



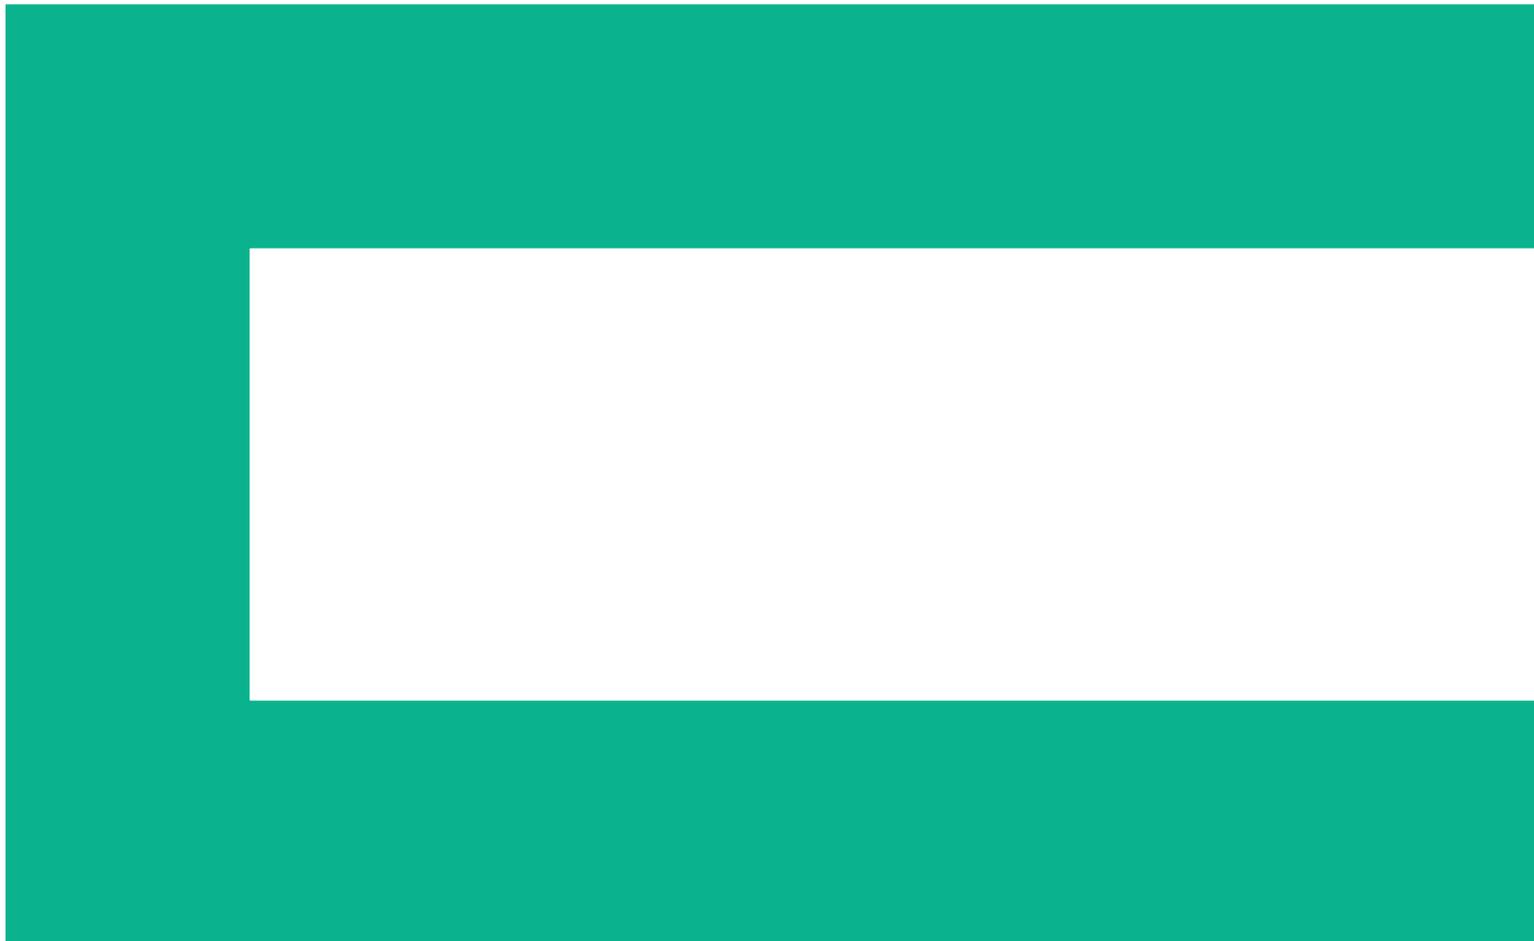
**Hewlett Packard
Enterprise**



Business white paper

HPE Composable Infrastructure

Bridging traditional IT with the Idea Economy



An IT era of unprecedented opportunity—and risk

Table of contents

3 The Idea Economy creates disruptive waves of new demands and opportunities

- 3 Time is not on the IT organization's side
- 4 Two ultimately incompatible ways to deliver applications
- 4 The challenge in moving forward
- 5 An infrastructure that serves both models
- 5 A new category of infrastructure is needed to power the Idea Economy
- 6 HPE Composable Infrastructure
- 7 Getting started
- 7 HPE Composable Infrastructure in action
- 8 Why Hewlett Packard Enterprise

CIOs and line-of-business executives are at the forefront of a major transformation aimed at leveraging the competitive advantages of the new hyper-connected enterprise. Industry research firm IDC describes the new IT for mobile devices, cloud services, social networks, and big data analytics as the third platform^{† 1}

At Hewlett Packard Enterprise, we call this evolution to the third platform the “Idea Economy” because it gives IT the ability to drive new business opportunities by quickly delivering revenue-generating products, services, and experiences. Instead of just providing technology to automate internal business processes, IT can now directly impact business strategy and revenues by creating software-based services that:

- Energize growth
- Strengthen profitability
- Boost productivity
- Enhance innovation
- Increase organizational agility
- Improve the customer experience
- Reduce risk

[†] Note: The first platform is defined as mainframe computer systems and the second platform as client/server systems.

¹ IDC, “The 3rd Platform: Enabling Digital Transformation,” November 2013.

The Idea Economy creates disruptive waves of new demands and opportunities

“Between 2013 to 2020, organizations’ investment in mobile, social, cloud, and big data technologies (third platform) will grow over 20 times faster than organizations’ investment in client/server technologies (second platform), with the third platform technologies accounting for 95% of the cumulative growth in investment.”²

– IDC predicts

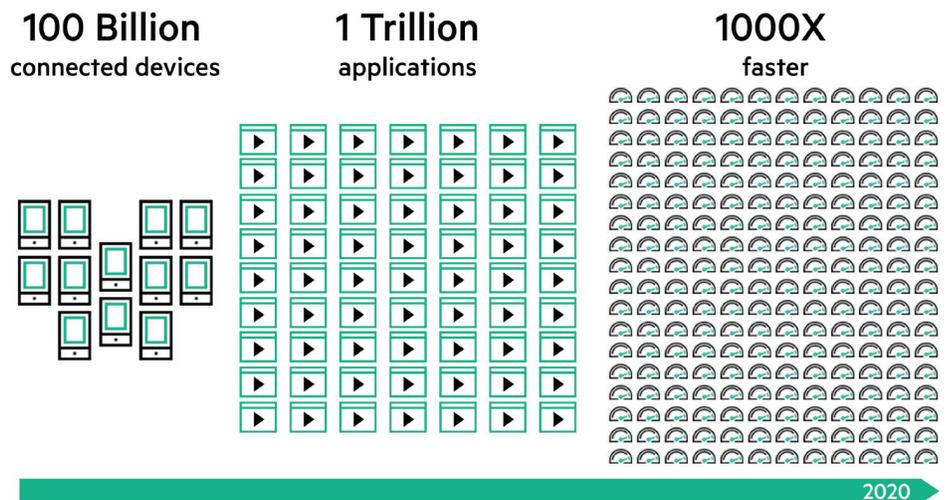


Figure 1: Thiving in the Idea Economy requires speed³

Time is not on the IT organization’s side

The pace of business is accelerating at an unprecedented rate. Today’s complex infrastructures make it difficult for IT to deliver on third platform priorities for a variety of reasons:

- **Update frequency:** Packaged applications release new revisions only a few times per year, limiting their ability to quickly respond to changing business needs.
- **Hardware centricity:** Technology vendors have developed infrastructure into rigid hardware and operational silos that are optimized for workloads versus applications and service delivery.
- **Fragmented management:** Infrastructure management software is layered on top of the hardware silos, adding to complexity. Often management scripts are used, which require continuous maintenance due to changing underlying components.

These disparate products and processes make change disruptive and require multiple points of manual coordination. The result is paralyzing digital sprawl: the proliferation of workload-specific hardware platforms, tools, and processes that undermine service delivery, speed, and quality as well as top-line business performance.

² IDC, “The 3rd Platform: Enabling Digital Transformation,” November 2013.

³ Source: HPE internal research. By 2020, more than a trillion applications will be exchanging 58 zettabytes of digital data over 100 billion devices, creating massive complexity and important opportunities for enterprises to transform their operations, enhance interactions with customers, and make workers more productive.

Two ultimately incompatible ways to deliver applications

Enterprises currently support two very different models for delivering applications. The first we'll call traditional (second platform) and the second the Idea Economy (third platform).

The traditional model runs packaged, back-office, and non-revenue generating applications designed to automate the business. Examples include enterprise resource planning (ERP), online transaction processing (OLTP), and communication and collaboration apps (e.g. email and Lync).

The traditional model operates in a conventional, steady-state environment that focuses on minimizing risk through standard methodologies, conventional vendors, and strong governance.

The Idea Economy model, on the other hand, operates mobile, big data, and cloud applications and services in a nonlinear, high-speed mode that requires low latency and can continuously adapt to high levels of uncertainty and change. The Idea Economy model is particularly adept at:

- Creating new sources of revenue and markets
- Monetizing data and services
- Enabling disruptive innovation
- Enriching end-user experiences
- Innovating, differentiating, and increasing customer loyalty
- Rapidly deploying new services
- Developing applications on a nearly continuous basis

The differences between these two models are numerous and fundamental.

The challenge in moving forward

IT leaders are between the proverbial rock and a hard place when it comes to both realizing the opportunities created by Idea Economy applications (designed to “be the business”) and maintaining and improving traditional applications (designed to “run the business”).

How can IT organizations balance the resources and investments needed to keep core legacy systems up and running while also creating value by integrating the new technologies needed to keep the business competitive and growing?

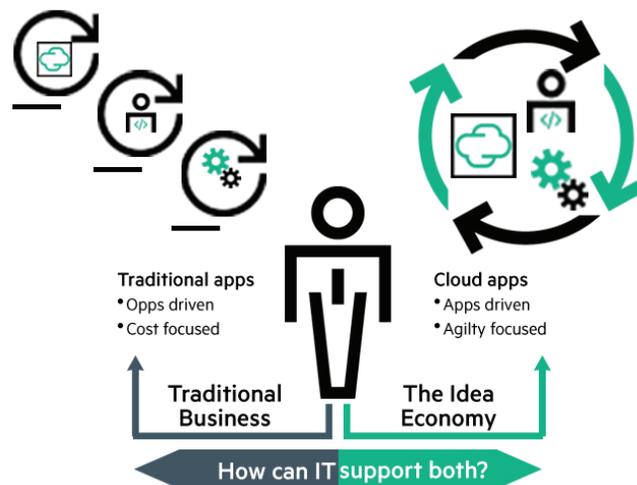


Figure 2: In the Idea Economy, IT is the business partner for value creation.

An infrastructure that serves both models

A two-mode approach is needed: one that allows IT organizations to align the infrastructure around the needs of both traditional and Idea Economy applications. For the Idea Economy’s nonlinear, high-speed needs, it would allow IT to implement the most advanced technology in areas that demand speed and flexibility, such as exploiting big data, managing unstructured data, and working with cloud applications and services. For the traditional applications, it would eliminate the drag of resource-intensive applications on daily operations, protecting their uninterrupted performance in the conventional, steady-state mode. It would also ease their ultimate transition to the cloud.

Business demands for performance and agility are already well beyond the ability of traditional, siloed infrastructures. Enterprises are adopting a two-mode IT approach to close the value gap by fast-tracking solutions and generating more value in less time. This approach is faster, more open, and that can continuously deliver applications and services.

A new category of infrastructure is needed to power the Idea Economy

A new architecture is needed – one that is designed to power innovation and value creation for the new breed of applications while running traditional workloads more efficiently. We call this new architectural approach Composable Infrastructure.

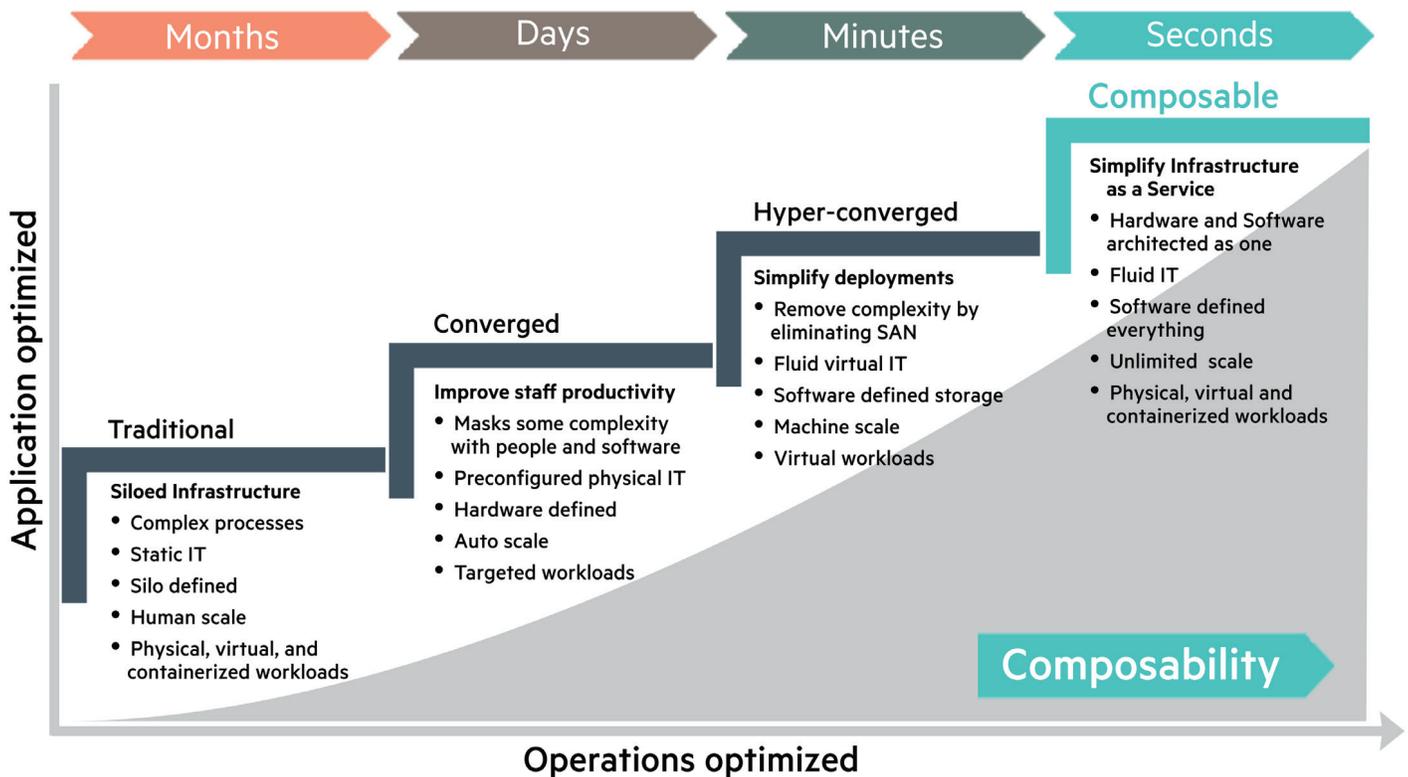


Figure 3: A new class of infrastructure is needed to power traditional and new applications

Run anything

Optimize any application and store all data on a single infrastructure with fluid pools of physical and virtual compute, storage, and fabric.

Move faster

Accelerate application and service delivery through a single interface that precisely composes logical infrastructures at near-instant speeds.

Work efficiently

Reduce operational effort and cost through internal software-defined intelligence with template-driven, frictionless operations.

Unlock value

Increase productivity and control across the data center by integrating and automating infrastructure operations and applications through a unified API.

HPE Composable Infrastructure

Hewlett Packard Enterprise has a well-established track record of infrastructure innovation and success. HPE Converged Infrastructure, software-defined management, and hyper-converged systems have consistently proven to reduce costs and increase operational efficiency by eliminating silos and freeing available compute, storage, and networking resources. Building on our converged infrastructure knowledge and experience, we have designed a new architecture that can meet the growing demands for a faster, more open, and continuous infrastructure.

We call this next-generation, two-mode approach composable infrastructure. It is designed around three core principles:



Fluid resource pools effortlessly meet each application's changing needs by allowing for the composition and recomposition of single blocks of disaggregated compute, storage, and fabric infrastructure.



Software-defined intelligence provides a single management interface to integrate operational silos and eliminate complexity. Workload templates speed deployment and frictionless change eliminates unnecessary downtime.



Unified API provides a single interface to discover, search, inventory, configure, provision, update, and diagnose the composable infrastructure. A single line of code enables full infrastructure programmability and can provision the infrastructure required for an application.

This new class of unified, two-mode infrastructure facilitates the move to a continuous services and application delivery model and enables applications to be updated as needed, rather than just once or twice a year.

Composable infrastructure enables IT to operate like a cloud provider to lines of business and the extended enterprise. It maximizes the speed, agility, and efficiency of core infrastructure and operations to consistently meet SLAs and provide the predictable performance needed to support core workloads—for both today and tomorrow.

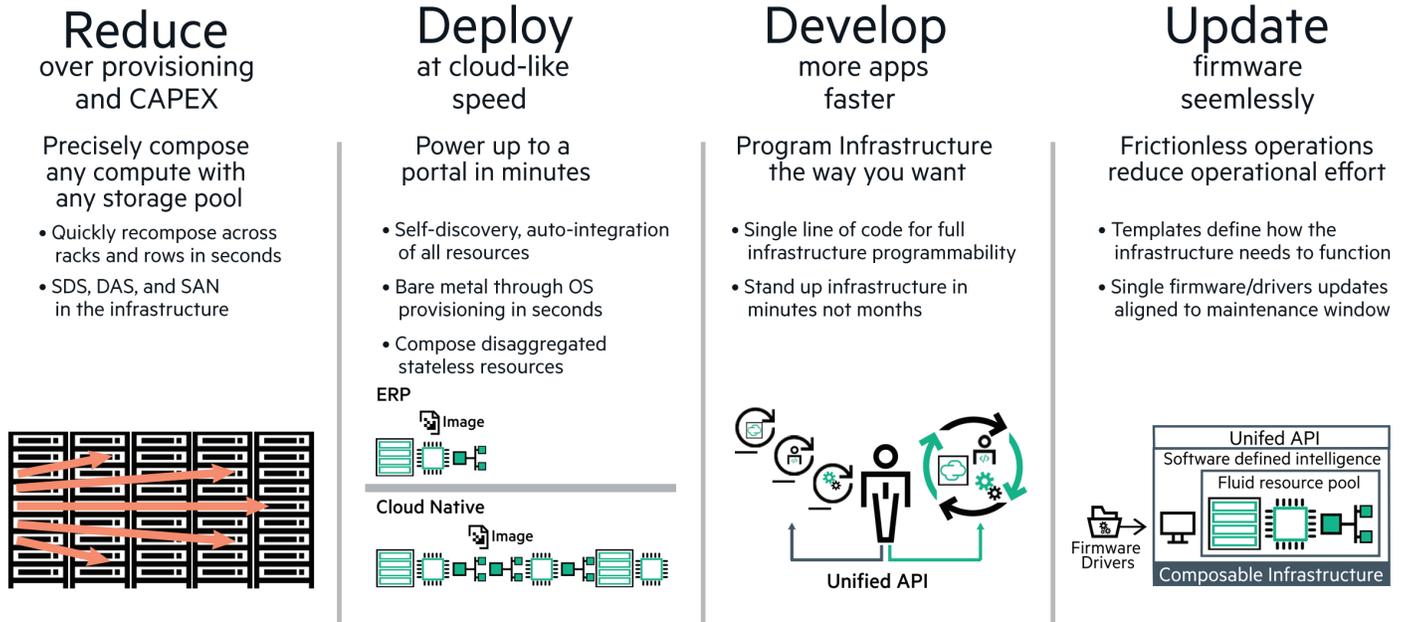


Figure 4: Composable infrastructure is the engine for the Idea Economy

Getting started

The journey to composable infrastructure can be started today with converged blocks that have composable attributes such as software defined automation and the unified API from HPE OneView:

- HPE BladeSystem/ConvergedSystem with HPE 3PAR
- HPE ProLiant with HPE StoreVirtual VSA.

These products enable continuity and investment protection of tools, integration, and processes.

HPE Composable Infrastructure in action

HPE Composable Infrastructure lets IT administrators and developers use infrastructure as code to control their internal environments. The unified API integrates dozens of popular management tools. Here’s an example of how it works using Chef™, a powerful automation tool that enables the rapid and reliable deployment and updating of infrastructure using recipes that can be versioned and tested just like application software. Chef needs an infrastructure to direct, such as a virtual machine, public cloud, or virtual cloud instance. The Composable Infrastructure API powered by HPE OneView creates, aggregates, and hosts internal IT resources that Chef can provision on-demand and pragmatically, without needing a detailed understanding of the underlying physical elements. By connecting Chef tools with HPE OneView, bare metal infrastructure can be directed in the same way as virtual and public cloud resources. To get started, HPE Datacenter Care-Infrastructure Automation (DC-IA) delivers enterprise-grade support, advice, guidance and best practices for infrastructure automation. The service includes enterprise editions of selected automation tools, including Enterprise Chef. The DC-IA Center of Excellence (CoE) is staffed with highly trained professionals with expertise in creating and managing an automated infrastructure and offers code coaching and script reviews.

Contact HPE

Program info:

ComposableAPIprogram@hpe.com

Technical support:

ComposableAPIsupport@hpe.com

Learn more at

HPE Composable Infrastructure
hpe.com/info/composable

HPE Composable Infrastructure
Partner Program
hpe.com/info/composableprogram

HPE Datacenter Care Services
hpe.com/services/dc

Provisioning with Chef is just one example of how composable infrastructure offers infrastructure as code. The same concepts allow:

- Other automation tools like Ansible™ and Puppet to provision bare metal infrastructure
- VMware® and Microsoft® technologies to create and flex virtualization clusters
- HPE Helion® to provision bare metal hosts or virtualization clusters for multi-tiered applications via the OpenStack® Ironic project
- Docker™ to deploy infrastructure to host containers

Developers and ISVs can programmatically control a composable infrastructure through a single, open API that is native in HPE OneView. By integrating with the HPE Composable Infrastructure API, they can support customer requirements for both traditional and Idea Economy IT environments. Continuous delivery of applications and services requires fast, policy-based automation of applications and infrastructure across development/testing/production environments. The HPE Composable Infrastructure API enables developers to integrate with development/testing/production automation tool chains and to drive a more aligned and responsive delivery of IT services.

Why Hewlett Packard Enterprise

While some technology vendors have begun talking about “composability” or “infrastructure as code,” none has our completeness of vision or ability to execute across the full compute, storage, and fabric infrastructure. Our approach to composable infrastructure is compelling for a number of reasons:

- Open Composable Infrastructure APIs that are unified across the entire infrastructure, including compute, fabric, and storage
- Runs virtually any application and stores virtually any data
- Composable physical, virtual, and container resource pools
- Composable data services (block, file, object are based on application requirements) with best-fit placement and optimized quality of service
- Support for both flash-optimized systems and software-defined storage with non-disruptive workload mobility via storage federation
- Frictionless change operations
- The Composable Infrastructure Partner Program provides a set of tools and resources that enables ISVs and developers to build interoperability between HPE OneView and other software for programmatic access to infrastructure
- Enterprise-grade support with single point of contact for access to a global network of specialists, 24x7
- HPE Flexible Capacity—A pay per use model for on premise infrastructure, providing the needed HPE Composable Infrastructure capacity for the datacenter with a buffer of additional capacity to use when needed.

Only HPE can bring the hardware, software, services, and partner ecosystem needed to deliver an end-to-end offering with a single point of accountability.



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 Hewlett Packard Enterprise 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. Hewlett Packard Enterprise 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. The HPE Helion Network will be launched in phases, and not all phases are available at this time. Details and features of the HPE Helion Network program are subject to change from time to time without notice. Microsoft is a U.S. registered trademark of the Microsoft group of companies. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

4AA5-8813ENW, December 2015, Rev. 2