables easy installation using the yum package manager. The software can be downloaded, and release packages are available to configure the repositories for the user. This document assumes that the user has the required repositories already configured on the hosts. Refer to the official Eucalyptus documentation for more information.

Eucalyptus configuration

Eucalyptus supports a wide variety of deployment options, ranging from single-server demo environments to larger-scale production installations. Here, we outline an example of a fixed capacity deployment that provides a Eucalyptus cloud for an AWS-compatible, on-premise environment.

Eucalyptus additionally supports a wide variety of underlying hardware options. We have chosen a baseline set of server, networking, and storage devices in support of this deployment, but note that other HPE server/networking/storage options are available, which results in a similar architecture as defined here.

In terms of networking, this Eucalyptus deployment requires only two distinct IP networks (shown here as two separate Ethernet/IP networks), which may already exist in the environment into which this architecture is being deployed. If one or both networks does not yet exist in the environment, consider using HPE Networking's 5900AF family of switches as the networking substrate for either the public network, the management network, or both. To facilitate the deployment of hosts connected to two networks per the reference architecture, some hosts have 2 x 10 GbE network interface cards to allow for bonded configurations.

In the following example, the DL360p Gen9 server (with 8 x SFF disks) is used throughout for the Eucalyptus cloud service components and the Node Controller KVM hosts. Alternative hardware with a similar or higher specification could be used interchangeably, providing RHEL or CentOS are supported.

The Node Controllers act as KVM virtual hosts and have the most demanding workload. As such, they are configured with high-end 12-core Intel Xeon E5 processors with 128 GB of memory. Storage for the operating system is provided by a mirrored RAID set of 300 GB disks. Instances (Virtual Machines) use local storage with a RAID-5 configuration of 4 x 600 GB 10k SAS disks designed to reduce contention for I/O with the host system itself.

The Storage Controller and Cluster Controller are not required to have a large amount of local storage, and as such, they use 4 x 300 GB 10k SAS disks in RAID-5 for the operating system and the Eucalyptus software.

The User-Facing Services and Management Console are stateless components that rely on good CPU performance and memory capacity and do not utilize local storage for critical system artifacts.

The Cloud Controller hosts the Eucalyptus database, where read and write performance is important together with good CPU and memory performance. The database itself is not significantly large, and as such, 4 x 600 GB 10k SAS disks are used in RAID-10 for peak performance and to provide a level of resilience, with the operating system residing on a mirrored RAID-1 configuration of 2 x 300 GB 10k SAS disks.

Networking

The reference architecture is based on the Eucalyptus Edge networking mode. All hosts, apart from the Storage Controller and Cluster Controller, have two network interfaces on two networks. These networks may be different VLANs. The network configuration of Eucalyptus requires a range of IP addresses on the public (user-facing, routable) network. This pool of available IP addresses will be used by instances (virtual machines) running in the cloud. These will be the instance public IP addresses and should be routable from client networks. A range of available IP addresses is also required on the private network to be used for instance private IPs. Eucalyptus components require an IP address on the networks to which they are connected.

Networking configuration defines two interfaces on Node Controllers: one for external/public network traffic and access to the instance public IP addresses, and another for instance private IPs and management traffic between Eucalyptus control components.

10 GbE is recommended for both the public and private network. The private network will be used for instance-to-instance communication, iSCSI traffic for EBS volumes, and image downloads during instance launch. Accordingly, the performance of the network should keep in mind these performance profiles.

The network configuration for Eucalyptus Edge mode is defined in a JSON file. Use the JSON snippet below (minus comments), save it to a file, and substitute with your environment-specific information:

```
{
  "InstanceDnsDomain": "eucalyptus.internal",    <- Internal DNS domain (default)
  "InstanceDnsServers": ["10.112.0.2"], <- A comma-separated list of DNS servers
  "PublicIps": [
    "10.112.0.10-10.112.0.200" <- A range of IP's on the public network
  ],
  "Subnets": [
  ],
  "Clusters": [
    {
      "Name": "cluster01", <- must match the name of your registered cluster
      "MacPrefix": "d0:0d",
      "Subnet": {
        "Name": "10.111.0.0", <- user-friendly name for the subnet
        "Subnet": "10.111.0.0", <- the private network subnet
        "Netmask": "255.255.0.0", <- netmask of the private subnet, as above
        "Gateway": "10.111.0.1" <- In single cluster, set to an arbitrary IP
      },
      "PrivateIps": [
        "10.111.101.10-10.111.101.200" <- range of IP's on the private net
      ]
    },
  ]
}
```

Detailed physical network configuration is outside of the scope of this solution document.

Eucalyptus using HPE DL360p Gen9 ProLiant rack servers

The models of server chosen below are for reference only and may be substituted with equivalent models of similar specifications.

Table 1. Server configurations

	QUANTITY	SERVER MODEL	PROCESSOR	MEMORY	STORAGE	NETWORK
Cloud Controller	1	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 8 Core 1.8 Ghz Intel Xeon E5-2630Lv3	128 GB	Smart Array P440ar + 2 GB FBWC 4 x 600 GB SAS 10k in RAID-10 2 x 300 GB SAS 10k in RAID-1	2 x 10 Gb Dual Port NIC
User-Facing Services and Management Console	1	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 8 Core 1.8Ghz Intel Xeon E5-2630Lv3	128 GB	Smart Array P440ar + 2 GB FBWC 4 x 300 GB SAS 10k in RAID-5	2 x 10 Gb Dual Port NIC
Cluster Controller and Storage Controller	1	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 8 Core 1.8Ghz Intel Xeon E5-2630Lv3	128 GB	Smart Array P440ar + 2 GB FBWC 4 x 300 GB SAS 10k in RAID-5	2 x 10 Gb Dual Port NIC
Node Controller (KVM Virtual Hosts)	1-32	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 14 core 2.3 Ghz Intel Xeon E5-2695v3	128 GB	Smart Array P440ar + 2 GB FBWC 4 x 600 GB SAS 10k in RAID-5 2 x 300 GB SAS 10k RAID-1 for OS	2 x 10 Gb Dual Port NIC
Riak storage node	5	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 8 Core 1.8Ghz Intel Xeon E5-2630Lv3	128 GB	Smart Array P440ar + 2 GB FBWC 8 x 1.2 TB 10k	2 x 10 Gb Dual Port NIC
Riak load balancer	2	ProLiant DL360p Gen9 (8 x SFF) or equivalent	2 x 8 Core 1.8 Ghz Intel Xeon E5-2630Lv3	128 GB	Smart Array P440ar + 2GB FBWC 2 x 300GB SAS 10k in RAID-1 for OS	2 x 10Gb Dual Port NIC
3PAR 7400	1				28 TB	

A diagram of this configuration is shown below in Figure 2:

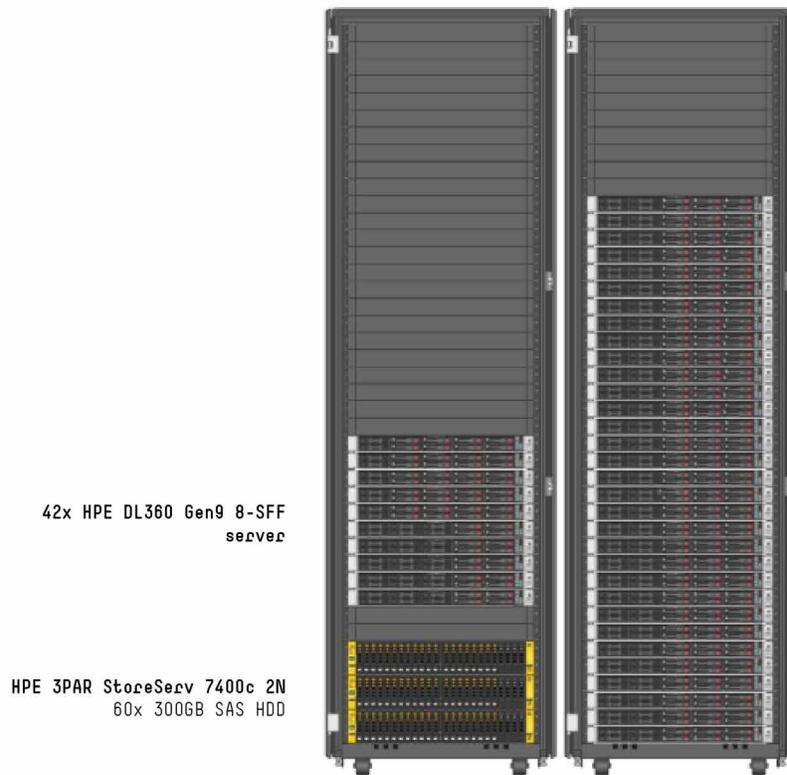


Figure 2. General purpose reference architecture (32 nodes with control components)

Additional clusters would be deployed as a single rack consisting of a further 32 nodes and a single cluster and storage controller, connected to the same 3PAR array for EBS.

A logical diagram of the deployment is shown in Figure 3:

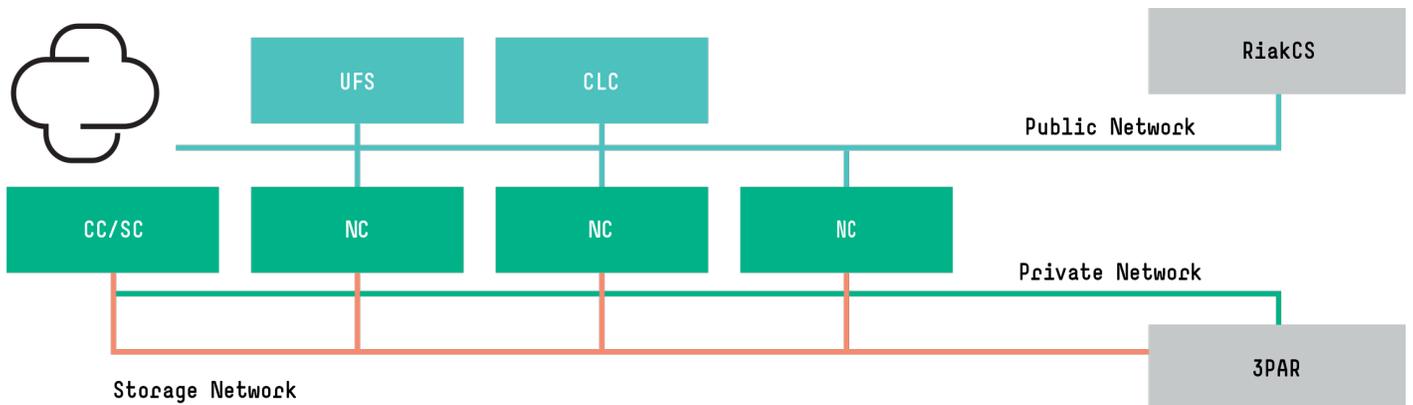


Figure 3. Eucalyptus Deployment Diagram

Capacity planning for Eucalyptus

For this deployment, there are a few choices to be made with regard to the maximum capacity of the cloud:

- Total amount of S3-compatible object storage available
- Total amount of EBS-compatible block storage available
- Total amount of running virtual machine instance capacity available

S3-compatible object storage capacity

To determine how much available AWS S3-compatible object storage is available to the entire cloud, the amount of local disk available to the RiakCS storage nodes should be adjusted. We have specified 8 x 1.2 TB drives per server, which would result in a total capacity. S3 object storage space is a cloud global service available to all users, where specific user/group quotas can be defined by the cloud administrator to limit the amount of global S3 storage space any particular user/group can use.

EBS-compatible block storage capacity

To determine how much available AWS EBS-compatible block storage is available to the entire cloud, the amount of local disk available to the RAID 5 volume on the Storage Controller (SC) server should be adjusted. We have specified 1.75 TB to start with, and this value can be adjusted based on the minimum and maximum amount of disk that the SC server can support. EBS block storage capacity is available to all users of the cloud, and users can allocate volumes in units of 1 GB. EBS volumes can be created and attached to virtual machines, or entire virtual machine root file systems can reside on EBS volumes. The cloud administrator can control how much EBS block storage is available on a per user/group basis using quotas.

Virtual machine instance capacity

The total virtual machine (VM) capacity of a given Eucalyptus cloud deployment is largely governed by the size and number of Node Controller (NC) servers that are available. AWS instances are minimally defined by the number of cores, RAM, and disk that is allocated for each instance. Because of this flexibility, determining the VM capacity of a given cloud deployment requires some calculation, as we can encounter cases where VMs can occupy resources in a way that leads to fragmentation.

We detail the ideal maximum capacity of the specified deployment per AWS instance type definition, but point out that in practice, placement of VMs can result in reduced maximum capacity when fragmentation occurs. In general, however, the overall VM capacity of a cloud can be controlled by the size and number of node controller servers that are part of the system, as well as the definition of the VM types (cores, memory, disk) that users are able to execute.

Here, we give maximum capacity based on current AWS instance type definitions; note that this table is not representative of the maximum capacity of Eucalyptus instance types since these are customizable by the cloud administrator. This table should be used to compare the relative VM capacity between AWS and a Eucalyptus deployment. Since the name, number of cores, amount of memory, and disk capacity per instance type is customizable by the cloud administrator, this should not be considered a table that describes the number of VMs of a given type that a fully configured deployment provides (which could result in more or less capacity for a given type depending on how the instance type is defined for the actual deployment).

Note

The values in this table assume 32 nodes and do not account for any form of overcommit in the environment. An increase in overcommit ratios for CPU and memory may be used to enhance compute capacity.

INSTANCE TYPE	INSTANCE TYPE PROPERTIES	PER NODE QUANTITY	TOTAL CLOUD QUANTITY
m1.small	vCPU=1 MEM=1.7 DISK=160	11	352
m1.medium	vCPU=1 MEM=3.75 DISK=410	4	128
c1.medium	vCPU=2 MEM=1.7 DISK=350	5	160
m2.xlarge	vCPU=2 MEM=17.1 DISK=420	4	128
m2.2xlarge	vCPU=4 MEM=34.2 DISK=850	2	64
t2.micro	vCPU=1 MEM=1 DISK=160	24	768
t2.small	vCPU=1 MEM=2 DISK=160	24	768
t2.medium	vCPU=2 MEM=4 DISK=160	12	384
m3.medium	vCPU=1 MEM=3.75 DISK=4	24	768
m3.large	vCPU=2 MEM=7.5 DISK=32	12	384
g2.2xlarge	vCPU=8 MEM=15 DISK=60	3	96
r3.large	vCPU=2 MEM=15.25 DISK=32	8	256
r3.xlarge	vCPU=4 MEM=30.5 DISK=80	4	128
r3.2xlarge	vCPU=8 MEM=61 DISK=160	2	64
r3.4xlarge	vCPU=16 MEM=122 DISK=320	1	32
i2.xlarge	vCPU=4 MEM=30.5 DISK=800	2	64
m1.small	vCPU=1 MEM=1.7 DISK=160	11	352

Learn more at
hpe.com/helion/eucalyptus



Sign up for updates

★ Rate this document



© Copyright 2015 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for HPE products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HPE shall not be liable for technical or editorial errors or omissions contained herein.

The OpenStack Word Mark is either a registered trademark/service mark or trademark/service mark of the OpenStack Foundation, in the United States and other countries and is used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community. Pivotal and Cloud Foundry are trademarks and/or registered trademarks of Pivotal Software, Inc. in the United States and/or other countries.

4AA6-2547ENW, November 2015