HPE Reference Configuration for Red Hat OpenShift Virtualization on HPE ProLiant DL325 and DL385 Gen11 servers using Red Hat OpenShift Container Platform 4.15

Rapid deployment of Red Hat OpenShift Virtualization on HPE ProLiant DL325 and DL385 Gen11 servers using Red Hat OpenShift Container Platform 4.15 automated approach

CONTENTS

Document type

EXECUTIVE SUMMARY

Enterprise organizations across all industries are embarking on a hybrid cloud journey. To support digital transformation, business innovation, and accelerated growth, organizations have certain key goals. Some of the predominant goals include speed, agility, simplicity, consistency, and cost-effectiveness.

However, current IT practices and various incompatible application deployment environments have created challenges for organizations to achieve these objectives. Some of the key challenges are as follows:

- Modernizing legacy apps to take advantage of the latest agile cloud-native innovations is difficult and time-consuming.
- Managing workloads that span multiple cloud environments is challenging.
- Provisioning a new environment is a slow process and can significantly stifle innovation as teams have to wait for the environment to be available.
- Vendor lock-in is a real concern, especially with but not limited to public cloud providers.
- Siloed infrastructure increases overhead costs including administrative overhead in addition to the price of additional infrastructure.
- Deploying a disconnected and secured end-to-end container platform quickly.

To unleash business opportunities through digital transformation, enterprises must overcome these restrictions and adapt to the cloud-native design principles and solutions of the next-generation IT practices. Hewlett Packard Enterprise and Red Hat® are collaborating to optimize Infrastructure by running and manage virtual machine workloads alongside container workloads Red Hat OpenShift Virtualization is an add-on to Red Hat **®** OpenShift® Container Platform 4.15 that leverages the HPE ProLiant DL325 and DL385 Gen11 servers to accelerate container application delivery.

IDC interviewed nine organizations that use Red Hat® OpenShift as their primary application development platform, finding that the solution yields significant business value. This executive summary shows the highlights from the report, which include a five-year ROI of 531% and average annual savings of $1.29 million per 100 developers.

Red Hat OpenShift Virtualization (RHOV) is an included feature of Red Hat OpenShift, providing a modern platform for organizations to run and deploy their new and existing virtual machine (VM) workloads. Here are some key points about it:

- Easy Migration: Red Hat OpenShift Virtualization includes a simple way to migrate existing virtual machines from other hypervisors using the included Migration Toolkit for Virtualization. You can even migrate VMs to the cloud. Red Hat Services also provides mentor-based consulting to support your migration process.
- Speed Up Time to Production: By streamlining infrastructure and application delivery, Red Hat OpenShift Virtualization allows developers to build, test, and deploy workloads faster, accelerating time to market.
- Unified Platform: Red Hat OpenShift Container Platform with OpenShift Virtualization simplifies operations by providing a single platform for VMs, containers, and serverless workloads. This standardizes infrastructure deployment and allows you to maintain all workloads using a common set of established enterprise tools.
- Infrastructure Modernization: When migrating VMs from other platforms to Red Hat OpenShift, you can maximize your existing virtualization investments while embracing cloud-native architectures, streamlined operations, and new development approaches.
- Windows Application Modernization: Red Hat OpenShift Virtualization fully supports Linux and Windows images, making it an excellent choice for migrating Windows applications to the platform. You can then modernize these applications by breaking them down into smaller pieces and deploying them as containers.

In summary, Red Hat OpenShift Virtualization enables you to run VMs alongside containers on the Red Hat OpenShift Container Platform, simplifying management and improving overall efficiency. It is a powerful solution for organizations looking to modernize their virtualization strategy while leveraging the benefits of a cloud-native application platform.

This Reference Configuration provides architectural guidance for deploying Red Hat OpenShift Virtualization on Red Hat OpenShift Container Platform 4.15 and HPE ProLiant DL325 Gen11 servers for Compute. The compute requirements can easily be scaled by adding more HPE ProLiant DL325 Gen11 servers with no workload downtime.

The Cloud Native Computing Foundation (CNCF) Operator Framework in this solution provides a cloud-native method of packaging, deploying, and managing Kubernetes-native applications that include:

1. Set up HPE ProLiant DL325 Gen11 servers.
2. To install and configure the Red Hat OpenShift Container Platform 4.x Cluster and Red Hat OpenShift Virtualization feature.
3. Validate the Red Hat OpenShift Virtualization and Red Hat OpenShift Container Platform installation.
4. Create, manage, or delete virtual machines on Red Hat OpenShift Virtualization
5. Migrate virtual machines from VMware or other supported providers to Red Hat OpenShift Virtualization

Significant reduction in the deployment time and efforts through the automated deployment process.

The Reference Configuration demonstrates a cost-effective yet reliable solution by leveraging the benefits of HPE ProLiant DL325 and DL385 Gen11 servers for compute, storage, networking, and Red Hat OpenShift Container Platform 4.15.

Target audience: This document is intended for Chief Information Officers (CIOs), Chief Technology Officers (CTOs), data center managers, enterprise architects, and implementation personnel who wish to learn more about Red Hat OpenShift Container Platform 4.x on HPE ProLiant DL325 Gen11 servers. This document assumes that the reader is familiar with HPE ProLiant DL325 Gen11 servers, Red Hat OpenShift Container Platform 4.15, core networking, and has a valid Red Hat OpenShift Container Platform Subscription.

Document purpose: This document describes the benefits and technical details of deploying Red Hat OpenShift Virtualization on Red Hat OpenShift Container Platform 4.15 using HPE ProLiant DL325 Gen11 and HPE ProLiant DL385 servers, the implementation details, and the processes. This guide is accompanied by a Deployment Guide which can be found at https://hewlettpackard.github.io/hpe-solutions-openshift/4.15-AMD-LTI/.

INTRODUCTION

This Reference Configuration provides guidance for installing & configuring Red Hat OpenShift Virtualization in Red Hat OpenShift Container Platform 4.15 (the solution), on HPE ProLiant DL325 and HPE ProLiant DL385 Gen11 servers. The solution consists of six (6) HPE ProLiant DL325 Gen11 servers: three (3) HPE ProLiant DL325 Gen11 servers used for the Red Hat Enterprise Linux (RHEL) KVM-based Head Nodes and three (3) HPE ProLiant DL325 Gen11 servers used for the solution worker nodes (out of which one node is used as a temporary Red Hat OpenShift Container Platform installer node). HPE ProLiant DL385 Gen11 servers can be added as a cluster for Red Hat OpenShift Data Foundation (ODF).

The persistent storage for this solution is provided by HPE Alletra 6k series storage array. For business-critical workloads, HPE Alletra 6070 delivers fast, consistent performance and industry-leading data efficiency.

SOLUTION OVERVIEW

This section provides an overview of the design and configuration of the solution. Figure 1 shows the high-level architecture of the solution.

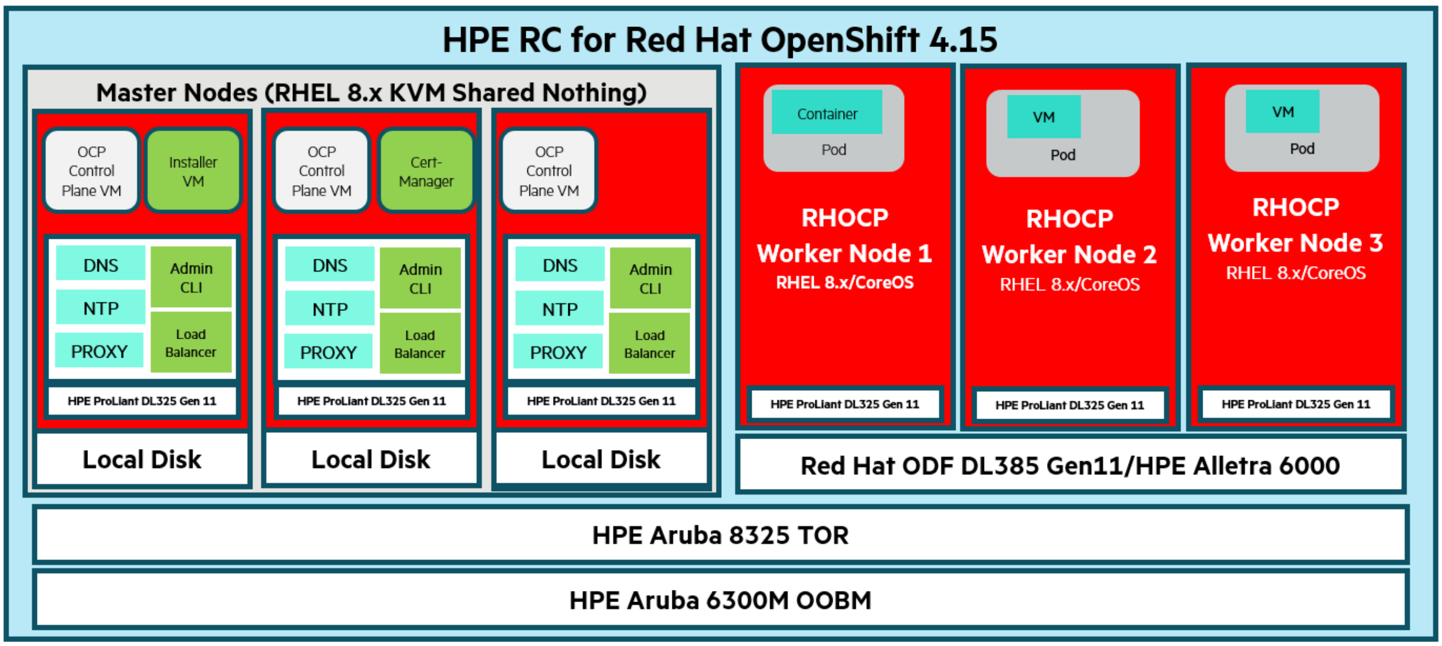

FIGURE 1. High-level architecture

This solution uses the Red Hat OpenShift User Provisioned Infrastructure method of installation to install Red Hat Enterprise CoreOS (RHCOS) and Red Hat Enterprise Linux® (RHEL) 8.9 on the HPE ProLiant DL325 Gen11 servers and configure the Red Hat OpenShift Container Platform cluster.

Design objectives

The objective of this Reference Configuration is to provide guidance that allows Hewlett Packard Enterprise customers to deliver value by providing a performance-oriented yet cost-effective solution offering for the Red Hat OpenShift Container Platform. HPE ProLiant DL325 and HPE ProLiant DL385 Gen11 servers and HPE Alletra Storage provide an intelligent foundation that delivers workload optimization, security, and automation.

Physical configuration

This solution uses a hybrid infrastructure configuration approach. The Red Hat OpenShift Container Platform Control Plane nodes are deployed as KVM virtual machines running Red Hat Enterprise CoreOS. These virtual machines are running RHEL8.9 and KVM on three (3) HPE ProLiant DL325 Gen11 servers. Three (3) HPE ProLiant DL325 Gen11 servers are deployed as solution worker nodes on bare metal. The temporary installer node is deployed on one of the worker nodes and later configured as a worker node. HPE ProLiant DL385 Gen11 servers can be added as a cluster for Red Hat OpenShift Data Foundation (ODF).

The solution uses the internal storage on the HPE ProLiant DL325 Gen11 servers for both the operating system and solution applications. The environment infrastructure support components (Installer machine, iPXE, DNS, DHCP, etc.) and a load balancer in this solution are deployed on virtual machines. The Red Hat OpenShift-installer tool is run to generate ignition files that contain information about the hosts that will be provisioned. The Red Hat Enterprise CoreOS for the nodes is then booted with the help of iPXE and the ignition files are passed with the operating system image during installation. HPE ProLiant DL325 Gen11 servers use HPE Alletra 6070 via iSCSI to provide persistent container volume for the solution application workload.

The rack diagram of the hardware components that form the solution is shown in Figure 2.

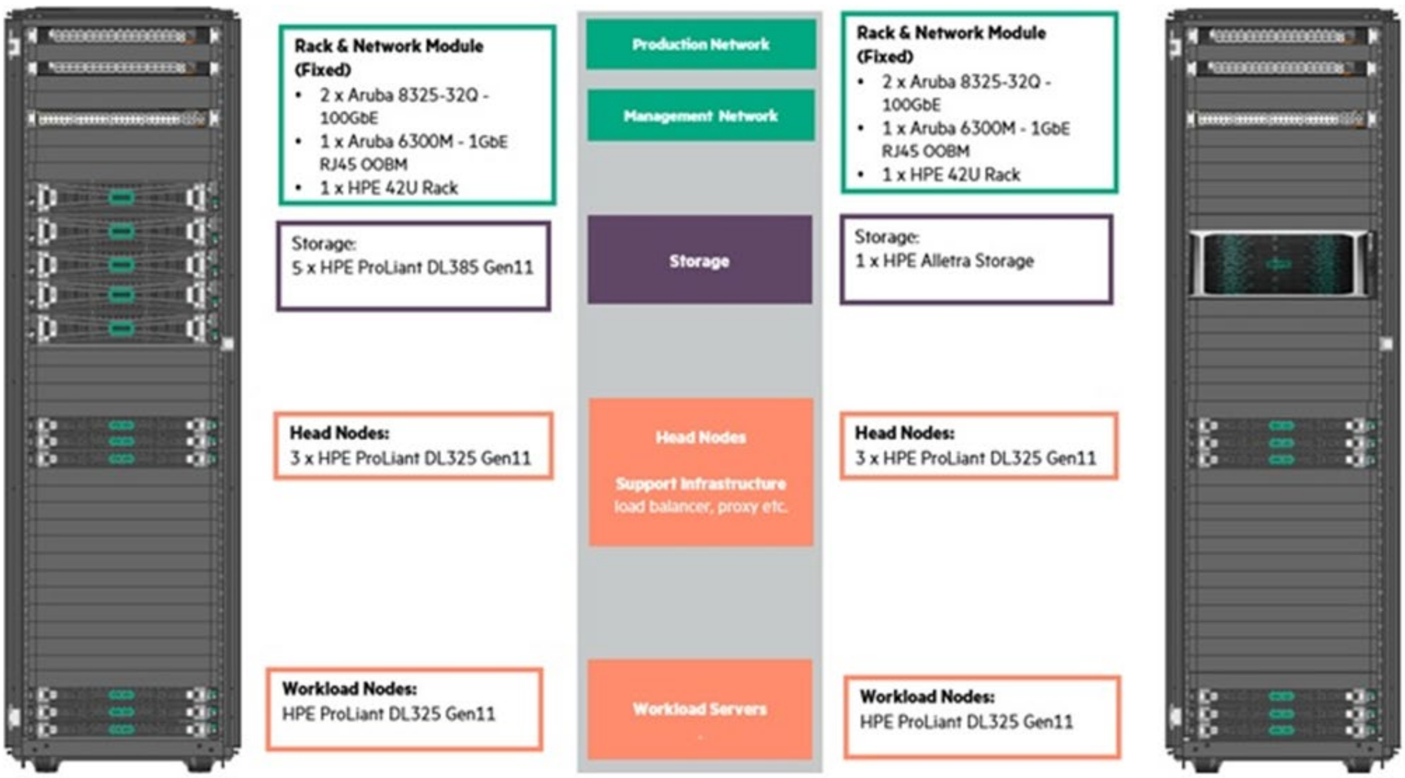

FIGURE 2. Solution components

NOTE

The figure depicts the hardware layout for the base configuration with three solution worker nodes and it is scalable.

Additional HPE ProLiant DL325 Gen11 servers can be added to this solution as per the customer's choice of configuration workload options.

SOLUTION COMPONENTS

This section provides the details of the hardware and software components used in the solution.

Hardware Components

Table 1 lists the various hardware components used in the solution.

TABLE 1. Hardware components utilized in this solution.

Component	Qty	Description
HPE ProLiant DL325 Gen11 servers	3	Provides capacity for head nodes with OpenShift control plane and bootstrap KVM VMS, HAProxy, DNS, Proxy
HPE ProLiant DL325 Gen11 servers	3	Provision as Red Hat OpenShift nodes
HPE Alletra 6070	1	External iSCSI storage for Persistent Volumes
HPE ProLiant DL385 Gen11 servers	5	Red Hat OpenShift Data Foundation nodes – Internal storage mode (optional)
HPE Aruba 8325 switch	2	A network switch for datacenter network
HPE Aruba 6300M switch	1	A network switch for iLO Management network

Hardware configuration

Table 2 lists the various hardware configurations used in this solution.

TABLE 2 Hardware configuration

Node	Operating System	vCPU	RAM	Storage
KVM Head Nodes	RHEL 8.9	64	512 GB	OS Disk: 2x 1.6 TB Data Disk: 2x 1.6 TB
Workload Nodes	RHEL 8.9 / RHCOS	64	512 GB	OS Disk: 2x 1.6 TB Data Disk: 2x 1.6 TB
Red Hat OpenShift Data Foundation Nodes	RHCOS	64	384 GB	OS Disk: 2x 1.6 TB Data Disk: 6x 3.2 TB

NOTE
For Red Hat OpenShift Virtualization worker nodes, it is recommended to have worker nodes with Red Hat Enterprise Linux CoreOS (RHCOS).

Red Hat OpenShift server roles configuration

Table 3 lists the various server roles and their configuration used in this solution.

TABLE 3. Server roles and configuration

Node	Operating System	vCPU	Virtual RAM	Storage
Bootstrap node	RHCOS	4	16 GB	120 GB
Control plane nodes	RHCOS	8	16 GB	250 GB

NOTE
The HAProxy load balancer was deployed on the KVM head node servers.

HPE ProLiant DL325 Gen11 server

The HPE ProLiant DL325 Gen11 server is a low-cost 1U 1P solution that delivers exceptional value balancing compute, memory, and network bandwidth at 1P economics. Powered by 4th Generation AMD EPYC™ Processors with up to 128 cores, increased memory bandwidth (up to 3 TB), high-speed PCIe Gen5 I/O and EDSFF storage, and supporting up to 2 GPUs at the front, this server is a superb low-cost, 1U 1P, performance solution for your virtualized workloads. The silicon root of trust anchors the server firmware, creating a fingerprint for the AMD Secure Processor that must be matched exactly before the server boot. The HPE ProLiant DL325 Gen11 server is an excellent choice for virtualized workloads such as software-defined compute, CDN, VDI, and secure edge apps that require balancing processor, memory, and network bandwidth.

Figure 3 shows the HPE ProLiant DL325 Gen11 server.

FIGURE 3. HPE ProLiant DL325 Gen11 server

Table 4 lists the hardware configuration used in this solution.

TABLE 4 Hardware configuration in each of the HPE ProLiant DL325 Gen11 servers

Component	Description
Processor	1x AMD EPYC 9554P 3.1GHz 64-core 360W Processor for HPE
Memory	12x 64GB (1x64GB) Dual Rank x4 DDR5-4800
Network	1x Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE 1x Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
Smart Array Controller	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
Disks	4x HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD

HPE ProLiant DL385 Gen11 server

The HPE ProLiant DL385 Gen11 server is an accelerator-optimized 2U 2P solution that delivers exceptional compute performance, upgraded high-speed data transfer rate and memory depth at 2P compute capability. Powered by 4th Generation AMD EPYC™ 9004 Series Processors with up to 96 cores, increased memory bandwidth (up to 6TB), high-speed PCIe Gen5 I/O, Gen5 EDSFF storage and the newly designed chassis supporting 8 single wide (SW) or 4 double wide (DW) GPUs. The HPE ProLiant DL385 Gen11 server is a perfect accelerator-optimized 2U 2P solution.

Figure 4 shows the HPE ProLiant DL385 Gen11 server.

FIGURE 4. HPE ProLiant DL385 Gen11 server

Table 5 lists the hardware configuration in each of the HPE ProLiant DL385 Gen11 servers used in this solution.

TABLE 5. Hardware configuration in each of the HPE ProLiant DL385 Gen11 servers

Component	Description
Processor	2 x AMD EPYC 9224 2.5GHz 24-core 200W Processor for HPE
Memory	24x HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
Network	1x Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE 1x Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
Array Controller	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
Disks	2x HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD 8x HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD

HPE Alletra 6K

HPE Alletra powers your data from edge-to-core with the cloud experience for all your apps. For business-critical workloads, HPE Alletra 6000 delivers fast, consistent performance and industry-leading data efficiency. It enables IT to shift from owning and maintaining data infrastructure to simply accessing and utilizing it on-demand, as-a-service. Eliminate performance and efficiency trade-offs with no knobs or configurations to adjust and always-on data services. Get resilient storage with intelligence and a no single point of failure platform that together deliver 6-nines availability guaranteed. Deliver recovery SLAs with fast, integrated app aware backup and recovery—on-premises and in the cloud.

FIGURE 5. HPE Alletra 6K

HPE iLO

HPE Integrated Lights Out (iLO) is embedded in HPE ProLiant DL325 and DL385 Gen11 platforms and provides server management that enables faster deployment, and simplified lifecycle operations while maintaining end-to-end security, thus increasing productivity.

HPE Aruba 8325-32C BF switch

The HPE Aruba CX 8325 switch is an enterprise-class, game-changing solution, offering a flexible approach to dealing with the new application, security, and scalability demands of the mobile, cloud, and IoT era. It provides the following benefits:

- Simplify your IT operations with AOS-CX
- Accelerate IT provisioning.
- Unparalleled visibility and analytics
- No downtime, even during upgrades

Figure 6 shows the HPE Aruba 8325-32C BF switch.

FIGURE 6. HPE Aruba 8325-32C BF switch

HPE Aruba CX 6300M OOBM switch

The HPE Aruba CX 6300 switch series is a modern, flexible, and intelligent family of AOS-CX stackable switches ideal for access, aggregation, and data center top-of-rack (TOR) deployments. With a cloud-centric design that combines a fully programmable OS with the HPE Aruba Network Analytics Engine, the HPE Aruba CX 6300 extends industry-leading monitoring and troubleshooting capabilities to the access layer. Support of Aruba Net Edit and the Aruba CX Mobile App verifies that configurations are flawless and easy to deploy.

A powerful HPE Aruba Gen7 ASICs architecture delivers fast, non-blocking performance, meaning your network is ready for tomorrow's unpredictable demands. HPE Aruba Virtual Stacking Framework (VSF) allows for the stacking of up to ten switches, providing scale and simplified management. This flexible series has built-in high-speed uplinks and supports high-density IEEE 802.3bt high-power PoE with HPE Smart Rate multi-gigabit Ethernet for high-speed APs and IoT devices.

Figure 7 shows the HPE Aruba 6300M OOBM switch.

FIGURE 7. HPE Aruba CX 6300M OOBM switch

Software components

Red Hat OpenShift Container Platform

Red Hat OpenShift Container Platform unites developers and IT operations on a single platform to build, deploy, and manage applications consistently across hybrid cloud and multi-cloud infrastructures. Red Hat OpenShift helps businesses achieve greater value by delivering modern and traditional applications with shorter development cycles and lower operating costs. Red Hat OpenShift is built on open-source innovation and industry standards, including Kubernetes and Red Hat Enterprise Linux.

Red Hat Enterprise CoreOS

Red Hat OpenShift Container Platform uses Red Hat Enterprise CoreOS (RHCOS), a new container-oriented operating system that combines some of the best features and functions of the CoreOS and Red Hat Atomic Host operating systems. Red Hat Enterprise CoreOS is specifically designed for running containerized applications from the Red Hat OpenShift Container Platform and works with new tools to provide fast installation, operator-based management, and simplified upgrades. For Red Hat OpenShift Container Platform 4.15 deployment on bare metal infrastructure, you must use Red Hat Enterprise CoreOS for all Red Hat OpenShift Container Platform control plane nodes, Bootstrap nodes, and worker nodes.

HPE Alletra Container Storage Interface

The HPE Container Storage Interface (CSI) Driver is a multi-vendor and multi-backend driver where each implementation has a Container Storage Provider (CSP). The HPE CSI Driver allows any vendor or project to develop its own Container Storage Provider (CSP) by using the CSP specification. This makes it very easy for third parties to integrate their storage solutions into Kubernetes as all the intricacies are taken care of by the HPE CSI Driver. The CSI specification includes constructs to manage snapshots as native Kubernetes objects and create a new Persistent Volume Claim (PVC) by referencing those objects. Other capabilities include PVC expansion, inline ephemeral volumes, and the ability to present raw block storage to pods.

Red Hat OpenShift Data Foundation

Red Hat OpenShift Data Foundation is software-defined storage that is optimized for container and virtual machine environments. It runs as an operator on Red Hat OpenShift Container Platform 4.15 to provide highly integrated and simplified persistent storage management for containers. Red Hat OpenShift Data Foundation supports a variety of storage types, including block storage for virtual machines and databases, shared file storage for continuous integration, messaging, and data aggregation, and object storage for archival, backup, and media storage.

Red Hat OpenShift Virtualization

Red Hat® OpenShift® Virtualization, an included feature of Red Hat OpenShift, provides a modern platform for organizations to run and deploy their new and existing virtual machine (VM) workloads. The solution allows for easy migration and management of traditional virtual machines onto a trusted, consistent, and comprehensive hybrid cloud application platform. Red Hat OpenShift Virtualization offers a path for infrastructure modernization, taking advantage of the simplicity and speed of a cloud-native application platform and aims to preserve existing virtualization investments while embracing modern management principles.

Table 6 lists the major software used in this solution.

TABLE 6. Software used in this solution.

Component	Versions	Usage
Red Hat Enterprise Linux CoreOS	4.x	Red Hat OpenShift control plane VMs
Red Hat OpenShift Container Platform	4.15	Red Hat OpenShift control plane nodes on KVM virtual machines and bare metal worker nodes
Red Hat Enterprise Linux	8.9	KVM head node and Red Hat worker nodes bare metal
Red Hat Enterprise Linux	8.9	Installer Machine required to execute automation scripts
HPE Alletra 6K	6.0.0.300-956221-opt	External storage
Red Hat OpenShift Data Foundation	4.x	Internal Storage
Red Hat OpenShift Virtualization	4.x	Red Hat OpenShift Virtualization feature on worker nodes

KUBEVIRT OVERVIEW

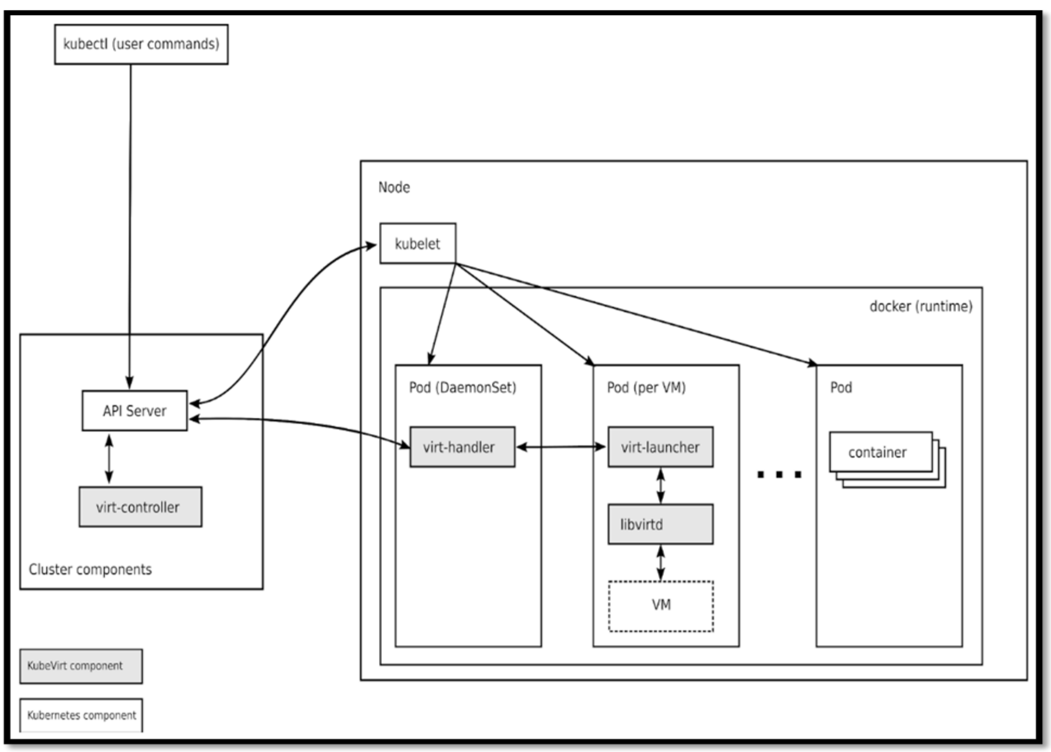

FIGURE 8. KubeVirt Architecture

KubeVirt is a technology that enables developers to run virtual machines (VMs) on Kubernetes platforms. As the trend toward containerized application stacks in orchestration platforms like Kubernetes grows, there will always be workloads that defy easy containerization. KubeVirt was designed specifically for these scenarios, offering a unified development platform where developers can create, manage, deploy, and scale applications—whether their components are containerized, virtualized, or a combination of both.

KubeVirt is provided by Red Hat as a feature of Red Hat OpenShift Container Platform called Red Hat OpenShift Virtualization. Red Hat OpenShift Virtualization is a feature in Red Hat OpenShift Container Platform that allows you to run and manage virtual machine workloads alongside container workloads this done using KubeVirt. Components that include the following:

- virt-controller: A cluster component that watches for VirtualMachine CRDs (Custom Resource Definitions).
- virt-api: The interface to the controller.
- virt-handler: A DaemonSet that runs on each node. It manages VMs by embedding them in pods running a hypervisor (libvirtd).
- VMs run within regular Kubernetes pods, giving them access to standard pod networking and storage.
- Red Hat OpenShift VMs can be managed using standard Kubernetes tools such as kubectl or 'oc' tool.
- KubeVirt introduces additional types (CRDs) to the Kubernetes API that include VirtualMachine and VirtualMachineInstance. Users interact with Kubernetes but modify VMIs (VirtualMachineInstances).

Note: Both controllers and daemons run as pods on top of the Kubernetes cluster refer

Red Hat OpenShift Virtualization is built on the KubeVirt project, which extends Kubernetes by adding virtualization resource types through Custom Resource Definitions (CRDs)

WHAT YOU CAN DO WITH RED HAT OPENSHIFT VIRTUALIZATION

Red Hat OpenShift Virtualization is an add-on to Red Hat OpenShift Container Platform that is based on KubeVirt which in terms allows you to run and manage virtual machine workloads alongside container workloads. Red Hat OpenShift Virtualization adds new objects into your OpenShift Container Platform cluster by using Kubernetes custom resources to enable virtualization tasks. The high level tasks include the following:

Creating and managing Linux and Windows virtual machines (VMs)

Create virtual machines (VMs) from custom operating system images by using one of the following methods:

- Importing the image as a container disk from a registry.
- Importing the image from a web page.
- Uploading the image from a local machine.
- Cloning a persistent volume claim (PVC) that contains the image.
- Managing virtual machine

Virtual Machine Wizard

The Virtual Machine Wizard in Red Hat OpenShift Virtualization is a convenient tool that simplifies the process of creating virtual machines. It is available through the web console and guides you through several steps to configure your virtual machine:

- General: Fill in essential details such as the virtual machine name, namespace, and template. Required fields are marked with an asterisk (*). You can also select a template to autofill some fields.
- Networking: Configure Network Interface Cards (NICs) for your virtual machine. By default, a nic0 NIC is attached, but you can add more if needed. NICs can also be created after the virtual machine is created.
- Storage: Add disks to your virtual machine. You can create additional disks and customize their settings. If you choose URL or Container as the source in the General step, a rootdisk disk is automatically created and attached as the Bootable Disk. You can modify the rootdisk but cannot remove it. A Bootable Disk is not required for virtual machines provisioned from a PXE source if no disks are attached.
- Review and Create: Review your configuration and create the virtual machine. The Results screen displays the JSON configuration file for the virtual machine.

NOTE

If you choose URL or Container as the source in the General step, a rootdisk disk is automatically created and attached to the virtual machine as the Bootable Disk. You can modify the rootdisk but cannot remove it. A Bootable Disk is not required for virtual machines provisioned from a PXE source if no disks are attached.

Prerequisite: Your virtual machine's storage medium must support Read-Write-Many (RWX) Persistent Volume Claims (PVCs).

Connecting to virtual machines through a variety of consoles and Command Line Interface tools

Red Hat OpenShift Virtualization provides different virtual machine consoles that you can use to accomplish different product tasks. You can access these consoles through the OpenShift Container Platform web console and by using command line interface (CLI) commands.

- Accessing virtual machine consoles in the OpenShift Container Platform web console
  - Connecting to the serial console
  - Connecting to the VNC console
  - Connecting to a Windows virtual machine with RDP
- Accessing virtual machine consoles by using CLI commands
  - Accessing a virtual machine via SSH by using virtctl

Importing and cloning existing virtual machines

You can create a VM by importing a custom image from a container registry or a web page, by uploading an image from your local machine, or by cloning a persistent volume claim (PVC).

- Importing the image from a web page.
- Importing the image as a container disk from a registry
- Uploading the image from a local machine.
- Cloning a persistent volume claim (PVC) that contains the image.

Managing network interface controllers and storage disks attached to virtual machines

Red Hat OpenShift Virtualization provides advanced networking functionality by using custom resources and plugins. Virtual machines (VMs) are integrated with Red Hat OpenShift Container Platform networking and its ecosystem. You can also configure a default storage class, storage profiles, Containerized Data Importer (CDI), data volumes, and automatic boot source updates.

- Red Hat OpenShift Virtualization networking overview
- Storage configuration overview

Live migrating virtual machines between nodes

Live migration is the process of moving a running virtual machine (VM) to another node in the cluster without interrupting the virtual workload. By default, live migration traffic is encrypted using Transport Layer Security (TLS).

An enhanced web console provides a graphical portal to manage these virtualized resources alongside the OpenShift Container Platform cluster containers and infrastructure.

Data Plane Development Kit and Single Root I/O Virtualization

Red Hat OpenShift Virtualization supports using Data Plane Development Kit (DPDK) with Single Root I/O Virtualization (SR-IOV). Let us break it down:

- DPDK and SR-IOV: DPDK provides a set of libraries and drivers for fast packet processing in user space. By leveraging DPDK, you can achieve lower latency and higher throughput for packet handling.
- SR-IOV allows you to share physical network interfaces among multiple virtual machines (VMs) while maintaining hardware-level isolation. It is particularly useful for high-performance workloads.

Graphics Processing Unit acceleration in Red Hat OpenShift Virtualization

Graphics Processing Units (GPU) are powerful resources that can significantly boost performance for various workloads. Red Hat OpenShift Virtualization allows you to create virtual GPUs (vGPUs) inside virtual machines (VMs). vGPUs can be shared among multiple consumers simultaneously, providing efficient resource utilization. You can choose between GPU passthrough and MIG to access containerized GPUs . NVIDIA GPU Operator and OpenShift Container Platform integrates with NVIDIA GPUs using the NVIDIA GPU Operator. This operator manages the full lifecycle of NVIDIA software components required for GPU-accelerated workloads

Migration Toolkit for Virtualization

The Migration Toolkit for Virtualization (MTV) enables you to migrate virtual machines from VMware vSphere, Red Hat Virtualization, or OpenStack to OpenShift Virtualization running on Red Hat OpenShift Virtualization platform.

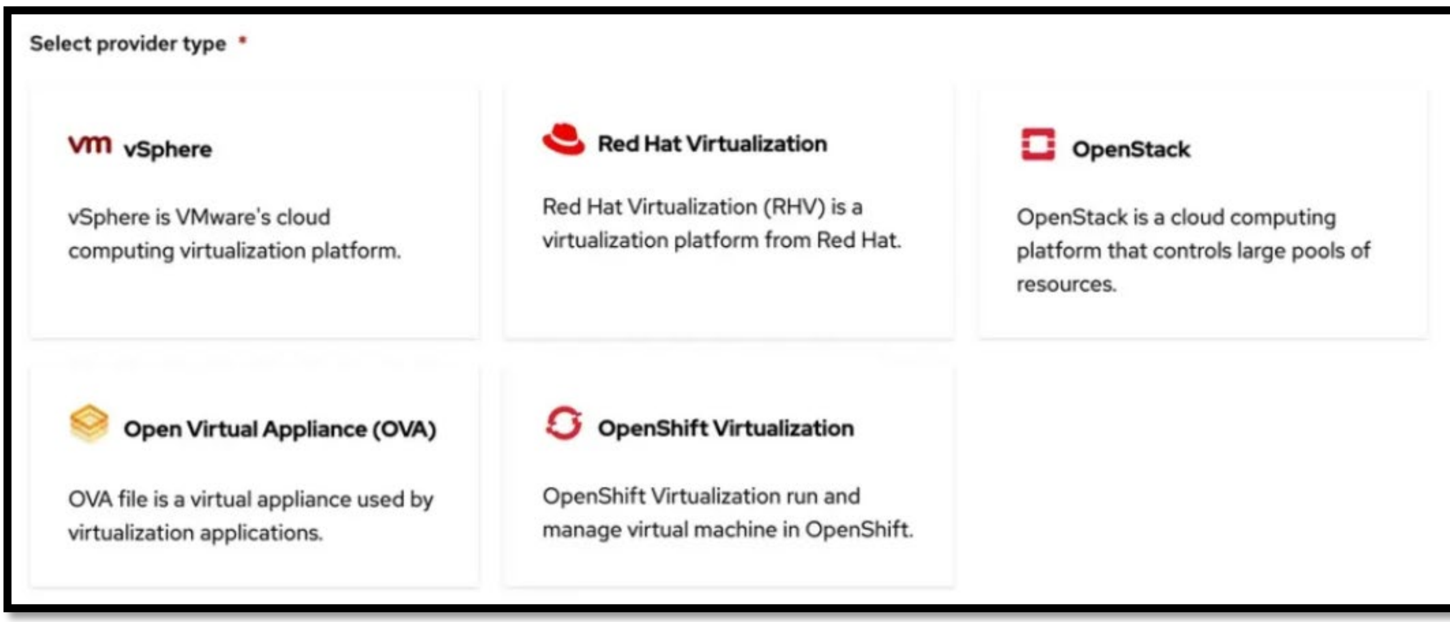

FIGURE 8. Migration Toolkit for Virtualization supported providers

MTV simplifies the migration process, allowing you to seamlessly move VM workloads to Red Hat OpenShift Virtualization

MTV supports Cold migration which is the default migration type. The source virtual machines are shutdown while the data is copied. Cold migration from

- VMware vSphere
- Red Hat Virtualization (RHV)
- OpenStack
- Remote Red Hat OpenShift Virtualization clusters

MTV supports warm migration from VMware vSphere and from RHV. In warm migration most of the data is copied during the precopy stage while the source virtual machines (VMs) are running. Then the VMs are shut down and the remaining data is copied during the cutover stage

Best Practice for Migration Toolkit for Virtualization

When planning a migration using the Migration Toolkit for Virtualization (MTV), following best practices can significantly improve the chances of a successful and smooth migration. Let us explore some key recommendations:

- Define Clear Objectives:
  - Clearly define the goals and objectives of your migration project. Understand why you are migrating, what you expect to achieve, and how success will be measured.

- Consider factors such as cost savings, performance improvements, scalability, or compliance requirements.

- Assess Your Environment:
  - Conduct a thorough assessment of your existing environment. Understand the source infrastructure, including VM configurations, storage, networking, and dependencies.
  - Identify any potential challenges, risks, or limitations that may impact the migration process.
- Plan Incrementally:
  - Break down the migration into smaller, manageable phases. Avoid attempting a big-bang migration.
  - Prioritize workloads based on criticality, complexity, and dependencies. Start with less critical workloads to gain experience and confidence.
- Test and Validate:
  - Set up a test environment to validate the migration process. Test various scenarios, including different VM types, storage configurations, and network setups.
  - Validate data integrity, application functionality, and performance after migration.
- Backup and Rollback Strategy:
  - Always have a backup strategy in place. Take snapshots or backups of VMs before migration.
  - Define a rollback plan in case issues arise during or after migration. Ensure you can revert to the previous state if needed.
- Communication and Training:
  - Communicate the migration plan to all stakeholders, including IT teams, application owners, and end-users.
  - Provide training on any changes or new processes resulting from the migration.
- Monitor and Optimize:
  - Monitor the migration process in real-time. Use monitoring tools to track performance, resource utilization, and any anomalies.
  - Optimize VM configurations, adjust resource allocations, and fine-tune settings post-migration.
- Document Everything:
  - Maintain detailed documentation throughout the migration process. Document configurations, decisions, and any issues encountered.
  - This documentation will be valuable for troubleshooting, future reference, and knowledge transfer.

Successful migrations require careful planning, testing, and collaboration. By following these best practices, you can minimize risks and achieve a successful outcome.

UNDERSTANDING RED HAT OPENSHIFT VIRTUALIZATION FEATURES

Red Hat OpenShift Virtualization is based on KubeVirt and KubeVirt in turn leverages the KVM hypervisor within Red Hat Enterprise Linux, Kernel, hence Red Hat OpenShift Virtualization supports almost all feature capabilities of KVM Hypervisor with a few exceptions.:

Memory Overcommit for Virtual Machines

If your virtual workload requires more memory than what is available, you can use memory overcommitment to allocate most or all the host's memory to your virtual machine instances (VMIs). Enabling memory overcommitment allows you to maximize resources that are typically reserved for the host. While it's possible to implement memory overcommitment in Red Hat OpenShift Virtualization, the point to keep in mind that memory overcommitment increases the potential for VM processes to be killed due to memory pressure (OOM killed). The likelihood of a VM being OOM killed depends on various factors, including node memory, available swap space, VM memory consumption, and other consideration. Hence the recommendation is not to enabling memory overcommitment for Virtual machines in Red Hat OpenShift Virtualization,

Disabling Guest Memory Overhead Accounting

By default, a small amount of memory is requested by each VM instance in addition to the amount you explicitly request. This additional memory is used for the infrastructure that wraps each VirtualMachineInstance process. While it is possible although it is not usually advisable, you can increase VM instance density on the node by disabling guest memory overhead accounting. However, this also increases the risk of VM processes being killed due to memory pressure (OOM killed).

Virtual Machine Snapshots

Red Hat OpenShift Virtualization snapshots, it is essential to understand the concept of crash-consistent snapshots. A crash-consistent snapshot captures the state of a virtual machine (VM) at a specific point in time, focusing on the data already written to the virtual disks. In other words, it represents the VM's state as if it had crashed abruptly, without any in-memory data or pending I/O operations being captured in the snapshot. Crash-consistent snapshots are useful for backup and disaster recovery scenarios. They allow you to roll back to a previous state quickly, but they do not guarantee data consistency within the VM. When taking a crash-consistent snapshot of a running VM, the system does not attempt to quiesce the VM's file system or ensure that all in-flight I/O operations are written to disk before capturing the snapshot. In summary, crash-consistent snapshots are faster and easier to create, but they may not capture all in-memory data or ensure data consistency.

Storage Migration between different Storage Class for Virtual Machine

StorageClass in Kubernetes focuses on dynamic provisioning and flexibility, while a datastore in virtualization environments represents a predefined storage location for VM files. Both concepts serve crucial roles in managing storage for different workloads. migrating data for a virtual machine across StorageClass in Red Hat OpenShift Virtualization is currently not supported.

Virtual Machine Encryption

The Red Hat OpenShift Virtualization supports encryption for virtual machine storage only if supported by the persistent volume providers. Red Hat OpenShift Virtualization has supported vTPM (Virtual Trusted Platform Module) as an option since OpenShift 4.13, with persistent storage capability added in OpenShift 4.14. With vTPM, VMs can securely manage encryption keys without exposing them directly to the host or other VMs.

CAPACITY AND SIZING

Red Hat OpenShift Container Platform defaults are designed to work well out of the box. Many pods and VMs scenarios and recommendations are documented in this section. Consider the following tested object maximums when running VMs on Red Hat OpenShift Virtualization. These values are based on the largest possible cluster size and reaching near the maximum values may reduce performance and increase response latency, carefully consider all of the multidimensional factors that limit the cluster scale. These maximums and minimums apply to Red Hat OpenShift Virtualization 4.x as a large-scale environment. The limits apply on a per-cluster basis. Keep in mind that some applications may work well in an overcommitted CPU environment. However, currently memory cannot be overcommitted for VMs. The following limits apply to Red Hat OpenShift Virtualization 4.x, per cluster.

TABLE 7 Scale limits for Red Hat OpenShift Virtualization 4.x, per cluster

Objective	4.x tested maximum	Theoretical limit
Virtual Machine Maximums	The following maximums apply to the OpenShift VM run on Red Hat OpenShift Virtualization	
Maximum virtual CPUs per virtual machine	216	710
Maximum memory per virtual machine	6 TB	16 TB
Minimum memory per virtual machine	512MB	
Maximum single disk size per virtual machine	6 TB	100 TB
Maximum number of hot-pluggable disks per virtual machine	255	
Host Maximums	The following maximums apply to the OpenShift hosts used for Red Hat OpenShift Virtualization	
Logical CPU cores or threads	Same as RHEL with KVM	
RAM	Same as RHEL with RHEL as KVM	
Simultaneous live migrations	Default to 2 outbound migrations per node, and 5 concurrent migrations per cluster	Depends on NIC bandwidth
Live migration bandwidth	There is no default bandwidth limit for live migrations	Depends on NIC bandwidth
Cluster Maximums	The following maximums apply to objects defined in Red Hat OpenShift Virtualization	
Maximum number of attached PVs per node		CSI storage provider dependent

Objective	4.x tested maximum	Theoretical limit
Maximum PV size		CSI storage provider dependent
Maximum number of hosts	500 (<100 recommended)	Same as OpenShift release
Maximum number of defined VMs	10,0004	Same as OpenShift release

For more information can be found in tuning guide which is a supplemental document for Red Hat OpenShift Virtualization. It focuses on fine-tuning cluster scalability and VM performance for several use cases and environments, for a variety of workloads and cluster sizes. This document provides guidance for scaling up and scaling out an environment, up to 100 OpenShift nodes based on this reference architecture, and for highly tuning each component of a VM definition.

For sizing a Red Hat OpenShift Container Platform 4.15 environment varies depending on the requirements of the organization and the type of deployment. This ensures the need for their environment is addressed based on Red Hat's published documentation around scalability and performance for each Red Hat OpenShift Container Platform release. For more information, see Red Hat OpenShift Container Platform scalability documentation.

BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

This section discusses the high-level cabling, networking, and storage layout of the solution hardware and software.

Network Overview

All the Red Hat OpenShift Container Platform control plane nodes and worker nodes in the cluster shall have the same network as that of the "Machine Config" server during boot to fetch ignition files. All the nodes in the cluster need to be assigned an IP address by the DHCP server.

The Red Hat OpenShift Container Platform 4.15 cluster also needs to have Internet access to perform the following tasks:

1. Access the Red Hat OpenShift Cluster Manager page to download the installation program and perform subscription management. If the cluster has internet access and you do not disable Telemetry, that service automatically entitles your cluster.
2. Access `Quay.io` to obtain the packages that are required to install your cluster.
3. Obtain the packages that are required to perform cluster updates.

Figure 8 lists the various networks used for this solution. All the cluster nodes and iPXE servers are connected to the same network.

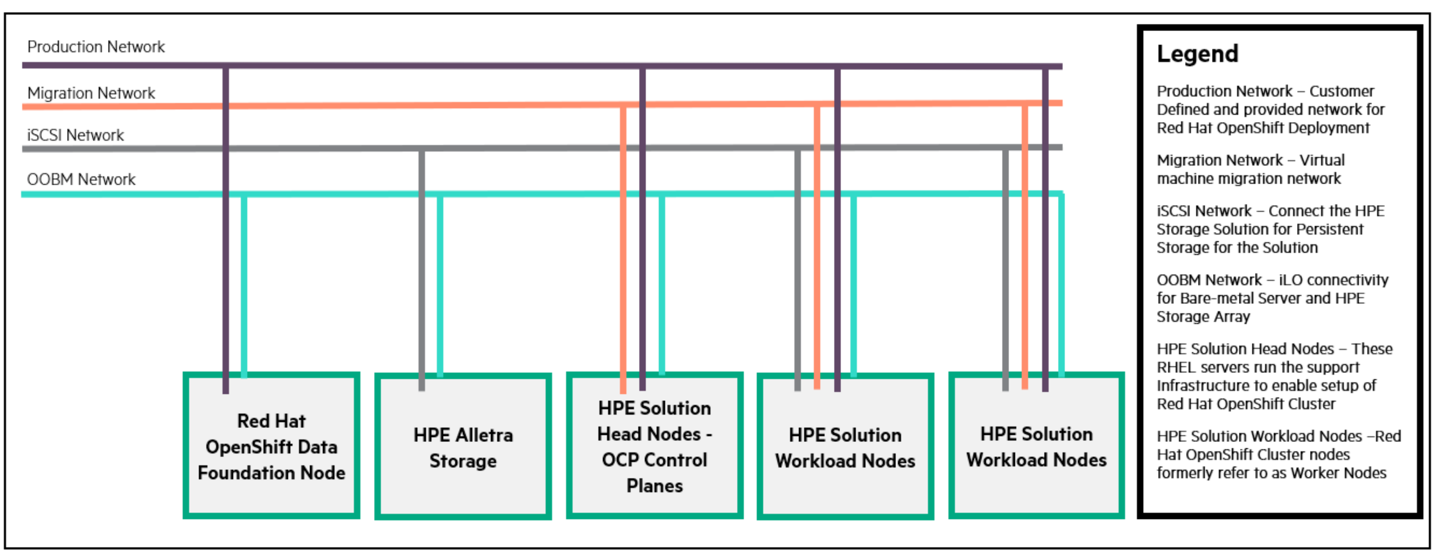

FIGURE 9. Networks for solution

NOTE
For OOBM High Availability, it is recommended to use 2x Aruba 6300 switches.

Storage

In the internal storage mode, the storage for the operating system and internal persistent volume is provided by the local storage disks (SSD) on the HPE ProLiant DL325 and HPE ProLiant DL385 Gen11 servers. Whereas in the external storage mode, the operating system storage is provided by local disks, and the container storage is provided by the HPE storage system such as HPE Alletra or by Red Hat OpenShift Container Platform that uses the local disks.

The Red Hat OpenShift Data Foundation (ODF) operator installation will be using the Local Storage operator. It provides persistent storage for services including OpenShift, monitoring, logging and registry, and other container-based applications that require persistent storage.

Table 7 lists all volumes used within the solution for the storage systems and highlights what storage provides the capacity and performance for each function.

TABLE 7. Details of the volume

Source	Volume/Disk Function	Hosts	Shared/Dedicated
Local storage on the servers	Red Hat OpenShift Container Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated
	Operating System	All Nodes	Dedicated
HPE Alletra Storage Array	iSCSI Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated
HPE ProLiant DL385 Gen11 (Red Hat OpenShift Data Foundation Internal storage)	Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated

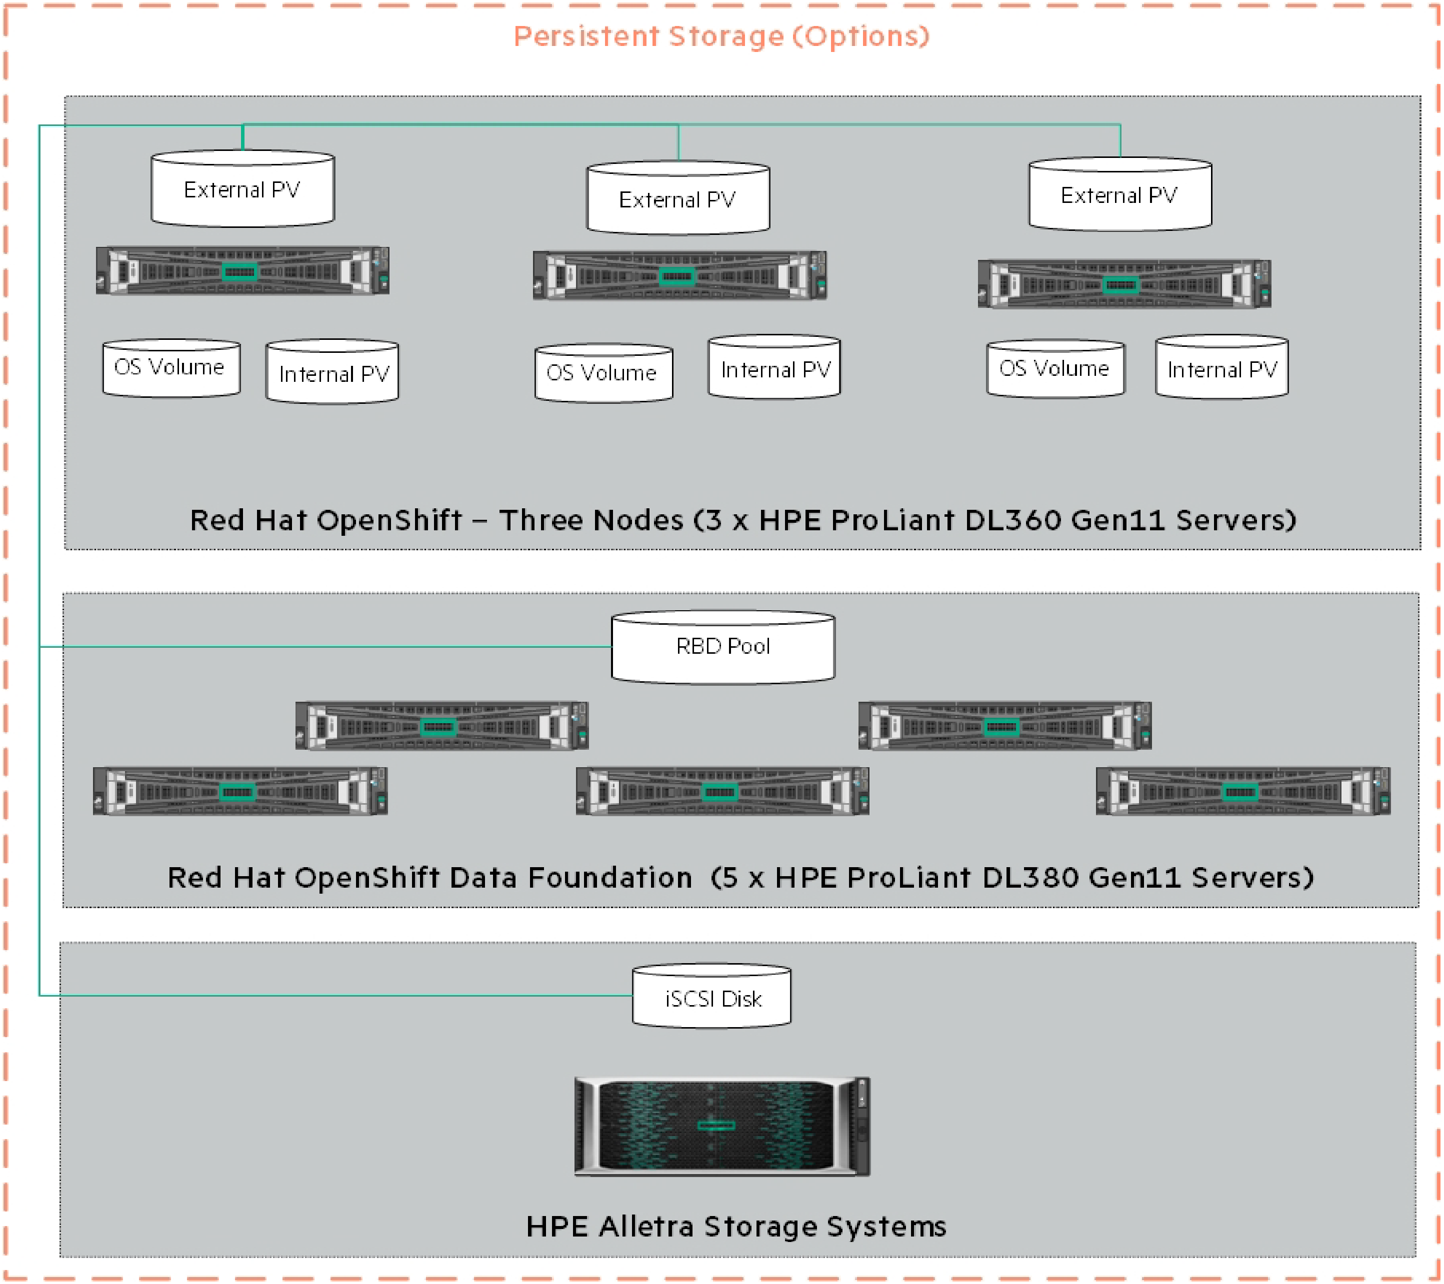

FIGURE 10. Logical storage layout for Red Hat OpenShift persistent volume options

---

NOTE

- HPE ProLiant DL385 Gen 11 servers can be added to this solution as per the customer's choice of configuration storage options.
- HPE Alletra 6K provides external storage to solution.
- In this solution, we have used Red Hat OpenShift Data Foundation internal and require Red Hat OpenShift Container Platform Plus (OPP) licenses.

DEPLOYMENT OVERVIEW

This section explains in detail the deployment of Red Hat OpenShift Container Platform 4.15 using internal and external storage mode. In the external storage mode, HPE Alletra 6070 is connected via the iSCSI network to the solution worker nodes.

Deploying the Red Hat OpenShift Container Platform 4.15 cluster using the User Provisioned Infrastructure

The Red Hat OpenShift Container Platform User Provisioned Infrastructure (UPI) deployment is a multi-step process. In this solution, most of the tasks are automated using the Hewlett Packard Enterprise developed automation scripts, whereas a few steps need manual intervention to complete the deployment.

The installer machine in the deployment environment uses the Red Hat OpenShift-installer program to create Red Hat Enterprise CoreOS ignition configuration files. These ignition files include the bootstrap ignition files, the solution control plane ignition files, and Workload ignition files. The ignition files are used to configure Red Hat Enterprise CoreOS on each of the solution control planes and worker nodes in the Red Hat OpenShift cluster. For detailed installation and configuration information, see the Deployment guide.

Figure 10 explains the Red Hat OpenShift Container Platform (RHOCP) 4.15 deployment process.

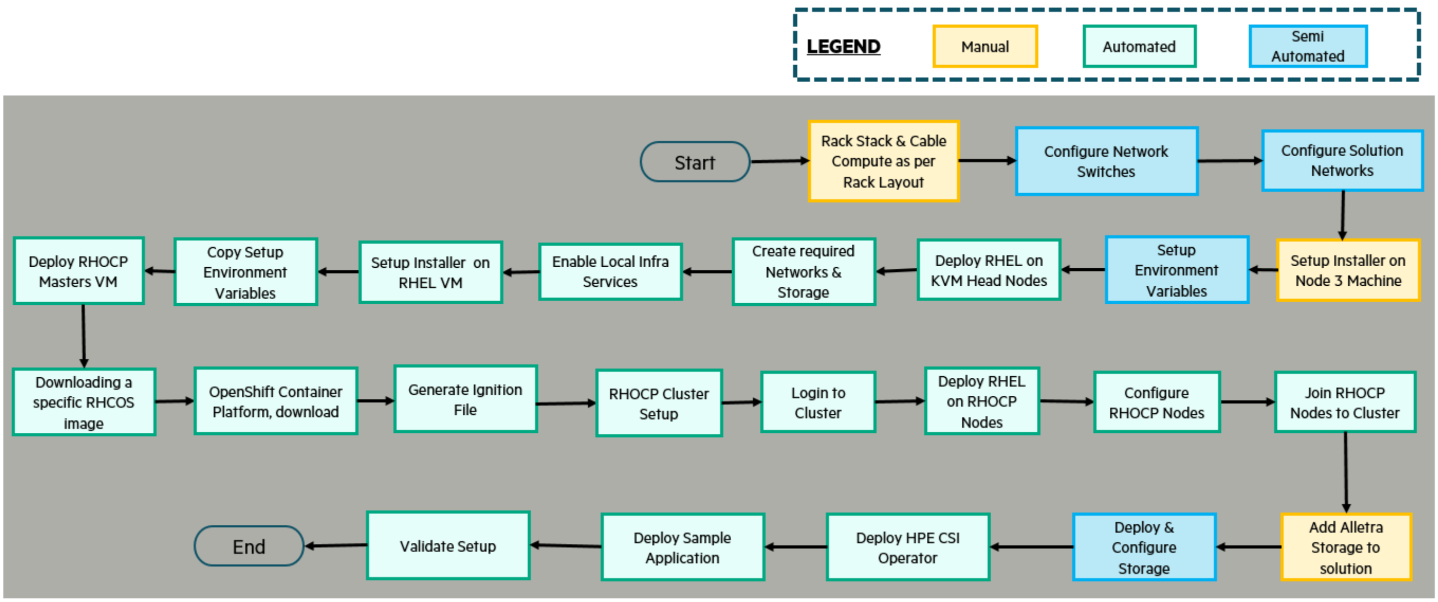

FIGURE 11. Deployment process for Red Hat OpenShift Container Platform 4.15 cluster using the UPI

NOTE
The load balancer described in this document is HAProxy.

Deploying Red Hat OpenShift Data Foundation

Red Hat OpenShift Data Foundation is deployed as an operator for internal storage mode with a minimal cluster of three (3) worker node servers. Spread the nodes across three different availability zones to ensure availability. Red Hat OpenShift Data Foundation can be set up as the default storage class in the Red Hat OpenShift Container Platform. The Red Hat OpenShift Data Foundation in our test environment was configured on the virtualized setup. The details of configuration and procedure on storage sizing are described in the Deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/4.15-AMD-LTI/.

Deploying HPE Alletra CSI

The HPE Container Storage Interface (CSI) Driver is a multi-vendor and multi-backend driver where each implementation has a Container Storage Provider (CSP). The HPE CSI Driver allows any vendor or project to develop its own CSP by using the CSP specification. This makes it very easy for third parties to integrate their storage solutions into Kubernetes as all the intricacies are taken care of by the HPE CSI Driver. The

CSI specification includes constructs to manage snapshots as native Kubernetes objects and create a new Persistent Volume Claim (PVC) by referencing those objects. The details of HPE Alletra CSI are described in the Deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/4.15-AMD-LTI/.

Deploying Red Hat OpenShift Virtualization

Red Hat® OpenShift® Virtualization, an included feature of Red Hat OpenShift, provides a modern platform for organizations to run and deploy their new and existing virtual machine (VM) workloads.

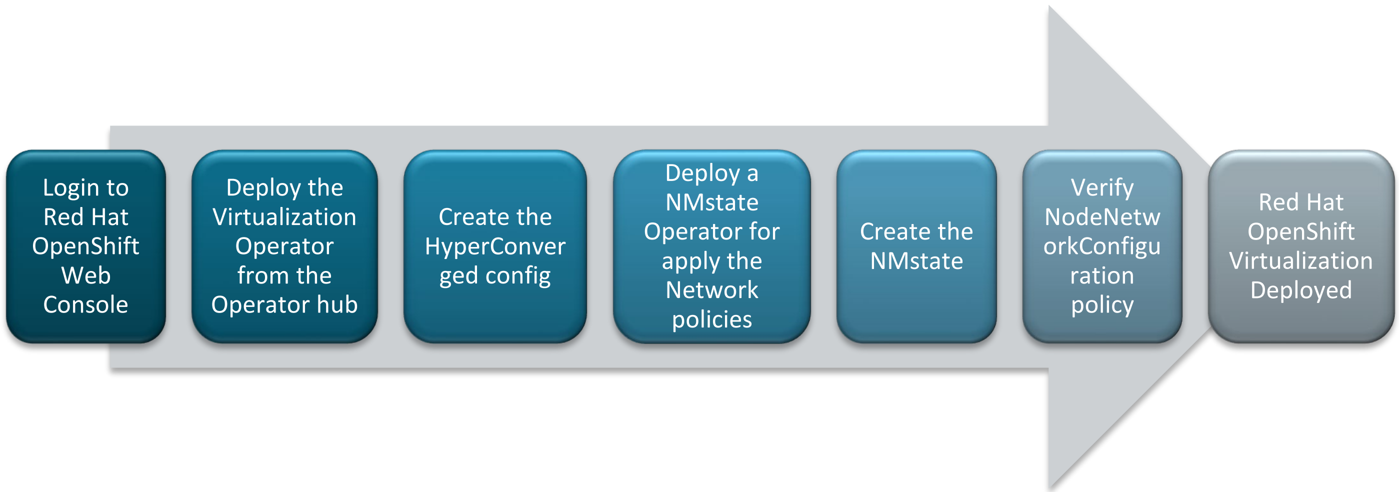

FIGURE 12. Red Hat OpenShift Virtualization deployment flow

Accelerating deployment

Automating the deployment fosters accuracy by decreasing the number of steps involved in setting up the solution. This solution leverages automation scripts developed by Hewlett Packard Enterprise to reduce the effort and time involved in deploying, configuring, and validating Red Hat OpenShift Container Platform 4.15. This in turn improves business productivity and promotes an "Idea Economy", where success is defined by the ability to turn ideas into value faster than the competition.

The graphs in this section quantify the time saved and the steps reduced in our lab setup. The graphs serve as a reference, and the time or the steps involved might differ depending on various environmental factors such as Infrastructure complexity and user proficiency with OpenShift. The key point in using automation scripts is to ensure improved business productivity.

Figure 12 depicts the time difference in forming a manual vs automated deployment of the Red Hat OpenShift Container Platform on bare metal servers using scripts mentioned in this document.

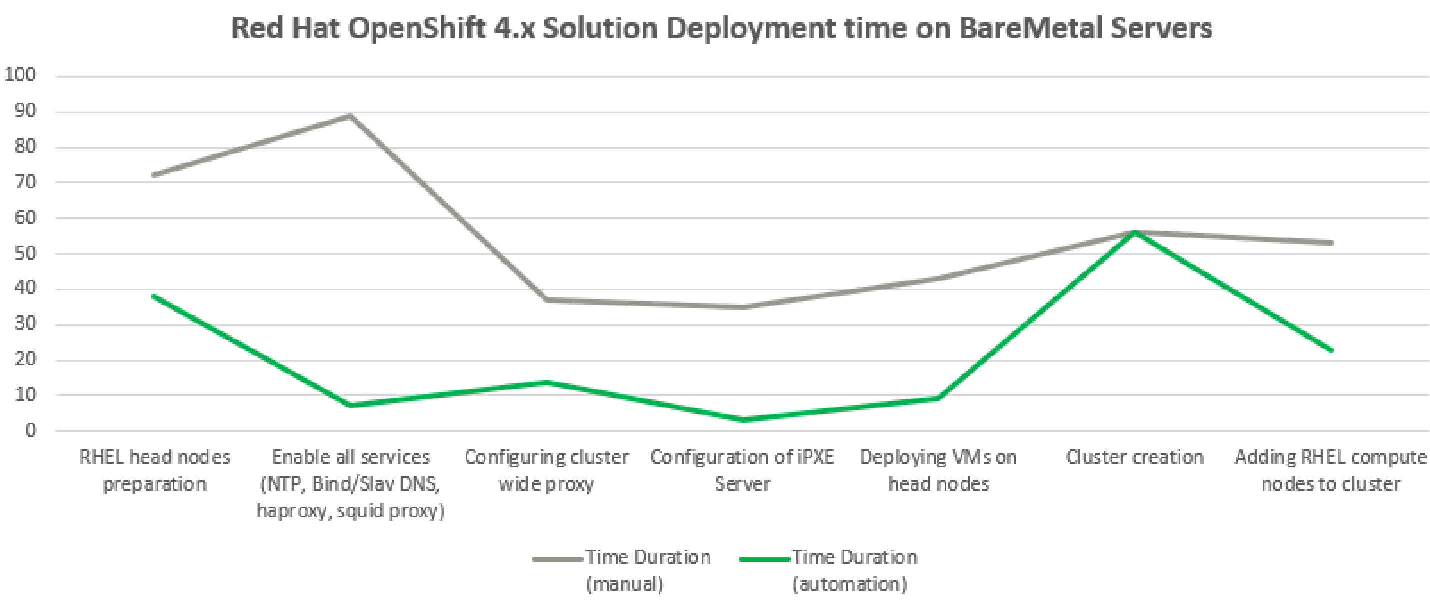

FIGURE 12. Red Hat OpenShift 4.x solution deployment manual and automation timelines on bare metal

Figure 13 depicts the steps involved in setting up a manual vs automated deployment of the Red Hat OpenShift Container Platform on bare metal using scripts mentioned in this document.

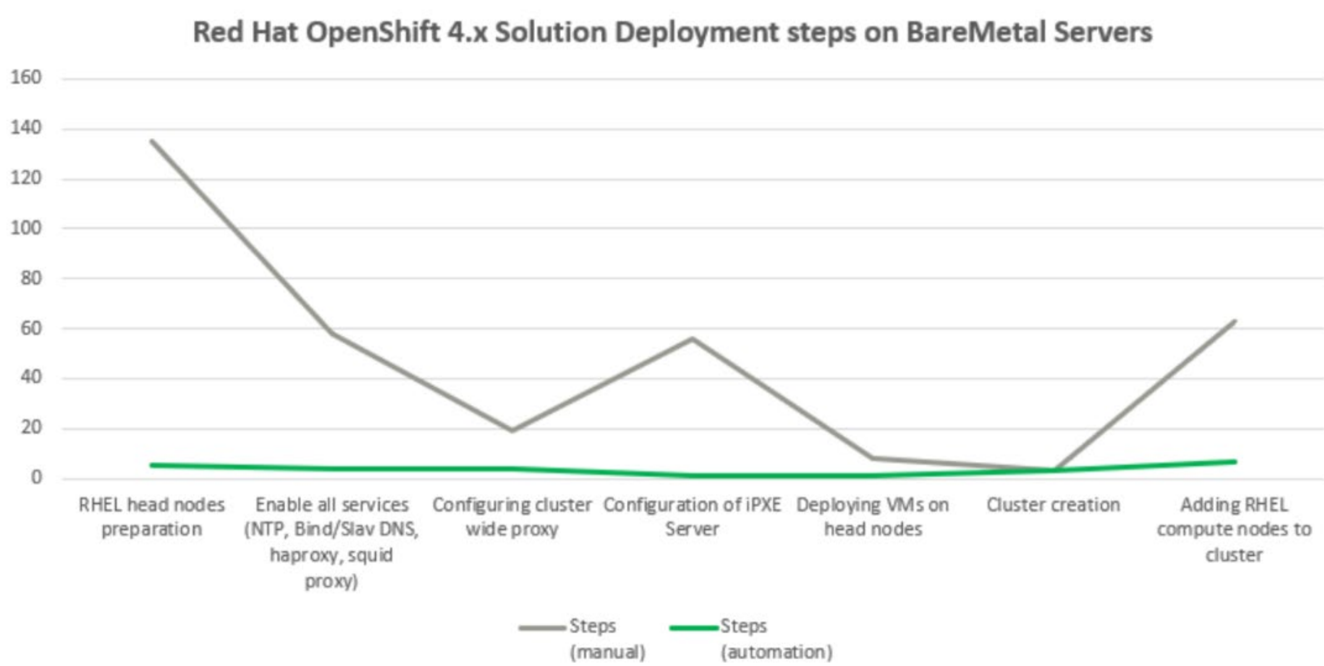

FIGURE 133. Red Hat OpenShift 4.x solution deployment manual and automation steps on bare metal

Securing Red Hat OpenShift with Red Hat Advanced Cluster Security for Kubernetes

Red Hat Advanced Cluster Security for Kubernetes is a Kubernetes-native security platform that equips you to build, deploy, and run cloud-native applications with more security. The solution helps protect containerized Kubernetes workloads in all major clouds and hybrid platforms, including Red Hat OpenShift, Amazon Elastic Kubernetes Service (EKS), Microsoft Azure Kubernetes Service (AKS), and Google Kubernetes Engine (GKE). This solution is included with Red Hat OpenShift Platform Plus, which provides a comprehensive set of tools to secure, protect, and manage your applications. Key Features and Benefits:

- DevSecOps Integration: Integrate with DevOps and security tools to mitigate threats and enforce security policies within your Kubernetes environment. Provide actionable guidelines to developers and streamline security analysis and remediation.
- Secure the Software Supply Chain: Continuously scan and ensure images by integrating with CI/CD pipelines and image registries. Shift security left to catch vulnerabilities and misconfigurations early.
- Kubernetes Infrastructure Protection: Ensure the underlying Kubernetes infrastructure remains hardened and protected, whether on-premises or in the cloud

Use Cases:

- Securing Ingress Controllers: You can dynamically secure the Ingress controller of a newly provisioned managed Red Hat
- OpenShift cluster using Red Hat Advanced Cluster Management (RHACM) policies in conjunction with Cert Manager running on the hub.
- Scanning OpenShift Image Streams: Use Red Hat Advanced Cluster Security to scan images held within OpenShift image streams (the OpenShift registry). This helps organizations identify policy violations and vulnerabilities before pushing images to external container registries.

Securing your Kubernetes environment is crucial, and solutions like Red Hat Advanced Cluster Security for Kubernetes play a vital role in achieving that.

BUSINESS CONTINUITY WITH DATA PROTECTION FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.15

Backup and restore is a management phase operational task for making periodic copies of configuration and application data to a separate or secondary device and then using those copies to recover the data and applications. This process is done to mitigate the risk if the original data and applications are lost or damaged due to a power outage, cyberattack, human error, disaster, or some other unplanned event. Traditional backup solutions have existed for a while in the ecosystem of the Enterprise Datacenter. These solutions need to evolve to address the needs of the new container infrastructure where Velero adds value. Velero is an open-source tool that is used to safely back up and restore, performs disaster recovery, and migrate Kubernetes cluster resources and persistent volumes.

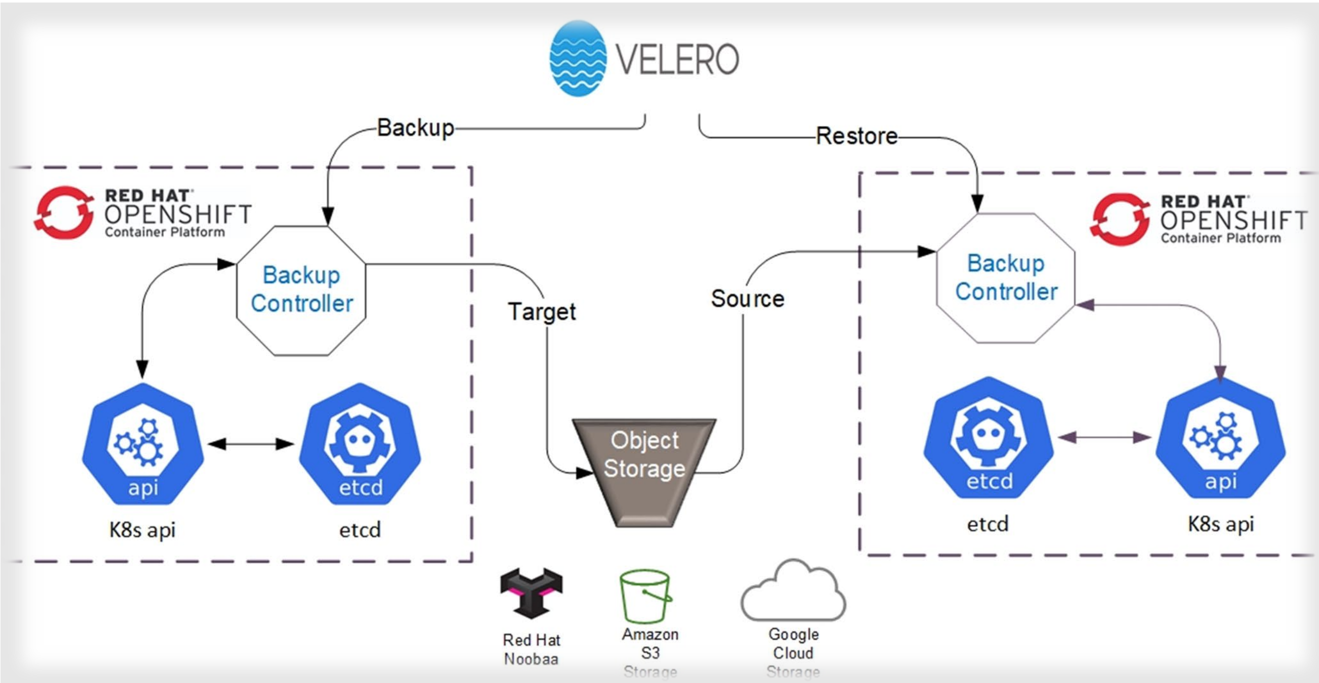

FIGURE 14. Velero backup and restore with Red Hat OpenShift Container Platform 4.15

Velero provides the following features to the Kubernetes-based container ecosystem:

- **Data Protection** - Offers key data protection features such as scheduled backups, retention schedules, and pre- or post-backup hooks for custom actions.
- **Disaster Recovery** - Reduces time to recovery in case of infrastructure loss, data corruption, and/or service outages.
- **Data Migration** - Enables cluster portability by easily migrating Kubernetes resources from one cluster to another.

In Red Hat OpenShift Container Platform 4.15, Velero uses a controller model where it monitors custom resources and takes actions.

Velero development consists of a server that runs in the Red Hat OpenShift Cluster and a command line client that runs locally on the management machine.

Prerequisites

- Red Hat OpenShift Container Platform cluster must be available with the administrator credentials.
- When using public cloud-based object storage, the appropriate Velero plug-in is required along with the access information and credentials.

Velero for Red Hat OpenShift setup overview

Figure 15 shows the overview of Velero for the Red Hat OpenShift setup.

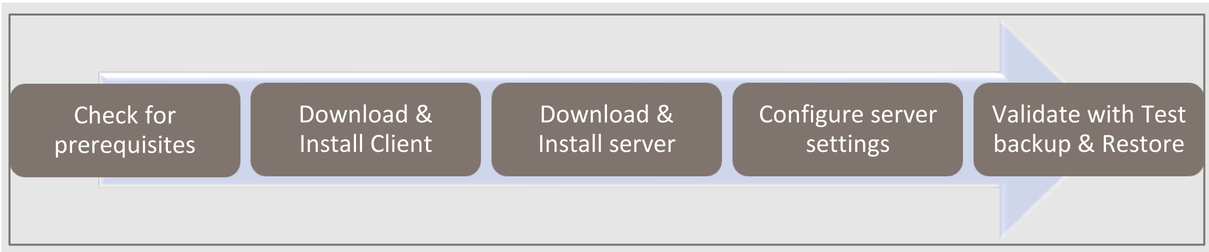

FIGURE 15. Velero for Red Hat OpenShift setup overview

Velero makes it simple to back up the Red Hat OpenShift configuration information and application data to a Cloud Object-based storage platform and restore it on demand. For more information, see the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/4.15-AMD-LTI/

BUSINESS CONTINUITY WITH DISASTER RECOVERY STRATEGIES FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.15

Stateful applications need a more sophisticated Disaster Recovery (DR) strategy than stateless applications, as a state must be maintained along with traffic redirection. Disaster recovery strategies become less generic and more application specific as application complexity increases. In this section, we shall see the various options available to provide disaster recovery for an application running on Red Hat OpenShift Container Platform 4.15 deployment. Recovery Time Objective (RTO) and Recovery Point Objective (RPO) are two key metrics that must be considered to develop an appropriate disaster recovery plan that can maintain business continuity after an unexpected event. RTO is the organization's tolerance for "App Downtime" and RPO is the organization's tolerance for "Data Loss".

Figure 16 shows the comparison of the Red Hat OpenShift disaster recovery strategies using RTO and RPO objectives.

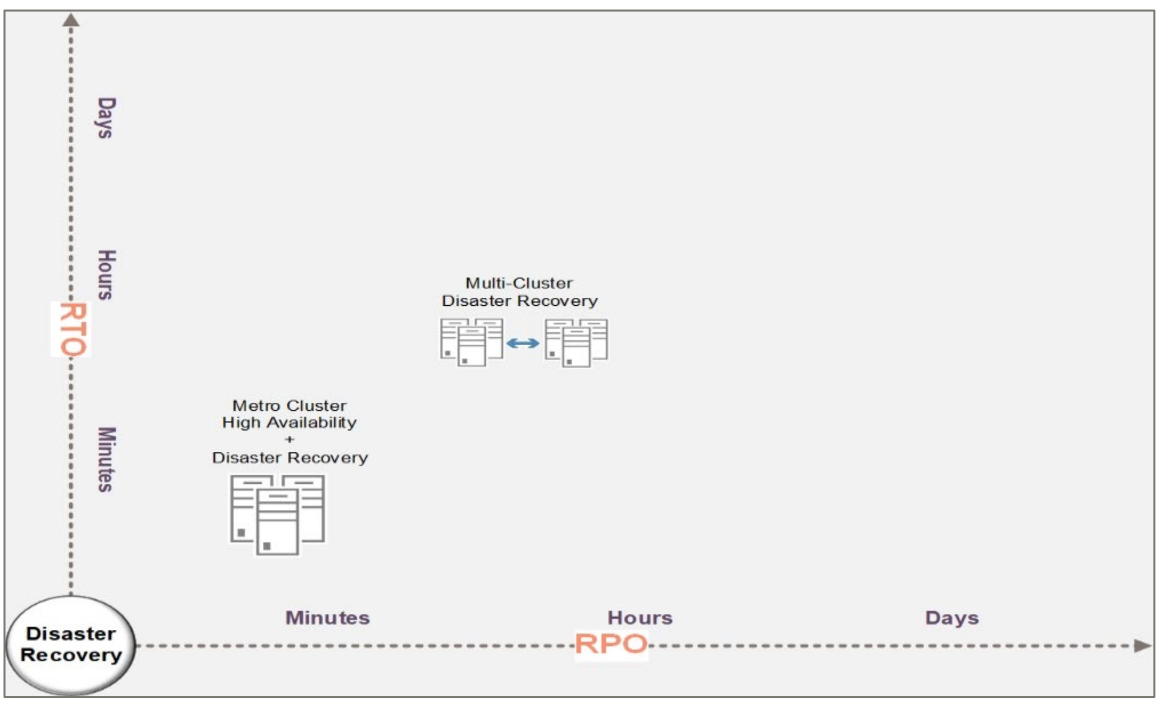

FIGURE 16. Red Hat OpenShift Disaster recovery strategies comparison using RTO and RPO objectives.

Metro Cluster High Availability and Disaster Recovery strategy

The Metro Cluster High Availability, also known as stretched or distributed clustering, is a high-availability configuration that allows one compute/storage cluster, such as a single Red Hat OpenShift cluster, to be stretched across two or more physically separate sites or data centers in an active/active DR strategy. It is recommended to use a minimum of three physically separate sites or data centers to meet generic application Service Level Agreements (SLA).

The following are the requirements for High Availability like automatic recovery along with no data loss data mirroring:

1. Synchronous High Availability-Disaster Recovery for localized data center failures.
   a. DR sites or Availability Zones (AZs) connected by MAN or campus networks.
   b. AZs are mapped to a fault domain (Heating, Ventilation and Air Conditioning, Power grids, etc.).
   c. An odd number of AZ or fault domains are required for the cluster quorum.
   d. Network latency between zones does not typically exceed 5 ms Round-Time Trip.
2. The solution ensures pods and nodes get scheduled across zones during deployment.
3. Red Hat OpenShift Data Foundation maintains consistent mirror copies across AZs resulting in less or no data loss.
4. Stretched solution cluster provides automatic and non-disruptive recovery for apps across AZs.
5. An application with a consensus protocol that allows it to determine which instances of the cluster are active and healthy.

Figure 17 shows an overview of the Red Hat OpenShift (OCP) Metro Cluster design.

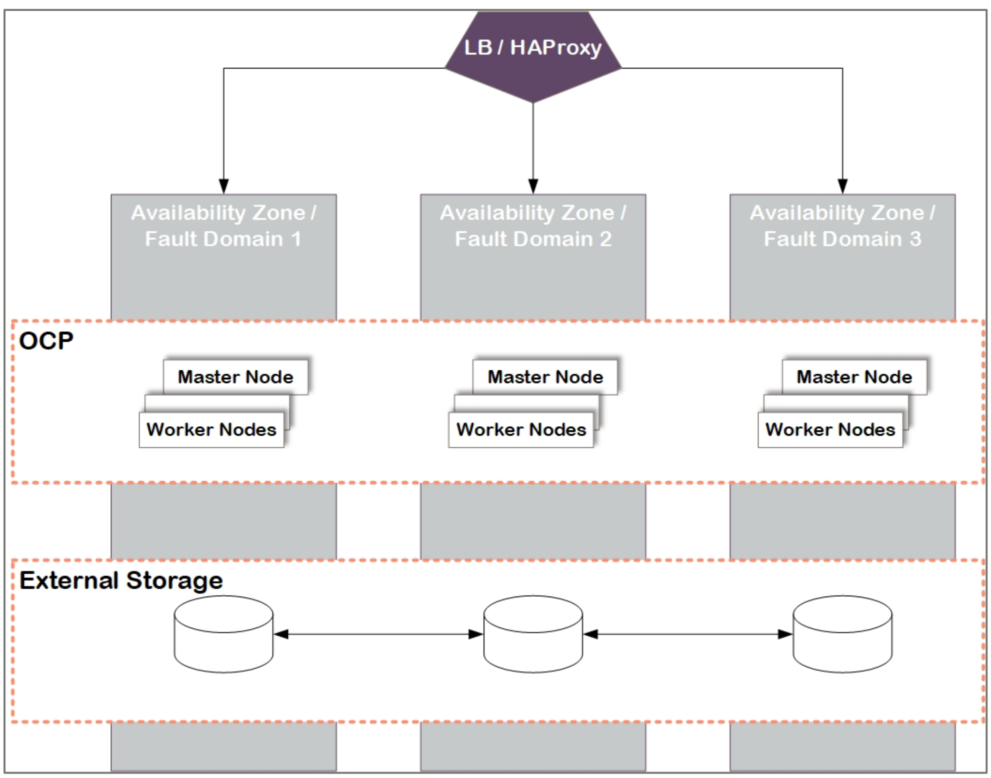

FIGURE 17. Red Hat OpenShift Metro Cluster design overview

When one of the AZs is down, no action needs to occur as both Red Hat OpenShift and the stateful workload will autonomously react to the situation. In particular, the stateful workload will sense the loss of one of the instances and will continue using the remaining instances. The same is true when the affected AZ is recovered. When the stateful instance in the recovered AZ comes back online, before the instance is allowed to join the cluster, it will need to resync its state. Again, this is handled autonomously and is part of the clustering features of some stateful workloads.

Multi-Cluster Disaster strategy
In this strategy, the multiple data centers (at least three) are geographically distributed. Each data center has its own independent Red Hat OpenShift clusters. A global load balancer balances traffic between the data centers. The stateful workload is deployed across the Red Hat OpenShift clusters. This approach is more suitable than the previous one for geographical, on-premises, and hybrid deployments. The compute and storage clusters are independent clusters, and the storage cluster is accessed using an external storage access framework from within the Red Hat OpenShift compute cluster. In this configuration, the members of the stateful workload cluster need to be able to communicate with each other across multiple clusters. Also, this entire strategy is dependent on the ability to replicate the state from the active site to another site. Each workload is different, so these various approaches should be chosen to meet SLA requirements according to cluster compute and storage configuration such as:

- Volume-level Replication
- Application-level Replication
- Proxy-level Replication

When one AZ is down, the global load balancer must be able to sense the unavailability of one of the data centers and redirect all traffic to the remaining active data centers. No action needs to occur on the stateful workload as it will self-reorganize to manage the loss of a cluster member.

Figure 18 shows the Red Hat OpenShift (OCP) Multi-cluster disaster recovery approach.

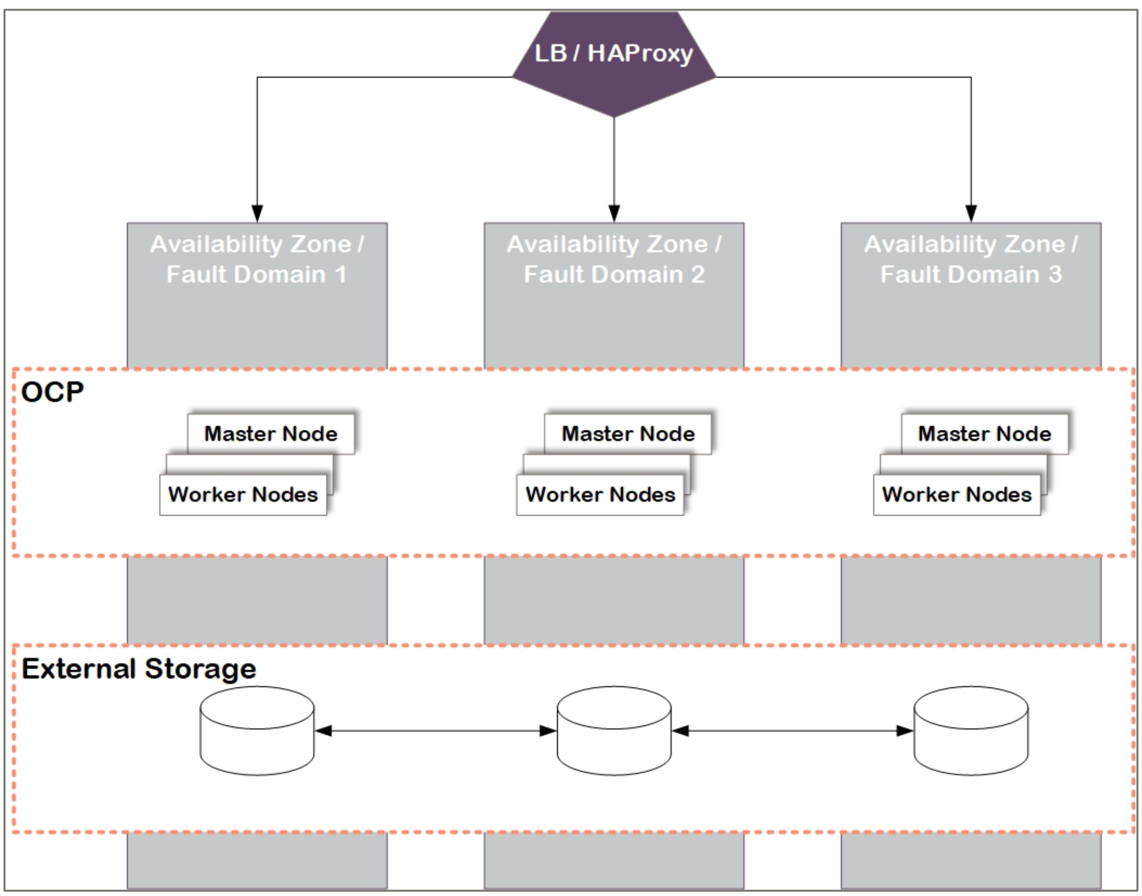

FIGURE 18. Red Hat OpenShift Multi-cluster disaster recovery approach

For more information, see Disaster Recovery Strategies for Applications Running on OpenShift.

RED HAT OPENSHIFT DATA FOUNDATION– INTERNAL MODE

The Red Hat OpenShift Data Foundation also formerly referred to as Red Hat OpenShift Container Storage (OCS) operator installation will be using Local Storage operator which will use file system storage of 10 GB for monitoring purposes and block storage of 500GB/2TB for Object Storage Daemon (OSD) volumes. These OSD are useful for configuring any application on top of Red Hat OpenShift container Platform cluster using Red Hat OpenShift Data Foundation configuration.

FIGURE 19. Red Hat OpenShift Data Foundation Flow Diagram

For more information, view the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/4.15-AMD-LTI/.

SUMMARY

In this solution installation and configuration of the Red Hat OpenShift Virtualization in Red Hat OpenShift Container Platform Version 4.15 on the HPE ProLiant DL325 configured as master and worker nodes with storage like HPE Alletra Storage Arrays and HPE ProLiant DL385 Gen11 servers as storage nodes for Red Hat OpenShift Data Foundation. This solution provides customers with greater efficiency, higher utilization, and bare-metal performance by “collapsing the stack” and eliminating the need for a separate virtualization platform. IT teams can manage multiple instances of containers and virtual machines using multiple Kubernetes clusters with multitenant container isolation and data access, for any workload from edge to core and cloud. The benefits of virtual machine & containers within a single platform goes beyond cloud-native microservices architected stateless applications can be extended by providing the ability to containerize monolithic stateful analytic applications with persistent data. Red Hat OpenShift Virtualization offers a holistic value proposition that empowers organizations to adapt, innovate, and thrive in a dynamic digital landscape.

Benefits include:

- Deploying the management, etc., and worker nodes on bare metal eliminates the overhead associated with hypervisors and thus optimizes performance.
- Deploying Red Hat OpenShift Container Platform 4.15 on HPE ProLiant DL325 and DL385 Gen11 servers using automation scripts saves significant efforts, resulting in quicker deployment.

APPENDIX A: BILL OF MATERIALS

The following BOMs contain electronic license-to-use (E-LTU) parts. Electronic software license delivery is now available in most countries. Hewlett Packard Enterprise recommends purchasing electronic products over physical products (when available) for faster delivery and for the convenience of not tracking and managing confidential paper licenses. For more information, please contact your reseller or a Hewlett Packard Enterprise representative.

Red Hat OpenShift Platform Plus subscription

Red Hat OpenShift Data Foundation and Red Hat Advanced Cluster Security for Kubernetes requires Red Hat OpenShift Platform Plus license (includes Red Hat OpenShift Data Foundation internal) and is socket based subscription.

NOTE

Part numbers are at the time of publication/testing and are subject to change. The bill for materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult your Hewlett Packard Enterprise Reseller or Hewlett Packard Enterprise Sales Representative. For more information, see hpe.com/us/en/services/consulting.html.

TABLE A1. Bill of materials

Component	Qty	Description
P9K08A	1	HPE 42U 600mmx1075mm G2 Kitted Advanced Shock Rack with Side Panels and Baying
P9K08A 001	1	HPE Factory Express Base Racking Service
P53921-B21	5	HPE ProLiant DL385 Gen11 8SFF Configure-to-order Server
P53921-B21 ABA	5	HPE DL385 Gen11 8SFF CTO Svr
P58540-B21	10	AMD EPYC 9224 2.5GHz 24-core 200W Processor for HPE
P50311-B21	120	HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P55082-B21	10	HPE ProLiant DL385 Gen11 8SFF Tri-Mode U.3 x1 BC Backplane Kit
P50230-B21	40	HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD
P63871-B21	10	HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD
P55097-B21	5	HPE ProLiant DL385 Gen11 x16 2U Secondary Riser Kit
P42044-B21	5	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE
P47789-B21	5	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
P42041-B21	5	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P44712-B21	10	HPE 1800W-2200W Flex Slot Titanium Hot Plug Low Halogen Power Supply Kit
E5Y43A	5	HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU
P57847-B21	5	HPE ProLiant DL385 Gen11 8SFF OROC x1 SAS/SATA Cable Kit
P57886-B21	5	HPE ProLiant DL385 Gen11 2U Standard/Performance FIO Air Baffle Kit
P58465-B21	30	HPE ProLiant DL3X5 Gen11 2U Performance Fan Kit
P50400-B21	5	HPE Gen11 2U Bezel Kit
P52351-B21	5	HPE DL3XX Gen11 Easy Install Rail 2 Kit
P57845-B21	5	HPE ProLiant DL385 Gen11 SFF Backplane Power Cable Kit
P58458-B21	10	HPE ProLiant DL3X5 Gen11 Standard 2U Heat Sink Kit
P54199-B21	6	HPE ProLiant DL325 Gen11 8SFF Configure-to-order Server
P54199-B21 ABA	6	HPE DL325 Gen11 8SFF CTO Svr
P53703-B21	6	AMD EPYC 9554P 3.1GHz 64-core 360W Processor for HPE
P50312-B21	72	HPE 64GB (1x64GB) Dual Rank x4 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit

Component	Qty	Description
P54999-B21	6	HPE ProLiant DL325 Gen11 8SFF x1 Tri-Mode U.3 Backplane Kit
P63871-B21	24	HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD
P42044-B21	6	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE
P47789-B21	6	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
P42041-B21	6	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P59668-B21	42	HPE ProLiant DL325 Gen11 Liquid Cooling Fan Kit
P03178-B21	12	HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
BD505A	6	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
S1A05A	6	HPE Compute Cloud Management Server FIO Enablement
P59619-B21	6	HPE ProLiant DL325 Gen11 8SFF x1 OCP2 Tri-Mode Cable Kit
P52351-B21	6	HPE DL3XX Gen11 Easy Install Rail 2 Kit
P58463-B21	6	HPE ProLiant DL325 Gen11 Closed-loop Liquid Cooling FIO Heat Sink Kit
R9F63A	1	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle
R9F63A B2B	1	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU
R9G06A	1	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable
R9G06A B01	1	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable
R9F61A	1	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit
R9F61A B2B	1	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit PDU
R9F57A	1	HPE Aruba Networking 1U Universal 4-post Rack Mount Kit
R9F59A	2	Aruba 4-post Rack Kit for HPE
R9F67A	2	HPE Aruba Networking 8325-32C Power-to-Port Airflow 6 Fans 2 Power Supply Units Bundle
R9F67A B2B	2	HPE Aruba Networking 8325-32C Power-to-Port Airflow 6 Fans 2 Power Supply Units Bundle PDU
R9F78A	12	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable
R9F78A B01	12	HPE Aruba Networking 100G QSFP28 to QSFP28 5m Direct Attach Copper Cable
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack include with Complete System Stabilizer Kit
R7A11AAE	6	HPE GreenLake for Compute Ops Management Enhanced 3-year Upfront ProLiant SaaS
R9G32AAE	1	HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 3y Subscription E-STU
R9G27AAE	2	HPE Aruba Networking Fabric Composer Device Management Service Tier 4 Switch 3y Subscription E-STU
HA113A1	1	HPE Installation Service
HA113A1 5BY	1	HPE Rack and Rack Options Install SVC
HA114A1	1	HPE Installation and Startup Service
HA114A1 5A6	5	HPE Proliant DL/ML Startup SVC
HU4A6A5	1	HPE 5Y Tech Care Essential Service
HU4A6A5 SVN	5	HPE One View w/Ilo Support
HU4A6A5 ZND	2	HPE Aruba Networking 8325-32 SW Support
HU4A6A5 ZSG	1	HPE Aruba Networking 6300M 48 SW Support
HU4A6A5 R2M	6	HPE iLO Advanced Non Blade Support
HU4A6A500DH	5	HPE DL385 GEN11 Support
HU4A6A500DE	6	HPE DL325 GEN11 Support

Component	Qty	Description
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5MR	1	HPE Tier 1 Storage Array Startup SVC
H33XSA1	20	HPE Learn Credits for Compute IT SVC
R4U31A	1	HPE Alletra 6070 Dual Controller Configure-to-order Base Array
R0R12A	2	HPE Alletra 6000 2x10/25GbE 2-port FIO Adapter Kit
R7S84A	1	HPE Alletra 6000 92TB (24x3.84TB) NVMe Flash Carrier FIO Flash Bundle
R9D23A	4	HPE C13 - C14 250V 10Amp 2m WW PDU FIO Power Cord
R9X15A	1	HPE Alletra Tier 1 Storage Array Standard Tracking
R7G13A	1	HPE Alletra 6000 4x 1600W FIO AC Power Supply Kit
Q8G27B	1	HPE Tier 1 Storage OS Default FIO Software
S0R29AAE	92	HPE Alletra Storage 6000 per TB 5-year Software and Support SaaS
HU4A6A5	1	HPE 5Y Tech Care Essential Service
HU4A6A5 ZUN	2	HPE Alletra 6000 2x10/25GbE 2p Kit Supp
HU4A6A5 ZUH	1	HPE Alletra 6070 Base Array Supp
HU4A6A5 ZV6	1	HPE Alletra 6000 AF 92TB 3.84 Flash Supp

NOTE
For high availability, 2x HPE Aruba 6300 switches are required.

RESOURCES AND ADDITIONAL LINKS

HPE Reference Architectures, hpe.com/info/ra

HPE servers, hpe.com/servers

HPE Storage, hpe.com/storage

HPE Networking, hpe.com/networking

HPE Technology Consulting Services, hpe.com/us/en/services/consulting.html

HPE ProLiant DL325 Gen11, hpe.com/servers

HPE ProLiant DL385 Gen11, hpe.com/servers

Red Hat OpenShift Container Platform, https://access.redhat.com/documentation/en-us/openshift_container_platform/4.15/

Red Hat OpenShift Container Storage, https://access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.15

HPE ProLiant Workload Solutions, https://www.hpe.com/us/en/servers/proliant-workload-solutions-ecosystem.html

To help us improve our documents, please provide feedback at hpe.com/contact/feedback.

a50010994enw, June 2024

Hewlett Packard Enterprise