HPE REFERENCE ARCHITECTURE FOR RED HAT OPENSHIFT ON HPE PROLIANT DL380 GEN10 AND HPE PROLIANT DL360 GEN10 SERVERS

Rapid deployment on HPE ProLiant DL servers using Red Hat OpenShift Container Platform 4.9

CONTENTS

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 environment has 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 are 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
- Quickly deploy disconnected and secured end-to-end container platform

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. To accelerate container application delivery, Hewlett Packard Enterprise and Red Hat® are collaborating to optimize Red Hat OpenShift® Container Platform 4.9 (RHOCP) on the HPE infrastructure which includes HPE ProLiant DL compute.

This Reference Architecture provides architectural guidance for deploying a Platform as a Service/Container as a Service (PaaS/CaaS) solution based on Red Hat OpenShift Container Platform 4.9 and HPE ProLiant DL servers. The compute and storage requirements can easily be scaled by adding more HPE ProLiant DL servers with no workload downtime.

The CNCF Operator Framework in the Reference Architecture provides a cloud-native method of packaging, deploying, and managing Kubernetes-native applications that include:

- Set up HPE ProLiant DL380 Gen10 and HPE ProLiant DL360 Gen10 servers
- To install and configure the Red Hat OpenShift Container Platform
- To install and configure the Red Hat OpenShift Data Foundation (ODF)
- Validate the Red Hat OpenShift Container Platform installation
- Set up and integrate HPE Ezmeral Data Fabric (EDF) for Persistent storage
- Validate OpenShift Virtualization to run and manage virtual machine workloads alongside container workloads
- Significant reduction in the deployment time and efforts through the automated deployment process

The solution demonstrates the cost-effective yet reliable solution by leveraging the benefits of HPE ProLiant DL servers, HPE storage solutions, and Red Hat OpenShift Container Platform 4.9.

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 on HPE ProLiant DL servers This document assumes that the target audience is familiar with HPE ProLiant DL servers, Red Hat OpenShift Container Platform 4, solutions, and core networking.

Document purpose: The purpose of this document describes the benefits and technical details of deploying Red Hat OpenShift Container Platform 4.9 on HPE ProLiant DL servers, the implementation details, and processes.

INTRODUCTION

This Reference Architecture provides guidance for installing Red Hat OpenShift Container Platform 4.9 on HPE ProLiant DL380 Gen10 and HPE ProLiant DL360 Gen10 servers. The solution consists of six (6) HPE ProLiant DL servers; three (3) HPE ProLiant DL360 Gen10 servers used for the OCP master nodes, three (3) HPE ProLiant DL380 Gen10 servers are used for the worker nodes (out of which one node will be used as the OCP bootstrap node). Later on, the bootstrap node will be reconfigured as an additional compute node (for Red Had CoreOS worker nodes only) and storage node.

Red Hat OpenShift Data Foundation (ODF) can be installed on the three (3) worker nodes. Each ODF node is equipped with a single SSD drive to be used as the ODF storage device. Optionally, the customer can configure the ODF in external storage mode. For this purpose, six (6) more HPE ProLiant DL380 Gen10 servers will be required to form a Red Hat Ceph storage cluster. HPE storage system can be used for the External storage mode.

SOLUTION OVERVIEW

This section provides an overview of Red Hat OpenShift Container Platform 4.9 User Provisioned Infrastructure (UPI) on bare metal HPE ProLiant DL380 Gen10 and HPE ProLiant DL360 Gen10 servers.

Figure 1 shows the high-level architecture of the solution. The storage can be internal or external depending upon the storage requirement. For OpenShift Data Foundation external mode, a six-node Red Hat Ceph cluster is required. For using HPE Ezmeral Data Fabric five-node cluster is needed

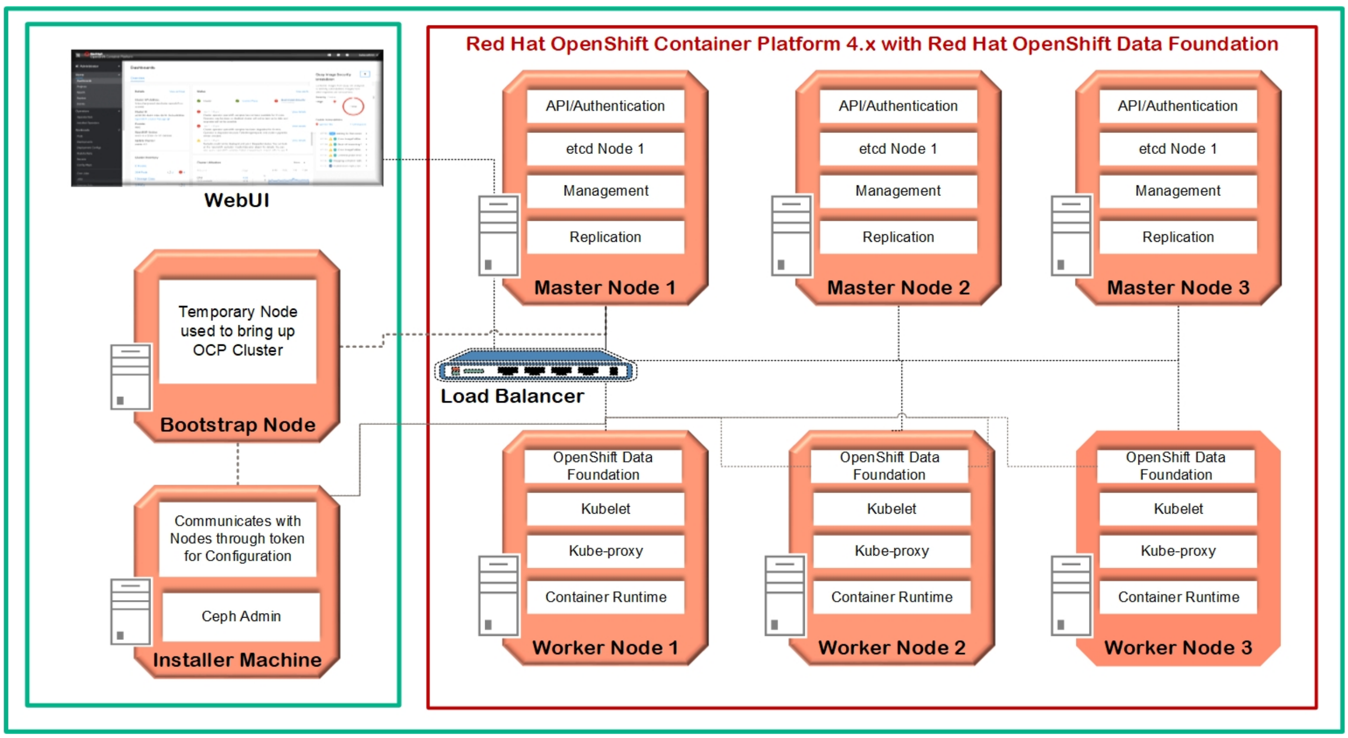

FIGURE 1. Red Hat OpenShift Container Platform 4.x

NOTE

The number of physical nodes represented in the solution layout is only a basic building block. This is subject to change based on customer requirements.

SOLUTION LAYOUT

This solution uses the OpenShift User Provisioned Infrastructure method of installation to install Red Hat Linux CoreOS on the HPE ProLiant DL servers and configure the OpenShift cluster

Physical bare-metal configuration

In a bare-metal configuration, the master nodes are deployed in a highly available configuration running on three (3) HPE ProLiant DL360 Gen10 servers. Red Hat OpenShift worker nodes are deployed on the bare metal with three (3) HPE ProLiant DL380 Gen10 servers. The temporary bootstrap node is deployed on one of the worker nodes and later on configured as a worker node.

The solution described in this Reference Architecture uses the internal storage on HPE ProLiant DL360 Gen10 and HPE ProLiant DL380 Gen10 servers. Optionally, six (6) more HPE ProLiant DL380 Gen10 servers can be added as a cluster for the Red Hat Ceph cluster for Red Hat OpenShift Data Foundation (ODF) external mode. The deployment environment (Installer machine, iPXE, DNS, DHCP, etc.) and a load balancer in this solution were deployed on the virtual machines. The OpenShift-installer tool is run to generate ignition files that contain information about the hosts that will be provisioned. The Red Hat Linux CoreOS (RHCOS) for the nodes is then booted with the help of iPXE and the ignition files are passed with the OS image during installation. Optionally HPE Storage systems such as HPE Nimble or HPE Ezmeral Data Fabric can be leveraged to provide support for persistent container volume.

Figure 2 provides an overview of the Red Hat OpenShift Container Platform 4 solution layout in a bare metal setup configuration.

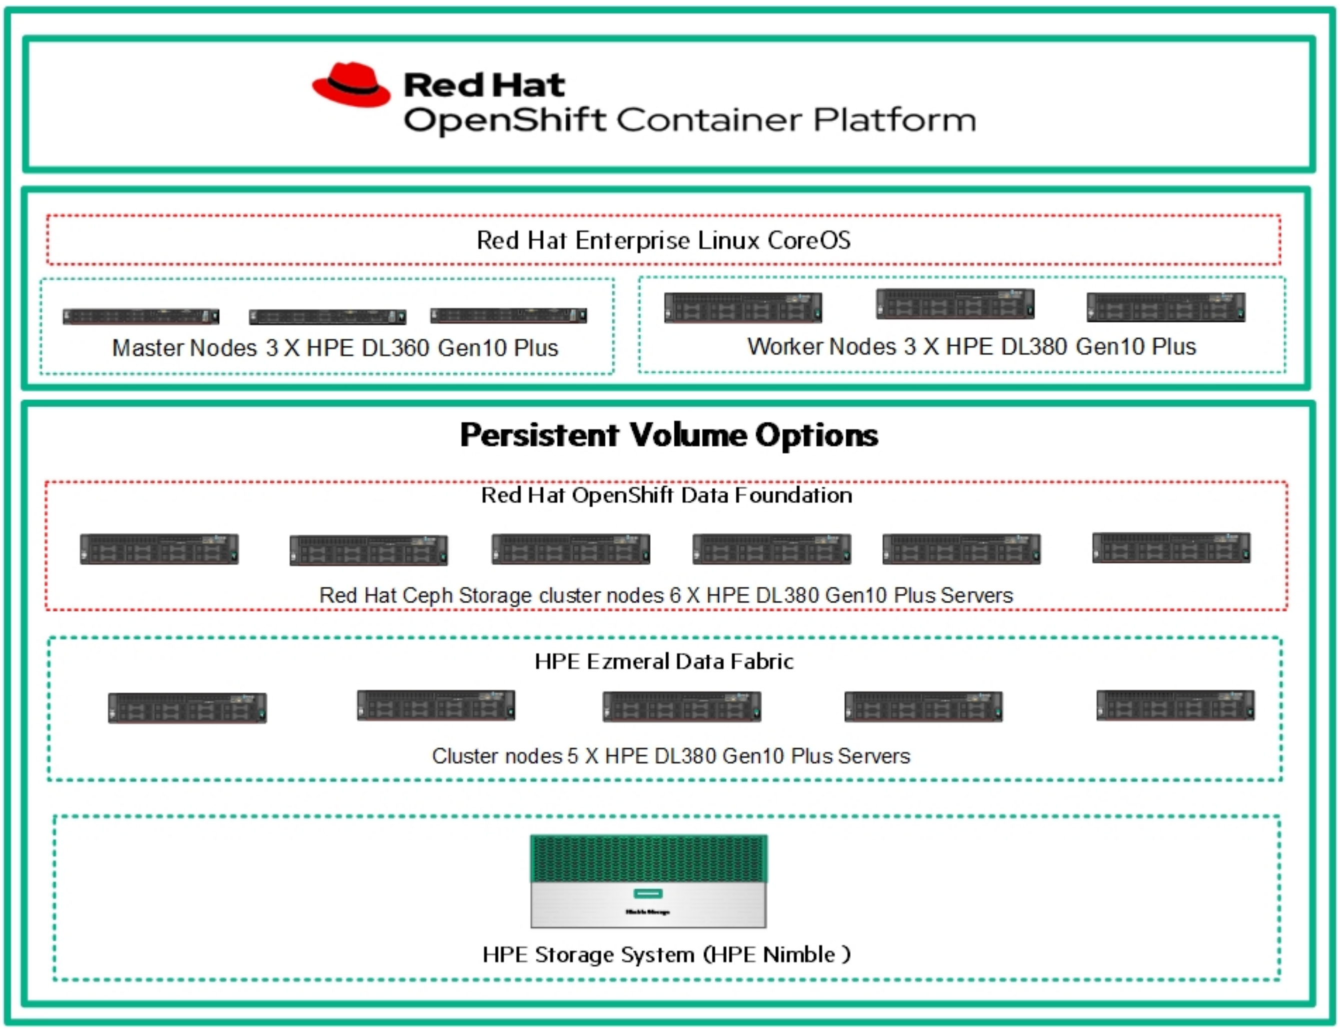

FIGURE 2. Solution Layout – Bare metal setup

Virtualized configuration

In a virtualized configuration, the OpenShift master nodes and worker nodes are deployed as VMware vSphere® virtual machines running on three (3) or more HPE ProLiant DL380 Gen10 servers run RHCOS version 4. The OpenShift install tool is run to generate ignition files that contain information about the hosts that will be provisioned. The CoreOS operating system for the nodes is then PXE booted and the ignition files are passed with the OS image during installation.

Figure 3 provides an overview of the Red Hat OpenShift Container Platform 4 solution layout in a virtualized setup configuration.

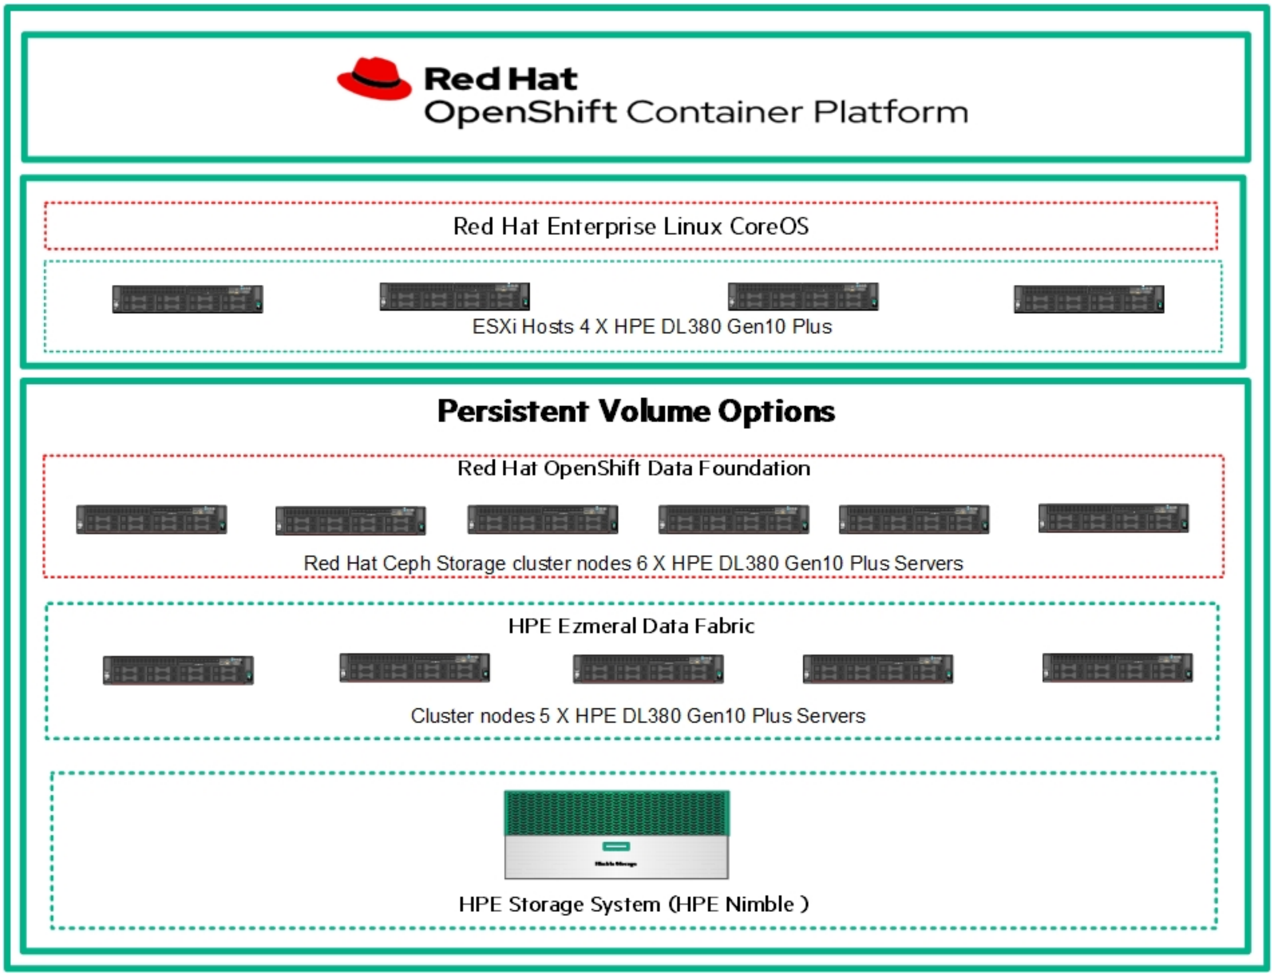

FIGURE 3. Solution Layout – Virtualized setup

The choices available for Persistent Volume storage in the solution are as follows:

- Install Red Hat OpenShift Data Foundation (Internal) as pods within the OpenShift Container Platform cluster
- Install Red Hat OpenShift Data Foundation (External) on OpenShift Container Platform cluster and connect to an external Red Hat Ceph Storage cluster
- Use the HPE CSI Operator for Kubernetes within OpenShift Container Platform cluster and connect to HPE Nimble or HPE Ezmeral Data Fabric

The solution assumes that the following infrastructure services and components are installed, configured, and function properly:

- LDAP or Active Directory (optional)

- DHCP
- DNS
- NTP
- TFTP
- iPXE
- External Load Balancer

The rack diagram of the hardware components that form the solution is shown here in Figure 4, inside a rack. This figure illustrates the rack view of the HPE ProLiant DL servers deployed in this solution. Additional HPE ProLiant DL360 Gen10 or HPE ProLiant DL380 Gen10 servers or both can be added to the solution as needed. For OpenShift Data Foundation external storage mode, which is optional, three (3) more HPE ProLiant DL380 Gen10 servers would be required to set up the Red Hat Ceph storage cluster. Optionally HPE Nimble storage can be used for External Persistent volumes using the HPE CSI Operator for Kubernetes or use HPE Ezmeral Data Fabric Cluster running, five (5) more HPE ProLiant DL380 Gen10 servers.

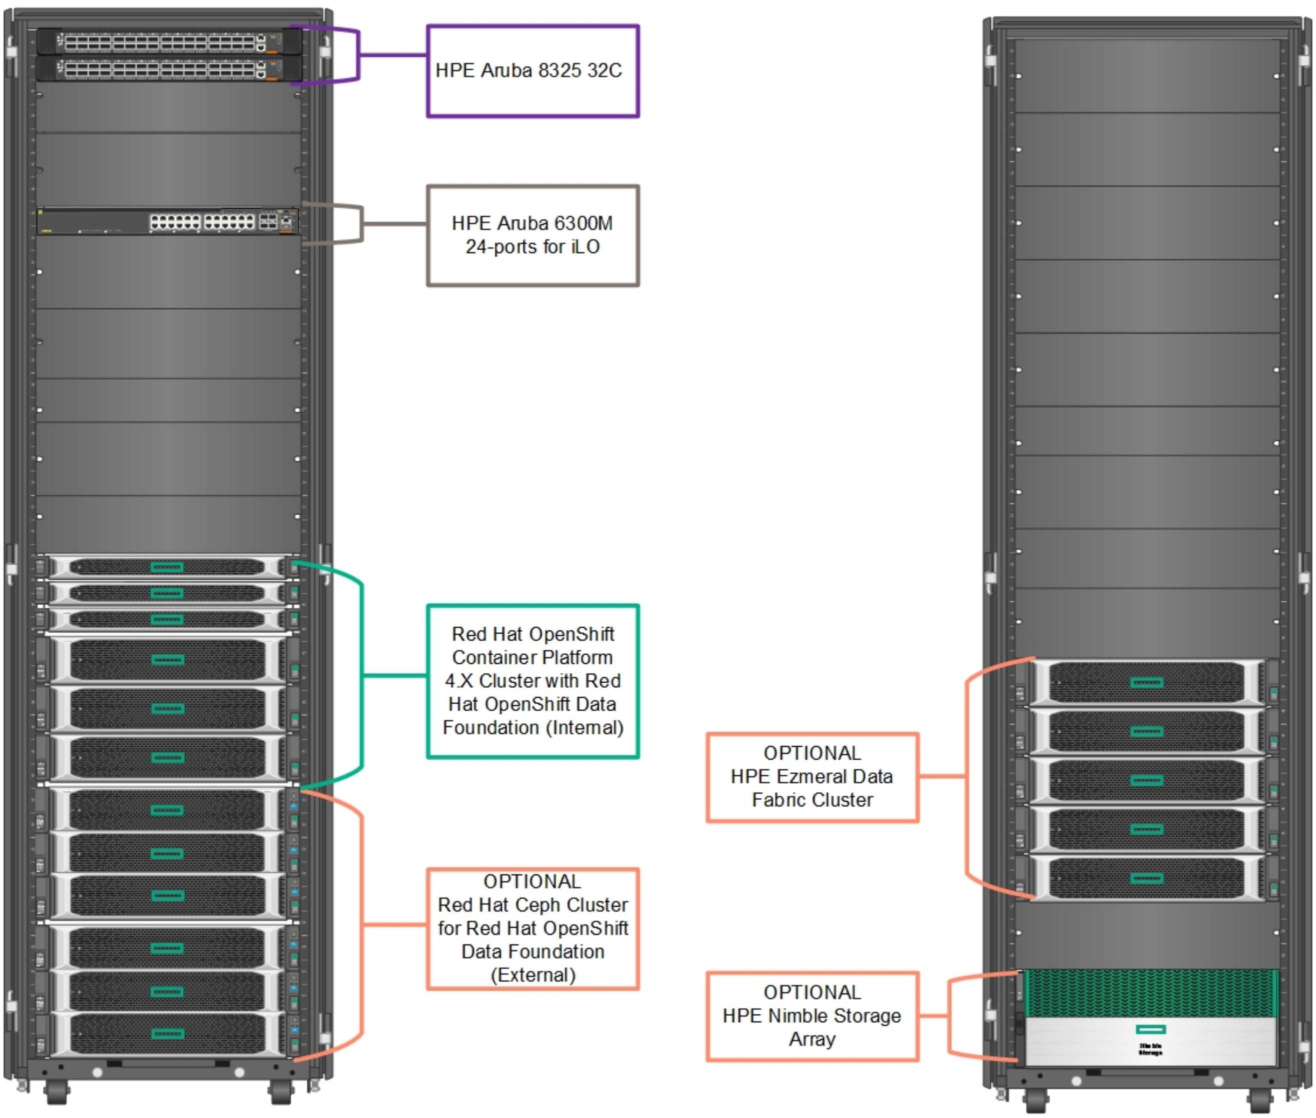

FIGURE 4. Front view of the HPE Rack and Solution Components featuring HPE ProLiant DL360 Gen10 or HPE ProLiant DL380 Gen10 servers, HPE Nimble Storage, and Network Switches

NOTE
The figure depicts the hardware layout in the test environment. However, this is subject to change based on the customer requirement. For providing Persistent storage for the OpenShift Container Platform the customer can choose to either to deploy ODF (internal or external) or optionally persistent volumes storage can be provided either using HPE Nimble or HPE Ezmeral Data Fabric using HPE CSI driver for Kubernetes.

SOLUTION COMPONENTS

This section provides the detail of the hardware and software components used in the solution being described.

Hardware

Table 1 shows the minimum resource requirements for the deployment of OCP 4.9. without ODF

TABLE 1. Minimum hardware requirement for OCP 4.9 cluster

Node	Operating System	vCPU	Virtual RAM	Storage
Bootstrap node	RHCOS	4	16 GB	120 GB
Control plane or Master nodes	RHCOS	4	16 GB	120 GB
Worker nodes	RHCOS	2	8 GB	120 GB

The configuration deployed for this solution is described in greater detail in this section.

Table 2 lists the various hardware components used in the solution.

TABLE 2. Hardware components utilized in this solution

Component	Qty	Description
HPE ProLiant DL360 Gen10 server	3	OpenShift master
HPE ProLiant DL380 Gen10 server	3	OpenShift worker nodes and bootstrap or compute node
HPE ProLiant DL380 Gen10 server	6	Red OpenShift Data Foundation nodes – External storage mode (optional)
HPE ProLiant DL380 Gen10 server	5	HPE Ezmeral Data Fabric (optional)
HPE Nimble storage (AF40)	1	External iSCSI storage (optional – Only for External storage mode)
HPE Aruba 8325 switch	2	A network switch for datacenter network
HPE Aruba 6300M switch	1	A network switch for iLO Management network
HPE ProLiant DL380 Gen10 server	1	Host the virtual machines for Installer Machine, HAProxy Load balancer (expected to be available in the customer environment)

TABLE 3. Hardware requirement for ODF Storage (external approach)

Node	Operating System	vCPU	RAM	Quantity	Hard Disks	Platform
MON, OSD	RHEL7.9	4	16 GB	3	/dev/sda – os volume /dev/sdb - 10 GB mon /dev/sdc -500 GB osd	HPE ProLiant DL380 Gen10
ADMIN	RHEL7.9	4	16 GB	1	/dev/sda – os volume	VM

Server roles and hardware configuration

Table 4 lists the various server roles and their configuration used in this solution.

TABLE 4. Server roles and configuration

Node	Operating System	vCPU	RAM	Storage
Bootstrap node	RHCOS	4	16 GB	Hard Disk 1 : 300 GB
Control plane or Master nodes	RHCOS	8	64 GB	Hard Disk 1 : 300 GB
Worker nodes	RHCOS	16	64 GB	Hard Disk 1 : 300 GB Hard Disk 2 : 300 GB Hard Disk 3 : 2.4 TB
HAProxy (Load Balancer)	RHEL 7	4	16 GB	300 GB

NOTE
The HAProxy load balancer was deployed on a virtual machine.

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

Figure 5 shows the HPE ProLiant DL380 Gen10 server and HPE ProLiant DL360 Gen10 server.

FIGURE 5. HPE ProLiant DL380 Gen10 server and HPE ProLiant DL360 Gen10 server

HPE ProLiant DL380 Gen10 server

Adaptable for diverse workloads and environments, the secure 2P 2U HPE ProLiant DL380 Gen10 delivers world-class performance with the right balance of expandability and scalability. Designed for supreme versatility and resiliency while being backed by a comprehensive warranty make it is ideal for multiple environments from Containers to Cloud to Big Data. Standardize on the industry's most trusted compute platform.

This section describes the HPE ProLiant DL380 Gen10 servers used in the creation of this solution. Table 5 describes the individual components. Each server was equipped with 192 GB RAM and dual Intel® Xeon® 2.8GHz 16 Core CPUs. Individual server sizing should be based on customer needs and may not align with the configuration outlined in this document.

The HPE ProLiant DL380 Gen10 servers used in this solution provide a robust platform to run containerized applications. The three (3) HPE ProLiant DL380 Gen10 servers used in this solution are deployed as OpenShift compute nodes and OpenShift Data Foundation nodes. One of the HPE ProLiant DL380 servers is initially used as the OpenShift bootstrap node. The bootstrap node provides resources used by the master nodes to create the control plane for the OpenShift cluster. The bootstrap node is a temporary role that is only used during the initial OpenShift cluster installation. After the OpenShift cluster bootstrap process is complete, the bootstrap node can be removed from the cluster and repurposed as an OpenShift compute node. In this Reference Architecture, the bootstrap node is used as one of the three OpenShift Data Foundation nodes.

Table 5 lists the hardware configuration in each of the HPE ProLiant DL380 Gen10 servers used in this solution.

TABLE 5. Hardware configuration in each of the HPE ProLiant DL380 Gen10 Server

Component	Description
Processor	2 x Intel Xeon-Gold 6242 (2.8 GHz/16-core/150 W)
Memory	12 x HPE 16GB 2Rx8 PC4-2933V-R Smart Kit
Network	HPE IB FDR/EN 40/50Gb 547FLR 2QSFP Adapter
Array Controller	HPE Smart Array E208i-a SR Gen10 Controller
Disks	HPE 960GB SATA MU SFF SC DS SSD

HPE iLO

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

HPE ProLiant DL360 Gen10 server

The HPE ProLiant DL360 Gen10 server supports the 3rd Generation Intel® Xeon® Scalable Processors with up to 40 cores, plus 3200 MT/s HPE DDR4 SmartMemory up to 4.0 TB per socket. Introducing PCIe Gen4 and Intel Software Guard Extensions (SGX) support on the dual-socket segment, the HPE ProLiant DL360 Gen10 server complements Gen10 reach by delivering premium compute, memory, I/O, and security capabilities for customers focused on performance at any cost.

This section describes the HPE ProLiant DL360 Gen10 servers used in the creation of this solution. Each server was equipped with 32GB RAM and dual Xeon 2.3 GHz 16 Core CPUs. Individual server sizing should be based on customer needs and may not align with the configuration outlined in this document.

The HPE ProLiant DL360 Gen10 servers in this Reference Architecture provide the OpenShift control plane and the OpenShift master nodes. The OpenShift master nodes are responsible for the OpenShift cluster health, scheduling, API access, and authentication. The etcd cluster runs on the OpenShift master nodes.

Table 6 lists the hardware configuration in each of the HPE ProLiant DL360 Gen10 servers used in this solution.

TABLE 6. Configuration of each Master Nodes (HPE ProLiant DL360 Gen10 server)

Component	Description
Processors	2 x Intel Xeon-Gold 5218 (2.3 GHz/16-core/125 W)
Memory	4 x HPE 8GB 1Rx8 PC4-2933V-R Smart Kit
Network	HPE IB FDR/EN 40/50Gb 547FLR 2QSFP Adapter
Array Controller	HPE Smart Array P408i-a SR Gen10 Controller
Disks	HDD1: 300GB HDD 2: 300GB HDD 3: HPE 2.4TB SATA MU SFF SC DS SSD

HPE Aruba 8325-32C BF

The Aruba CX 8325 Switch is an enterprise-class, game-changing solution, offering a flexible approach to deal with the new application, security, and scalability demands of the mobile, cloud, and IoT era. With 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

Aruba 6300M 24-port

The 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 Aruba Network Analytics Engine, the Aruba CX 6300 extends industry-leading monitoring and troubleshooting capabilities to the access layer. Support of Aruba NetEdit and the Aruba CX Mobile App verify that configurations are flawless and easy to deploy.

A powerful Aruba Gen7 ASIC architecture delivers fast, non-blocking performance, meaning your network is ready for tomorrow's unpredictable demands. Aruba Virtual Stacking Framework (VSF) allows for stacking of up to 10 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. HPE Aruba 6300M 24-port

HPE Nimble Storage

HPE Nimble Storage AF40 is used to provide persistent, block storage in this solution. The HPE Nimble Storage array for Docker data provides the storage volume, hosts the repository, stores container images, and provides persistent volume for applications. Using the HPE-CSI-driver for Kubernetes, HPE Nimble Storage can provide persistent volume for the Kubernetes cluster in this solution. With HPE Nimble Storage, DevOps, IT teams, and others can be assured that their environment exceeds enterprise-grade requirements for supporting backup or protection and disaster recovery of data.

Software

Red Hat OpenShift Container Platform

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 Linux CoreOS

Red Hat OpenShift Container Platform uses Red Hat Enterprise Linux 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. RHCOS is specifically designed for running containerized applications from the RHOCP and works with new tools to provide fast installation, operator-based management, and simplified upgrades. For Red Hat OpenShift Container Platform 4.9 deployment on bare metal infrastructure, you must use RHCOS for all control plane machines, Bootstrap node, and worker nodes.

Red Hat OpenShift Data Foundation 4.9

Red Hat OpenShift Data Foundation (ODF) is software-defined storage that is optimized for container environments. It runs as an operator on Red Hat OpenShift Container Platform 4 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 databases, shared file storage for continuous integration, messaging, and data aggregation, and object storage for archival, backup, and media storage.

Table 7 lists the major software used in this solution.

TABLE 7. Software used in this solution

Component	Versions
Red Hat CoreOS	4.9
Red Hat OpenShift Container Platform	4.9
Red Hat OpenShift Data Foundation	4.8

CAPACITY AND SIZING

Sizing for a Red Hat OpenShift Container Platform 4.9 environment varies depending on the requirements of the organization and the type of deployment. The installer should read and understand Red Hat's recommendations around scalability and performance before installation. This ensures the need for their environment is addressed. Red Hat publishes documentation around scalability and performance for each OpenShift Container Platform release. For more information on the documentation for OpenShift Container Platform 4.9, see https://docs.openshift.com/container-platform/4.9/scalability_and_performance/recommended-host-practices.html.

Analysis and recommendations

This solution described in this document provides the guidance for deploying the RHOCP 4.9 cluster with three master nodes and three worker nodes using bare metal HPE ProLiant DL360 Gen10 and HPE ProLiant DL380 Gen10 servers. However, based on the customer requirements on the workloads or setting up an external cluster, etc. the architecture and the sizing can change.

The solution described in this paper uses the internal storage on HPE ProLiant DL360 Gen10 and HPE ProLiant DL380 Gen10 servers. Optionally, six (6) more HPE ProLiant DL380 Gen10 servers can be added to the cluster for the Red Hat Ceph cluster for Red Hat container external storage. HPE Ezmeral Data Fabric or HPE Nimble storage can be used for external storage. Deploying the OCP cluster on bare metal eliminates the overhead associated with hypervisors and thus optimizes performance. Additionally, the use of automation scripts saves significant efforts, resulting in quicker deployment.

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 master and worker nodes in the RHOCP 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 RHOCP 4.9 cluster also need to have internet access to perform the following tasks:

- 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
- Access Quay.io to obtain the packages that are required to install your cluster

- Obtain the packages that are required to perform cluster updates

Table 8 lists the various networks used for this solution. All the cluster nodes and iPXE servers are connected on the same twentynet network. Twentynet satisfies the network requirements for the deployment of OCP 4.9.

TABLE 8. Networks defined for the OCP 4.9

Network Name	Type	VLAN Number	Purpose
Management	Ethernet	3111	Solution management
Data_Center	Ethernet	2111	Application access, authentication, and other user networks
iPXE_Boot	Ethernet	2111	Boot for compute
iSCSI_VLAN_A	Ethernet	1111	iSCSI_VLAN_A
iSCSI_VLAN_B	Ethernet	1111	iSCSI_VLAN_B

Solution cabling

Figure 8 below shows the wiring between the compute servers and network switches used in this solution. The various links include iLO management links for servers, IRF links between the network switches, and the internal network links.

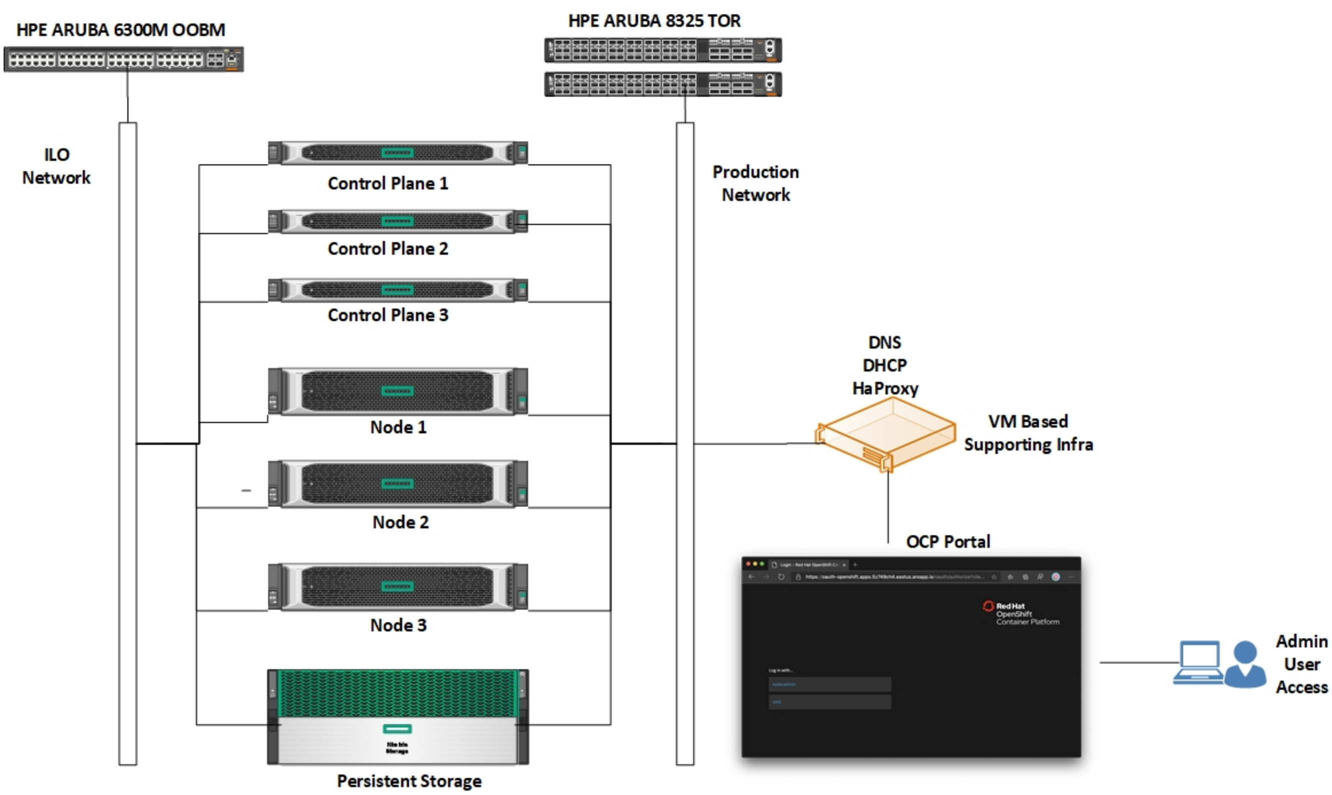

FIGURE 8. Wiring between the servers and the network switches

Storage

In the internal storage mode, the storage for the OS and internal persistent volume is provided by the local storage disks (SSD/HDD) on the HPE ProLiant DL360 Gen10 or HPE ProLiant 380 Gen10 Plus servers. Whereas in the external storage mode, the OS storage is provided by local disks, and the container storage is provided by the HPE Storage system such as HPE Nimble or HPE Ezmeral Data Fabric or by Red Hat OpenShift Container Platform that uses the local disks.

Table 9 lists all volumes used within the solution for the storage systems and highlights what storage provides the capacity and performance for each function.

TABLE 9. Details of the volume

Source	Volume/Disk Function	Hosts	Shared/Dedicated
Local storage on the Servers	OpenShift Container Volume	OpenShift worker nodes	Dedicated
	Operating System	All Nodes	Dedicated
HPE Nimble	iSCSI Persistent Volume	OpenShift worker nodes	Dedicated
HPE ProLiant DL380 Gen10 (External storage ODF)	Ceph Persistent Volume	OpenShift worker nodes	Dedicated

Figure 9 shows the Logical storage layout.

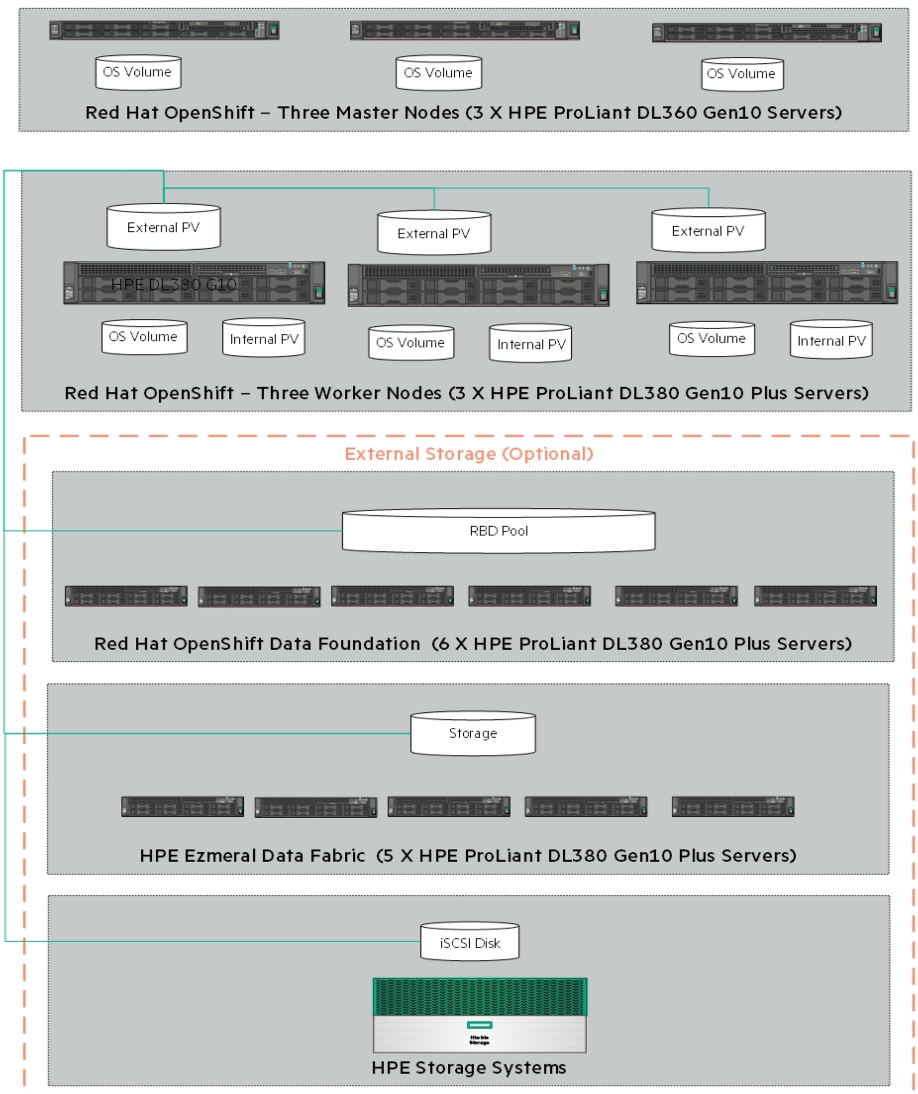

FIGURE 9. Logical storage layout in Solution

DEPLOYMENT OVERVIEW

This section explains in detail the deployment of OCP 4.9 using both the internal storage mode and external storage mode. In the internal storage mode, the local storage for OS and internal Persistent volume is provided by the local disks on the HPE ProLiant DL360 Gen10 or HPE ProLiant DL380 Gen10 servers.

In the external storage mode, a Ceph cluster is set up using six (6) HPE ProLiant DL380 Gen10 servers and integrated with the ODF. The external persistent storage could also be provided by HPE external storage such as HPE Nimble connected through the iSCSI network to the OCP cluster.

Deploying the OCP 4.9 cluster using the User Provisioned Infrastructure (UPI)

The OpenShift Container Platform 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 openshift-installer program to create RHCOS ignition configuration files. These ignition files include bootstrap ignition file, master ignition file, worker ignition file. The ignition files are used to configure RHCOS on each of the master and worker nodes in the OpenShift cluster. For detailed installation and configuration information, refer to the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/. For the Red Hat OpenShift document, see https://access.redhat.com/documentation/en-us/openshift_container_platform/4.9/html/installing_on_bare_metal/installing-on-bare-metal.

Figure 10 explains the OCP4.9 deployment process.

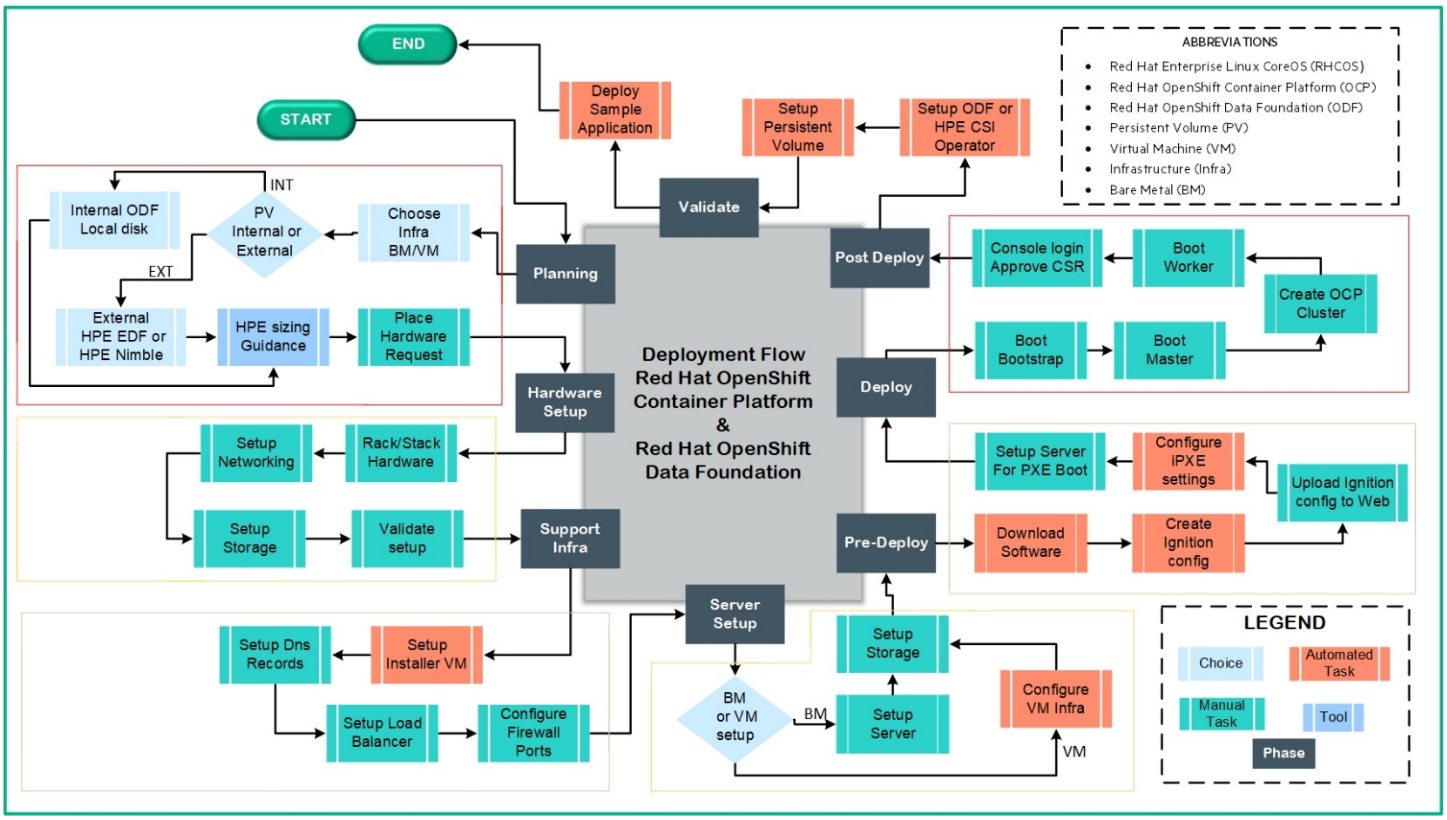

FIGURE 10. Deployment process flow diagram for OCP 4.9 cluster using the User Provisioned Infrastructure

---

NOTE

The load balancer described in this document is HAProxy.

HPE Solution Differentiator with User Provisioned Infrastructure

Red Hat OpenShift Container Platform 4.9 on HPE ProLiant DL Servers and HPE Storage systems solution leverages automation scripts to reduce the effort involved in setting up a manual User Provision Infrastructure based Red Hat OpenShift Container Platform 4.9 solution. The automation script help save time and also fosters accuracy by decreasing the number of steps involved in setting up the solution. 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 graph shown in Figure 11 depicts the graphical time difference in forming a manual vs automated deployment of Red Hat OpenShift Container Platform on bare metal servers using scripts mentioned in this document.

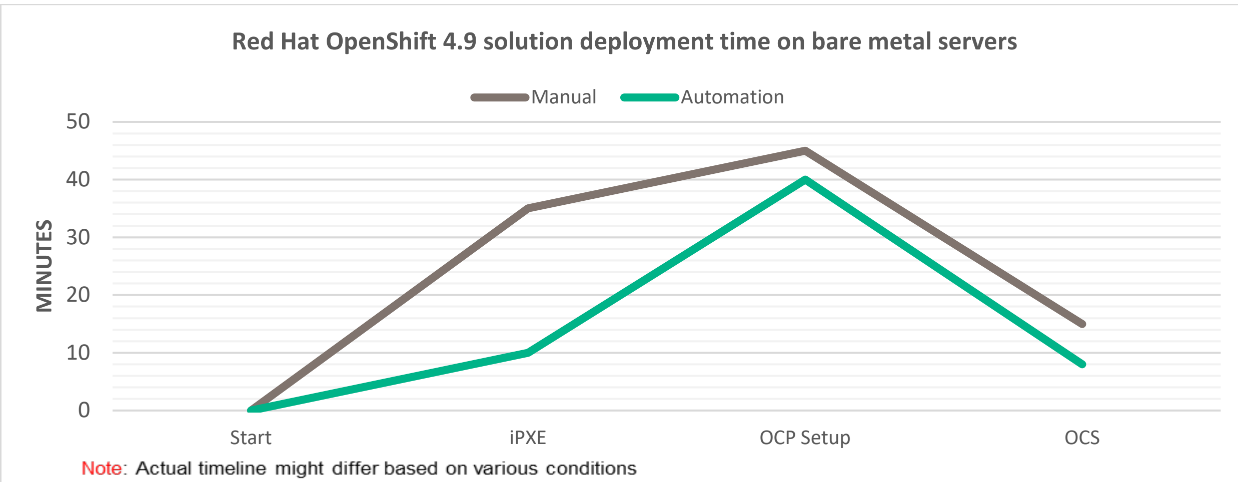

FIGURE 11. Red Hat OpenShift 4.9 solution deployment manual and automation timelines on bare metal

The graph shown in Figure 12 depicts the graphical time difference in forming a manual vs automated deployment of Red Hat OpenShift Container Platform on virtual machines using scripts mentioned in this document.

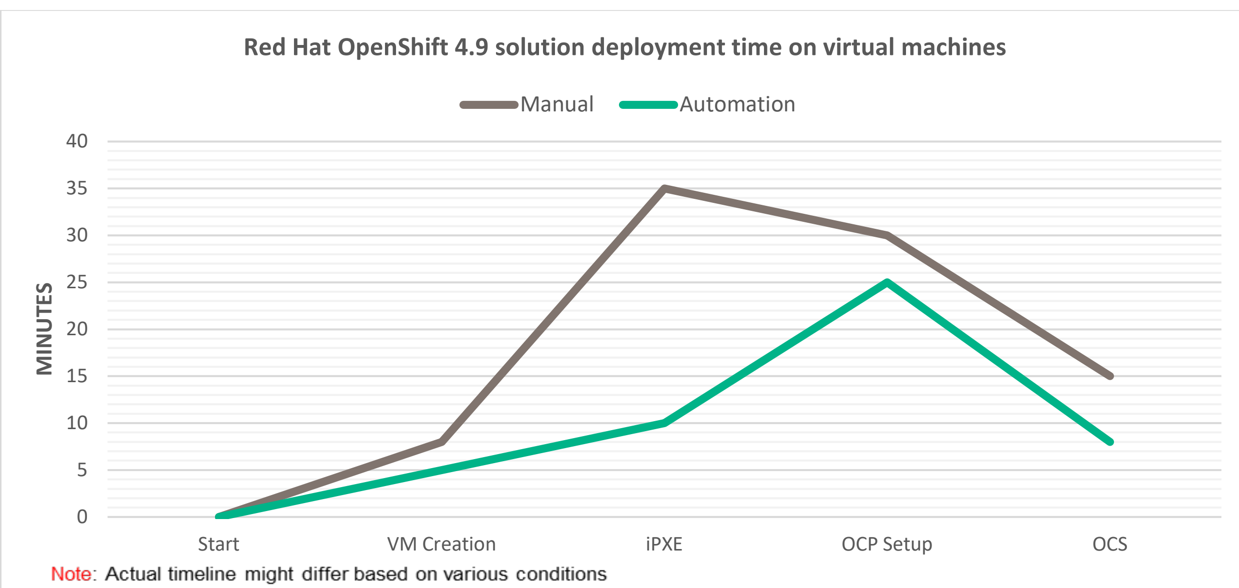

FIGURE 12. Red Hat OpenShift 4.9 solution deployment manual and automation timelines on virtual machines

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

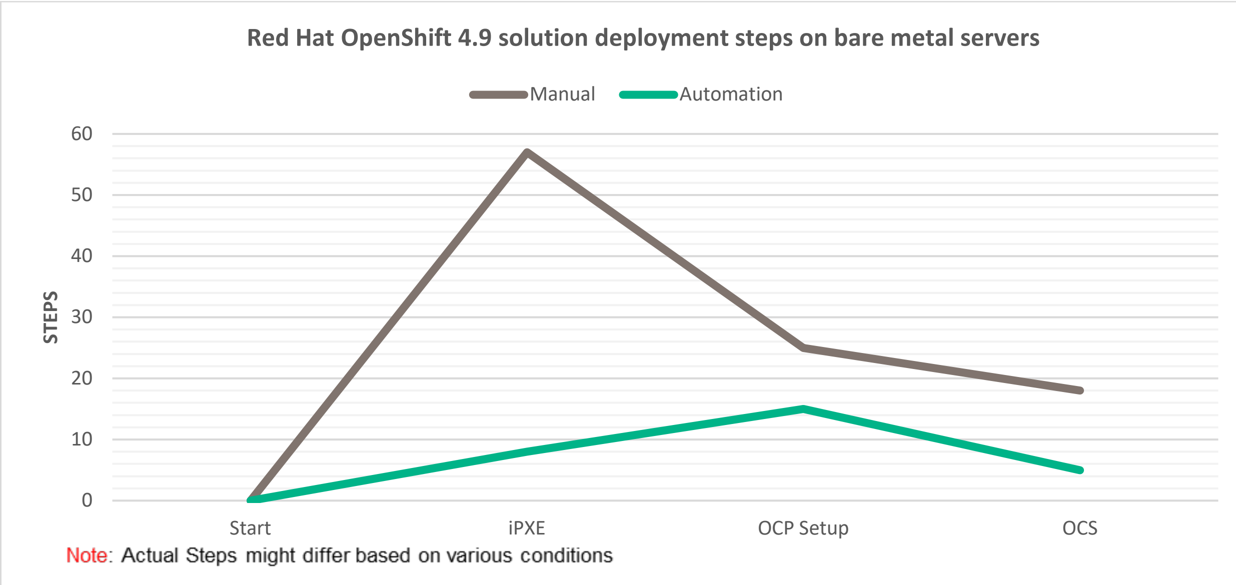

FIGURE 13. Red Hat OpenShift 4.9 solution deployment manual and automation steps on bare metal

The graph shown in Figure 14 depicts graphically the steps involved in setting up a manual vs automated deployment of Red Hat OpenShift Container Platform on virtual machines using scripts mentioned in this document.

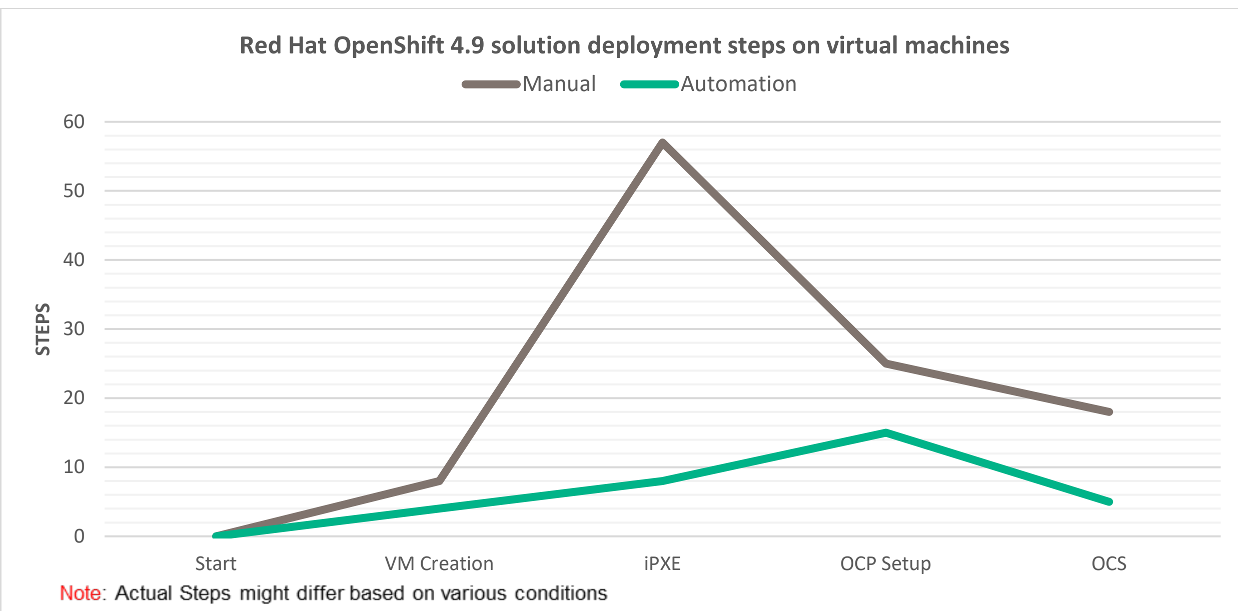

FIGURE 14. Red Hat OpenShift 4.9 solution deployment manual and automation steps on virtual machines

The graphs are shown in this section quantify the time saved and the steps reduced in our lab setup. The graph serves as a reference and the time or the steps involved might differ depending on various environmental factors such as Infrastructure complexity, user proficiency with OpenShift, etc. The key is using the automation scripts along with this Reference Architecture and Deployment Guide will ensure improved business productivity.

Deploying the OCP 4.9 cluster using UPI in a restricted network

A restricted or disconnected network deployment sometimes also referred to as the air-gapped network is a network deployment that has security implemented on one or more network segments within a data center to ensure there is restricted or no network access, to and from unsecured or less secured networks, such as the public Internet or a general local area network. Red Hat OpenShift Container Platform 4.9 can be deployed in such a restricted network on User Provisioned Infrastructure (UPI) only. To set up a restricted network installation, the key is having a registry that mirrors the contents of the OpenShift Container Platform registry on the internet and contains the installation media files required. This registry on a mirror host can be set up on a host that has access to the restricted network. The required files for setting up the mirror registry could be made available using methods such as file copy/transfer that meet the environment restrictions can be followed. Red Hat OpenShift Clusters in restricted networks have the following additional limitations and restrictions:

- The Cluster version status includes an unable to retrieve available updates error
- By default, you cannot use the contents of the Developer Catalog because you cannot access the required ImageStreamTags

Figure 15 shows a high-level solution overview of deploying Red Hat OpenShift 4.9 in a restricted network.

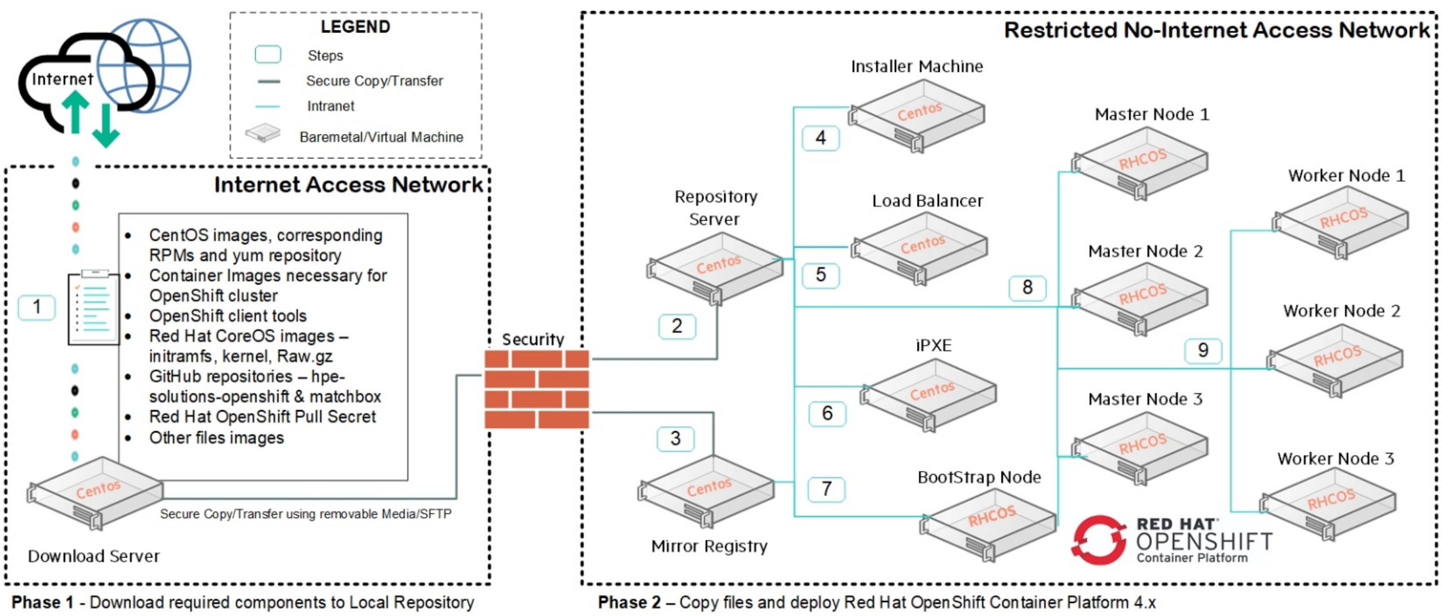

FIGURE 15. Red Hat OpenShift Container Platform deployment in a restricted network – Logical layout

The high-level deployment process is divided into two phases:

- Phase 1 – Download required components to a local repository that has access to the internet.
- Phase 2 - Copying/transferring files to a repository server in the restricted network, setting up the registry, and then deploying the OCP.

Figure 16 shows a high-level flow for deploying Red Hat OpenShift Container Platform 4.9 in a restricted network.

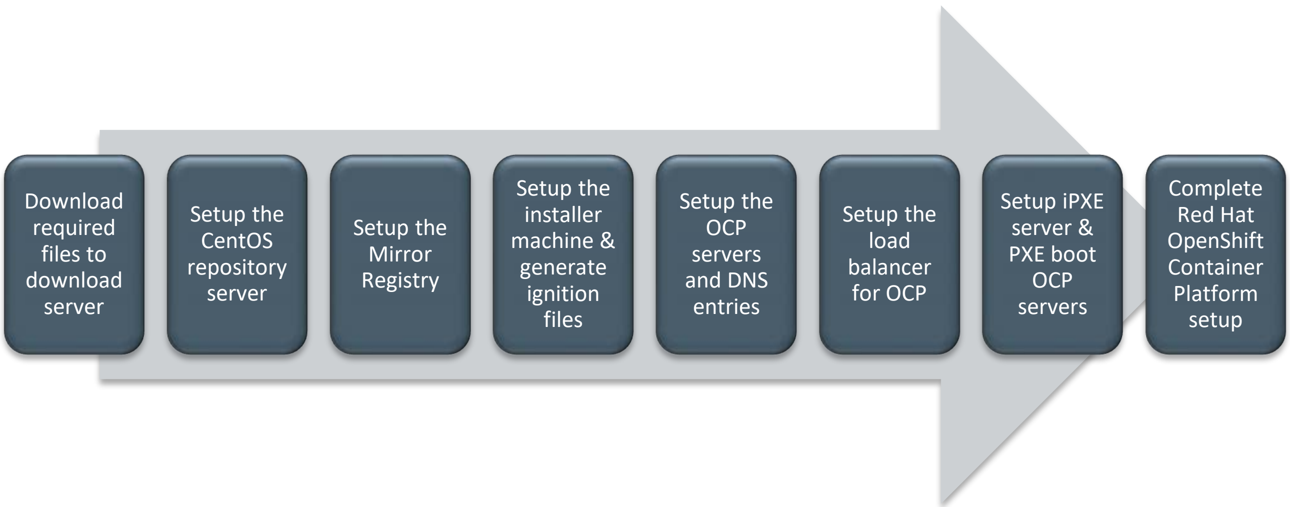

FIGURE 16. Red Hat OpenShift Container Platform deployment in a restricted network – High-level setup flow

For more information, see the Deployment Guide.

Red Hat OpenShift Data Foundation

Red Hat OpenShift Data Foundation is deployed as an operator with a minimal cluster of three(3) worker nodes. 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 OpenShift Data Foundation in our test environment was configured on the virtualized setup. The details of ODF configuration and procedure on storage sizing are described in the Deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

HPE Ezmeral Data Fabric CSI

The CSI Storage Plugin is a volume driver that uses the industry-standard container-storage interface to expose the HPE Ezmeral Data Fabric to workloads on container-orchestration systems. Once deployed HPE Ezmeral Data Fabric CSI manages HPE Ezmeral CSI Drivers on Kubernetes and OpenShift for static and dynamic provisioning of persistent volumes on the HPE Ezmeral Data Fabric platform.

HPE 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 by using the CSP specification. This makes it very easy for third parties to integrate their storage solution 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.

Figure 17 shows the HPE CSI Driver for Kubernetes architecture.

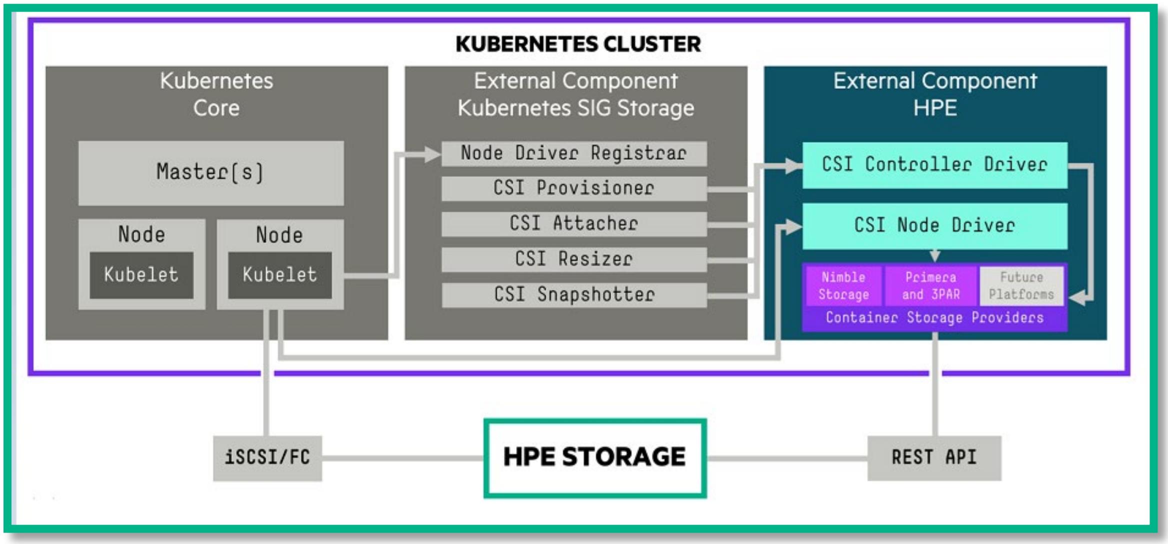

FIGURE 17. HPE CSI driver for Kubernetes architecture

For more information, see https://scod.hpedev.io/csi_driver/index.html.

Physical worker node labeling in OpenShift

In this document so far, we have discussed setting up the OCP cluster and then setting up the persistent volume. This is followed by discovering the node properties and advertising them through node labels that can be used to control workload placement in an OpenShift cluster. OpenShift does not label nodes by default with any hardware configuration information. If IT wants to use hardware configuration to optimize scheduling, the capabilities of the underlying platform must be manually uncovered and labeled by administrators to use the hardware configuration in scheduling decisions. An OpenShift cluster can have many nodes. Each node in turn can run multiple pods which, at scale, means that this process is both tedious and error-prone. With OpenShift running on the HPE server platform, organizations can automate the discovery of hardware properties and use that information to schedule workloads that benefit from the different capabilities that the underlying hardware provides. Using HPE iLO and its REST or Redfish® API-based discovery capabilities (Proliantutils), the following properties can be discovered about the nodes:

- Presence of GPUs
- Underlying RAID configurations
- Presence of disk by type
- Persistent-memory availability
- Status of CPU virtualization features
- SR-IOV capabilities
- CPU architecture
- CPU core count
- Platform information including model, iLO, and BIOS versions
- Memory capacity
- UEFI security settings
- Health status of compute, storage, and network components

After these properties are discovered for the physical worker nodes, node labeling can be applied to facilitate grouping nodes based on the underlying features of those hosts. Labels do not provide uniqueness. In general, it is expected that many objects will carry the same label(s). Using a label selector, the administrator can identify a set of objects with similar properties. This labeling can be used as either a hard or soft constraint for scheduling application pods on the desired node based on application requirements. For example, if the compute module in the HPE ProLiant DL Server Infrastructure must support Intel TXT, which is specifically designed to harden the platform from the emerging threats of hypervisor attacks, malicious rootkit installations, or other software-based attacks. Administrators can use this information to restrict confidential data or sensitive workload to nodes that are better controlled and have their configurations more thoroughly evaluated using the Intel TXT-enabled platform. For more information about node labeling configuration, see the HPE solutions for Red Hat OpenShift GitHub at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

Securing and monitoring OpenShift with Sysdig SaaS

To address the security challenges that exist in containerized environments, this solution leverages the Sysdig SaaS Platform to secure and monitor the Red Hat OpenShift Container Platform, an enterprise-ready Kubernetes platform that is installed and configured on HPE Compute Infrastructure. After the configuration is deployed, access to the Red Hat OpenShift cluster is granted to the Sysdig SaaS Platform. The Sysdig SaaS Platform is a cloud-based service where security and monitoring services will be available to the user based on the choice of subscription. For security and monitoring of Red Hat OpenShift Containers, it is required to install the Sysdig Agent on the OpenShift cluster. This means Sysdig Agents that are lightweight entities will be installed within each node in the OpenShift cluster. These agents run as a daemon to enable Sysdig Monitor and Sysdig Secure functionality. Sysdig Monitor provides deep, process-level visibility into a dynamic, distributed production environment. Sysdig Secure provides image scanning, run-time protection, and forensics to identify vulnerabilities, block threats, enforce compliance, and audit activity across an OpenShift cluster.

The key benefits are:

- Faster incident resolution using Sysdig Monitor for OpenShift cluster
- Simplified compliance for the entire solution
- Service-based access control for container security and monitoring
- Less time spent on managing platforms, containers, and vulnerabilities

The implementation of Sysdig in this solution uses the Software as a Service (SaaS) deployment method. The playbooks deploy Sysdig Agent software to capture the data from every node in the OpenShift deployment and the captured data is relayed back to the Sysdig SaaS Cloud portal. The deployment provides access to a 90- day try-and-buy fully-featured version of the Sysdig software. For more information about Sysdig Agent deployment in the OpenShift setup, see the HPE solutions for Red Hat OpenShift GitHub at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

NOTE
The Sysdig functionality is not turned on by default in this solution. For more information on how to enable Sysdig, see the Sysdig configuration section listed in HPE solutions for Red Hat OpenShift Container Platform GitHub at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

Post-validation deployment

After the successful deployment of Red Hat OpenShift Container Platform 4.9 on the identified nodes and access to the OpenShift console is successful, the following are deployed from the operator hub for monitoring the logs of the applications deployed on the cluster.

- Install EFK operators in OCP
- Install the Elasticsearch and Cluster Logging operators (Fluentd and Kibana)
- Install Cluster Logging operator
- Create cluster logging Instance
- Launch Kibana

Figure 18 shows the Kibana Dashboard.

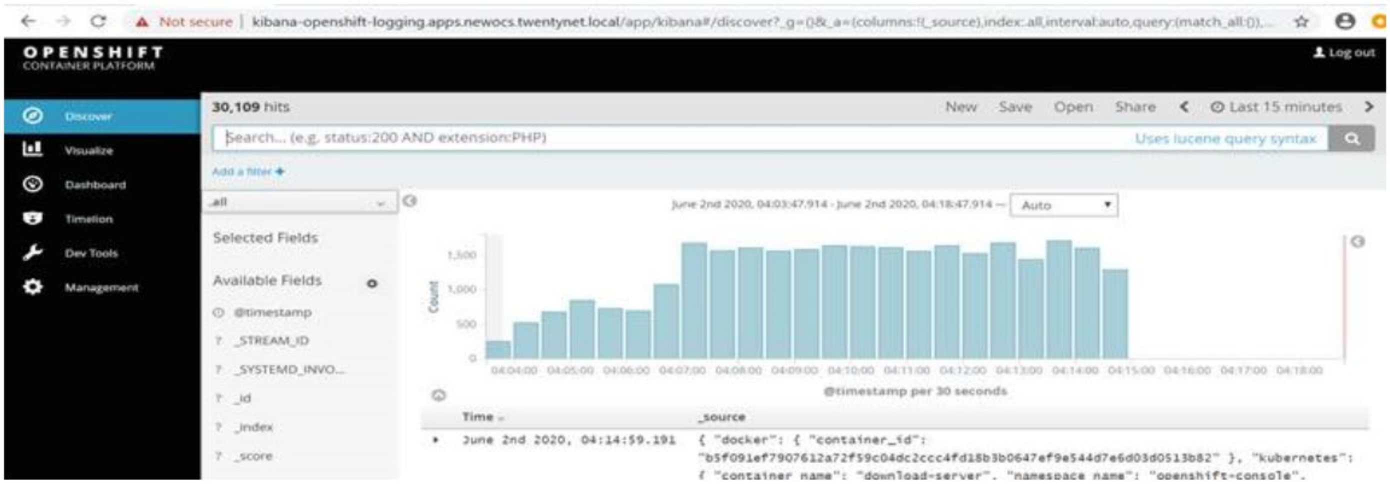

FIGURE 18. View of Kibana Dashboard

The details of installation and configuration of EFK operators, cluster operators, Kibana, etc. are described in the Deployment guide at hhttps://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

NOTE

Kibana is a browser-based console to query, discover, and visualize your logs through histograms, line graphs, pie charts, heat maps, built-in geospatial support, and other visualizations.

RED HAT ADVANCED CLUSTER MANAGEMENT FOR KUBERNETES

Red Hat Advanced Cluster Management for Kubernetes provides end-to-end management visibility and control to manage the Kubernetes environment. It allows the administrator to take control of the application modernization program with management capabilities for cluster creation, application lifecycle, and provide security and compliance for all of them across data centers and hybrid cloud environments. Clusters and applications are all visible and managed from a single console, with built-in security policies. The administrator can perform day-to-day operations from anywhere that Red Hat OpenShift runs, and manage any Kubernetes cluster. The following are components for Red Hat Advanced Cluster Management for Kubernetes:

- Hub cluster - The hub cluster is the common term that is used to define the central controller that runs in a Red Hat Advanced Cluster Management for the Kubernetes cluster. The hub cluster provides access to the management console. The hub cluster aggregates information from multiple clusters by using an asynchronous work request model.
- Managed cluster - The managed cluster is the term that is used to define additional clusters with the Klusterlet, which is the agent that initiates a connection to the Red Hat Advanced Cluster Management for the Kubernetes hub cluster.
- Cluster lifecycle - The process of creating, importing, and managing clusters across public and private clouds from the hub cluster console.
- Application lifecycle -The processes that are used to manage application resources on managed clusters.
- Governance and risk - The processes that are used to manage security and compliance from a central interface page.

Figure 19 shows an overview of the Red Hat Advanced Cluster Management.

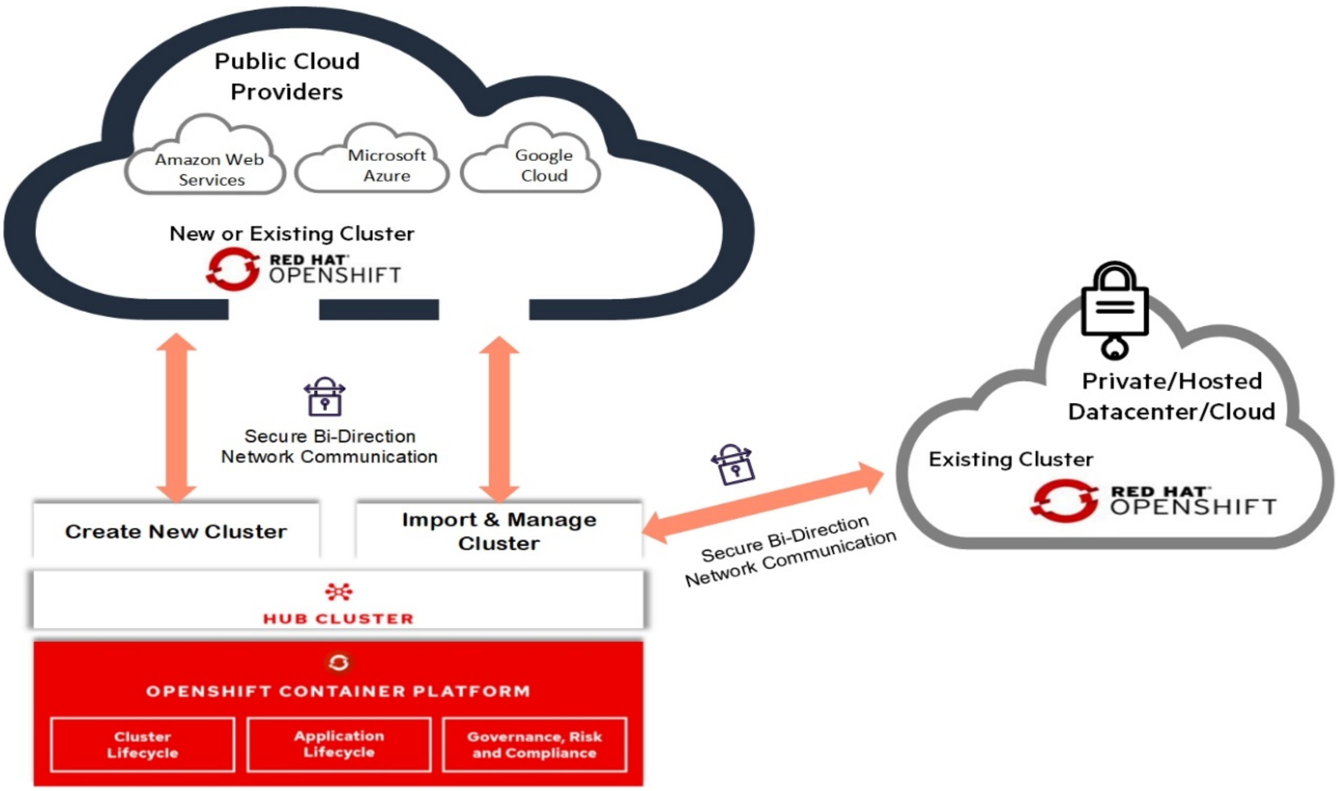

FIGURE 19. Red Hat Advanced Cluster Management Overview

Installing Red Hat Advanced Cluster Management for Kubernetes

High-level installation flow:

- Have a supported version of Red Hat OpenShift Container Platform installed and configured
- Install the Operator for Red Hat Advanced Cluster Management for Kubernetes from the catalog
- Configure Red Hat Advanced Cluster Management for Kubernetes

For more information, see the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

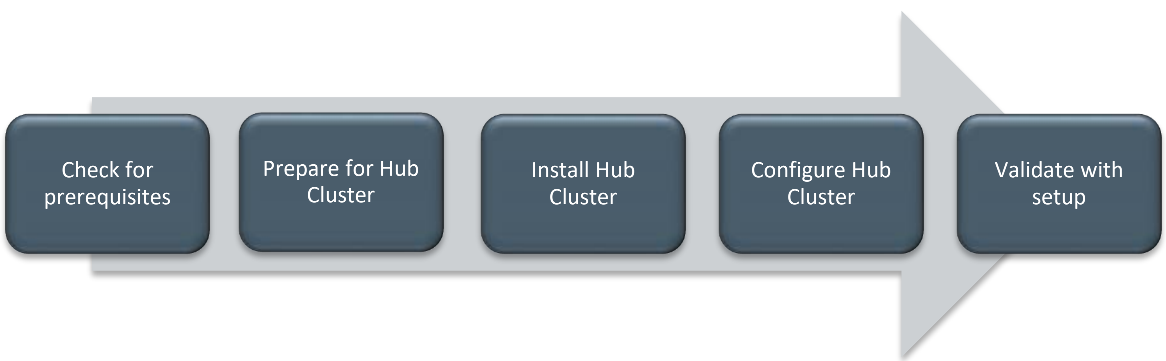

FIGURE 20. Red Hat Advanced Cluster Management Installation Overview

With Red Hat Advanced Cluster Management for Kubernetes:

- Work across a range of environments, including multiple data centers, private clouds, and public clouds that run Kubernetes clusters
- Easily create Kubernetes clusters and offer cluster lifecycle management in a single console
- Enforce policies at the target clusters using Kubernetes-supported custom resource definitions
- Deploy and maintain day-two operations of business applications distributed across your cluster landscape

BUSINESS CONTINUITY WITH DISASTER RECOVERY STRATEGIES FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.9

Stateful Applications require 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 4.9 deployment. Recovery Time Objective (RTO) is the organization's tolerance for "App Downtime" and Recovery Point Objective (RPO) is the organization's tolerance for "Data Loss" are two key metrics that must be considered to develop an appropriate disaster recovery plan that can maintain business continuity after an unexpected event.

Figure 21 shows the comparison of the Red Hat OpenShift Disaster recovery strategies using RTO and RPO objectives.

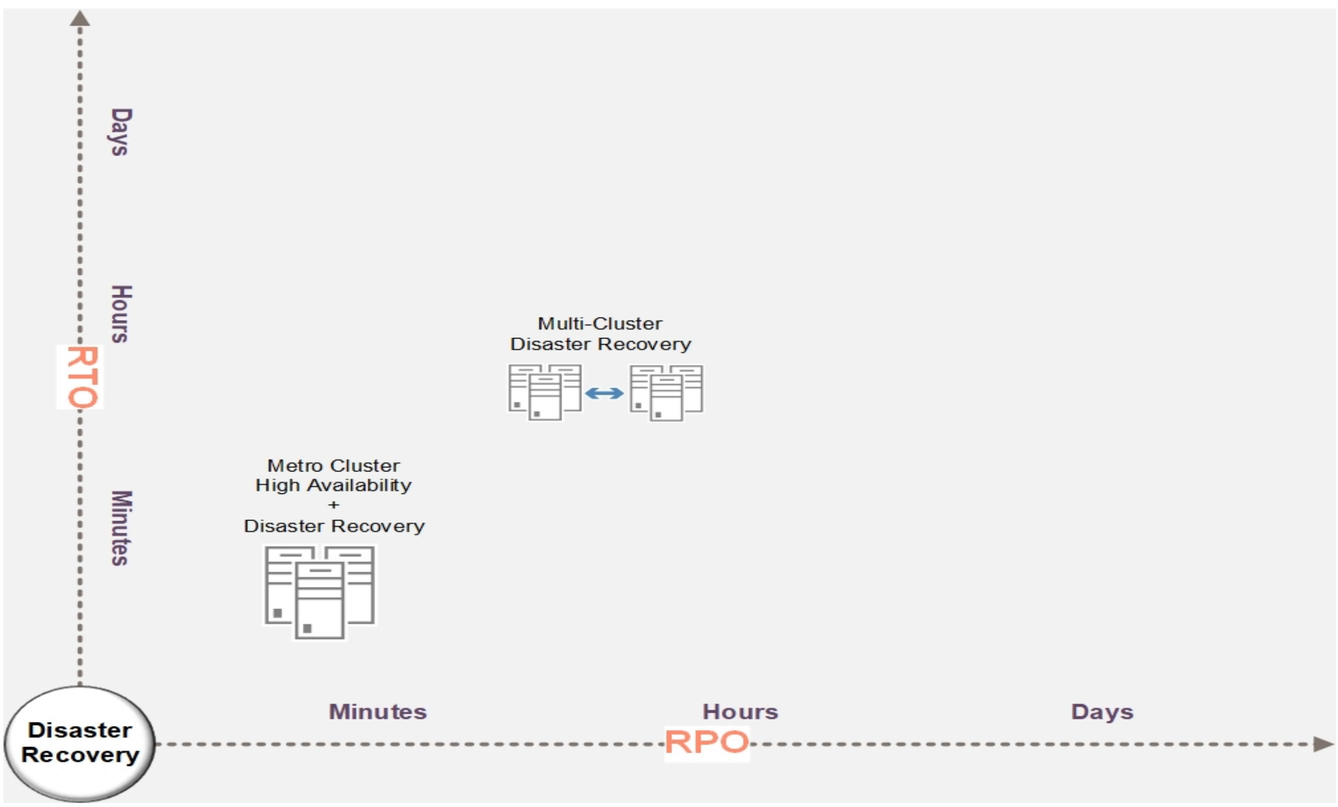

FIGURE 21. Red Hat OpenShift Disaster recovery strategies comparison using RTO and RPO objectives

Metro Cluster High Availability and Disaster Recovery Strategy

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

Following are the requirements for HA like Automatic Recovery along with No Data loss Data Mirroring:

- Synchronous HA-DR for localized data center failures
  - DR sites/Availability zones (AZs) connected by MAN or campus networks
  - AZs are mapped to a fault domain (HVAC, Power grids, etc.)
  - An odd number of AZ/Fault Domains are required for the cluster quorum
  - Network latency between zones does not typically exceed 5 ms
- OCP ensures Pods, nodes get scheduled across zones during deployment
- ODF ensures maintains consistent mirror copies across AZs => no data loss
- Stretched OCP cluster provides automatic and non-disruptive recovery for apps across AZs
- An application with a consensus protocol that allows it to determine which instances of the cluster are active and healthy

Figure 22 shows an overview Red Hat OpenShift Metro Cluster design.

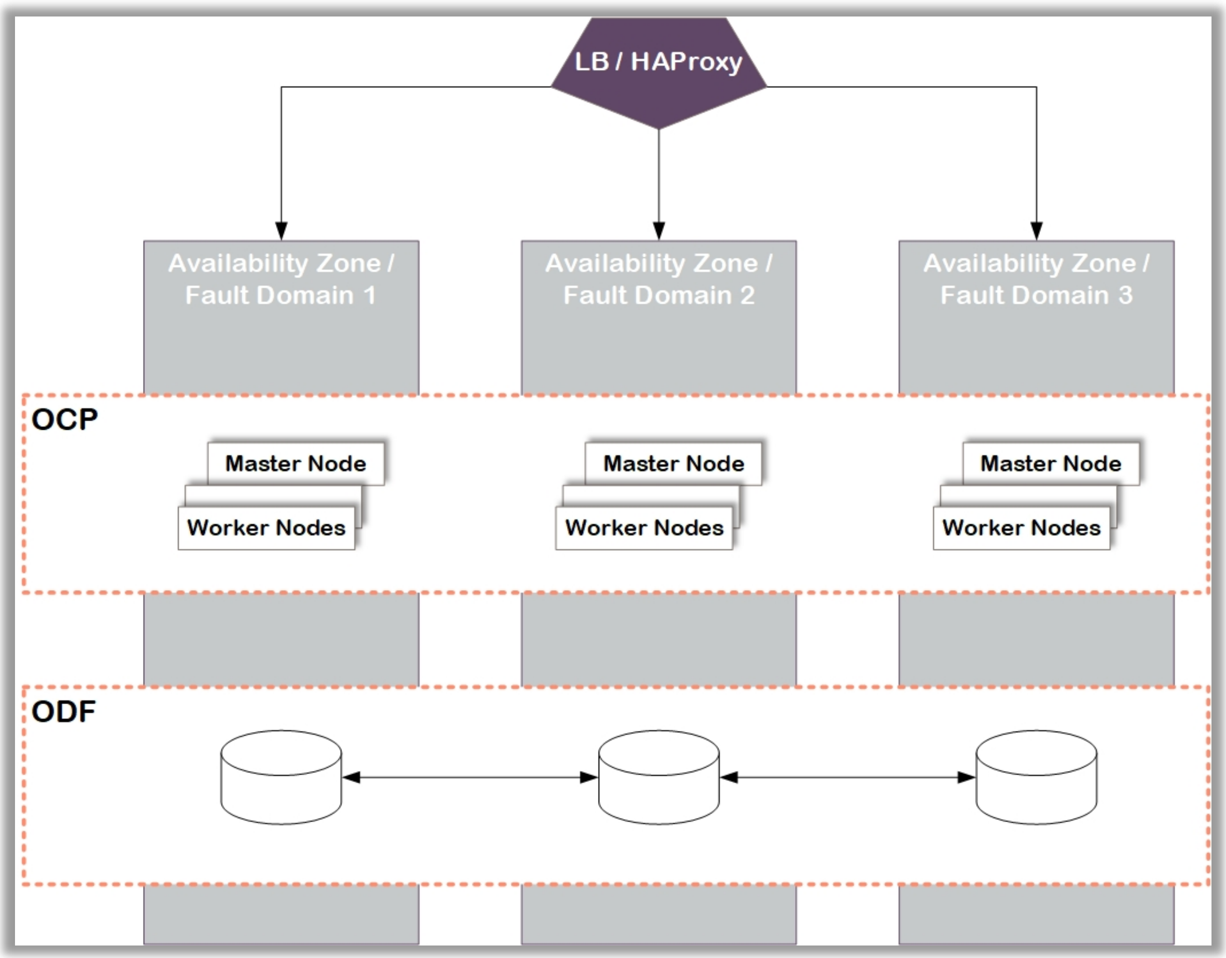

FIGURE 22. Red Hat OpenShift Metro Cluster design overview

When one of the AZs is down, no action needs to occur as both 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 potentially geographically distributed in each data center and have their own independent OpenShift clusters. A global load balancer balances traffic between the data centers. The stateful workload is deployed across the 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 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 23 shows the Red Hat OpenShift Multi-cluster disaster recovery approach.

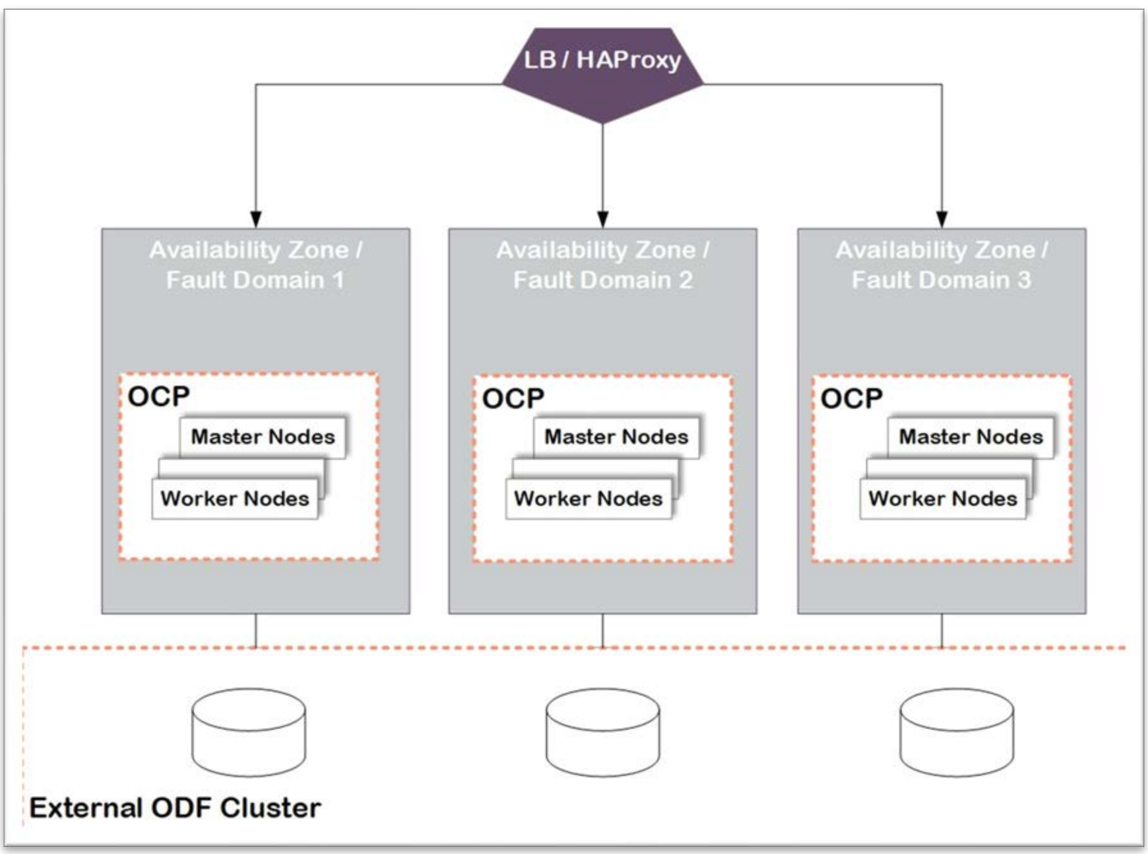

FIGURE 23. Red Hat OpenShift Multi-cluster disaster recovery approach

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

BUSINESS CONTINUITY WITH DATA PROTECTION (BACKUP) FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.9

Backup and restore is a day 2 operational task of 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 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 a while in Enterprise Datacenter's ecosystem, these solutions need to evolve to address the needs of the new container infrastructure this where Velero adds value. Velero is an open-source tool to safely backup and restore, perform disaster recovery, and migrate Kubernetes cluster resources and persistent volumes.

Velero provides the following feature to the Kubernetes bases 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 4.9, Velero uses a controller model where it monitors custom resources and takes actions.

Figure 24 shows the overview of Velero backup and restore with OCP 4.9.

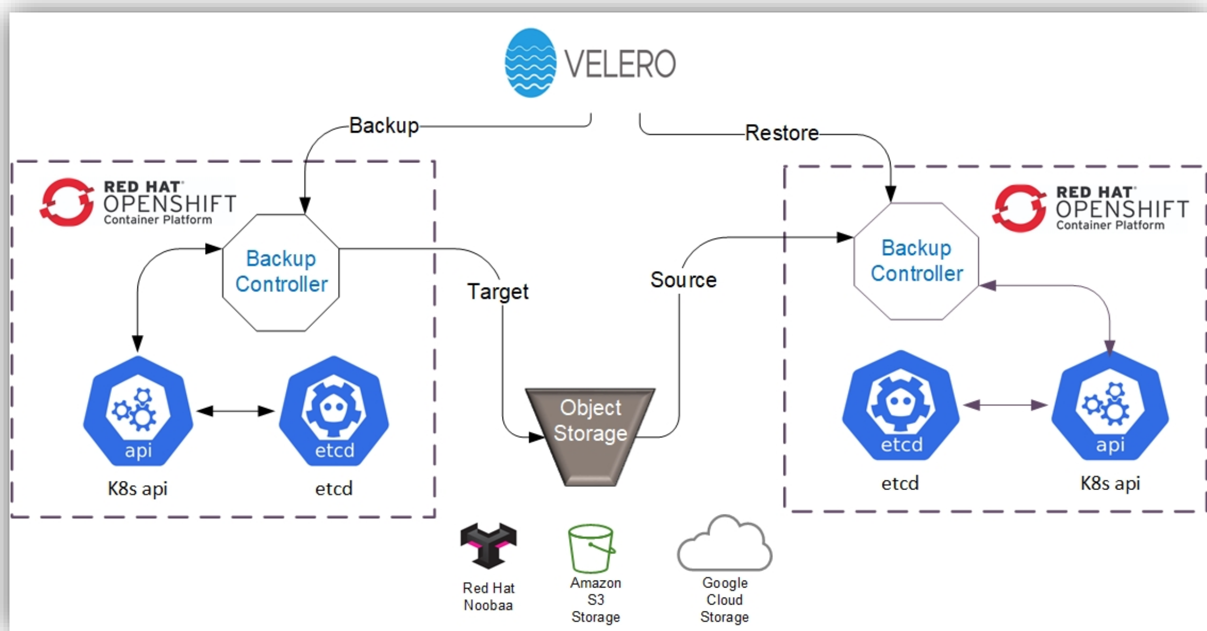

FIGURE 24. Velero backup and restore with OCP 4.9

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

Prerequisites

- OpenShift Container Platform cluster should be available with the administrator credentials
- When using NooBaa as a target we need external DNS of NooBaa and access Id and key

- When using public cloud-based object storage require the appropriate Velero plugin along with the access information and credentials

Velero for OpenShift setup overview

Figure 25 shows the overview of Velero for the OpenShift setup.

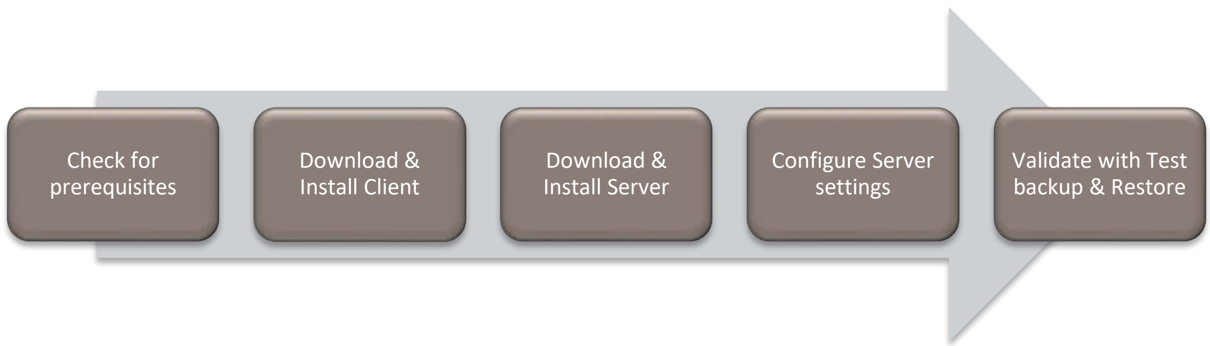

FIGURE 25. Velero for OpenShift setup overview

Velero makes it simple to back up the 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/49-dl/.

OPENSHIFT VIRTUALIZATION

The new paradigm shift in application development is moving to containers and is increasingly serverless as well, however, existing investment in applications that run as virtual machines in many organizations poses management challenges. Many of the existing virtual machines (VMs) provide vital services to new and existing containerized applications at times are not good candidates for containerization. OpenShift Virtualization also referred to as container-native virtualization lets developers bring virtual machines (VMs) into containerized workflows by running a virtual machine within a container where they can develop, manage, and deploy virtual machines side-by-side with containers and serverless all in one platform. This is a built-in feature of OpenShift Container Platform, OpenShift Virtualization combines two technologies of virtualization and containerization into a single management platform so organizations take advantage of the simplicity and speed of containers and Kubernetes while still benefiting from the applications and services that have been architected for virtual machines.

Figure 26 shows an overview of OpenShift virtualization.

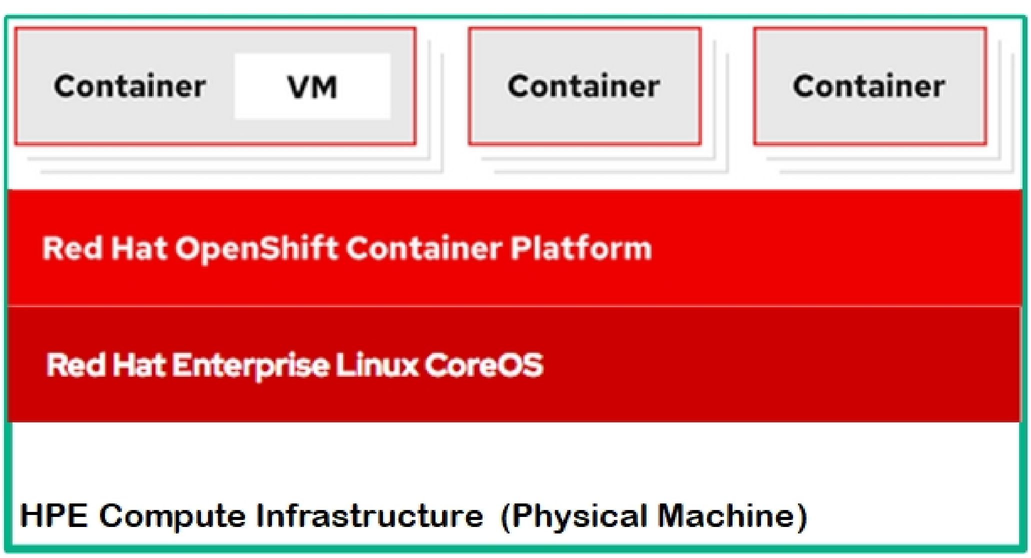

FIGURE 26. OpenShift virtualization overview layout

OpenShift virtualization adds new objects into your OpenShift Container Platform cluster via Kubernetes custom resources to enable virtualization tasks. These tasks include:

- Creating and managing Linux and Windows® virtual machines
- Connecting to virtual machines through a variety of consoles and CLI tools
- Importing and cloning existing virtual machines
- Managing network interface controllers and storage disks attached to virtual machines
- Live migrating virtual machines between nodes
- An enhanced web console provides a graphical portal to manage these virtualized resources alongside the OpenShift Container Platform cluster containers and infrastructure
- OpenShift Virtualization is tested with OpenShift Data Foundation (ODF) and designed to use with ODF features for the best experience
- OpenShift Virtualization allows the usage with either the OVN-Kubernetes or the OpenShiftSDN default Container Network Interface (CNI) network provider
- OpenShift Virtualization 2.4 is supported for use on OpenShift Container Platform 4.9 clusters

Enabling OpenShift Virtualization

Figure 27 shows an overview of OpenShift virtualization.

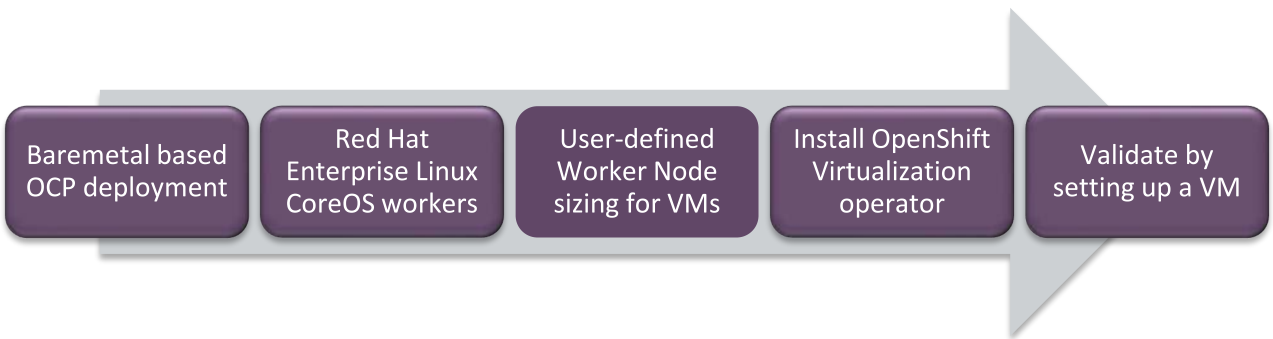

FIGURE 27. OpenShift virtualization setup

OpenShift Virtualization adds the ability to easily create and manage traditional virtual machines in OpenShift alongside standard container workloads. Its deep integration into the OpenShift UI makes the first step very easy and intuitive. For more information, see OpenShift Virtualization documentation.

RED HAT OPENSHIFT DATA FOUNDATION– EXTERNAL MODE

Red Hat OpenShift Data Foundation can use an externally hosted Red Hat Ceph Storage cluster as the storage provider. Red Hat Ceph Storage is designed for cloud infrastructure and web-scale object storage. Red Hat Ceph Storage is a scalable, open, software-defined storage platform that combines an enterprise-hardened version of the Ceph storage system with a Ceph management platform, deployment utilities, and supports services.

Figure 28 shows a high-level solution overview of Red Hat OpenShift Data Foundation– External Mode.

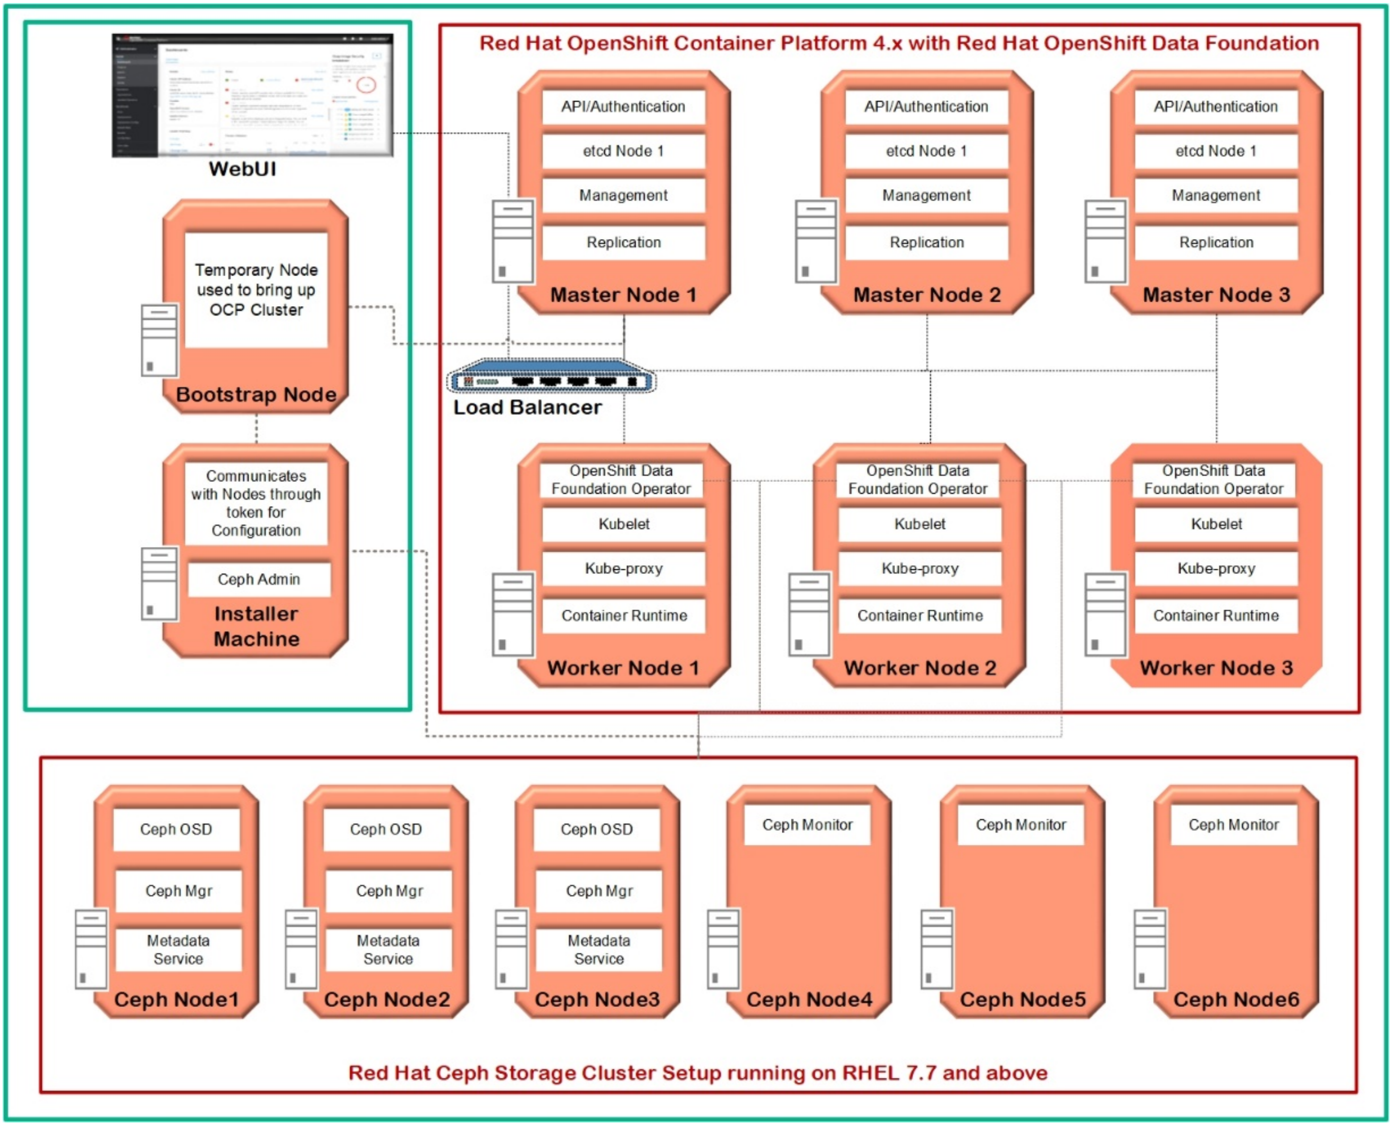

FIGURE 28. Red Hat OpenShift Data Foundation– External Mode Solution Layout

Setting up the Red Hat OpenShift Data Foundation in external mode allows an enterprise to disintegrate the container setup platform from the container compute platform. The disintegration allows on-demand independent growth of storage and compute clusters. Before configuring the Red Hat OpenShift Data Foundation in external mode the Red Hat Ceph Storage pool must be set up and configured appropriately. For this Reference Architecture, a six-node Red Hat Ceph storage cluster running on Red Hat Enterprise Linux 7.7 is used. The high-level steps involved in setting up Red Hat OpenShift Data Foundation in external is as follows:

- Setting up a prerequisite for Red Hat Ceph installation on three (3) bare metal servers
- Deploy and register Ceph Nodes
- Setup Ceph storage pool
- Configure OpenShift Data Foundation in external mode
- Validate the creation of PV on Red Hat OpenShift Data Foundation in external mode

Figure 29 shows the high-level setup.

FIGURE 29. Red Hat OpenShift Data Foundation– External mode solution setup flow

For more information, view the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

A single Red Hat Ceph Storage could be used to provide persistent volume to multiple OpenShift Container Platform clusters.

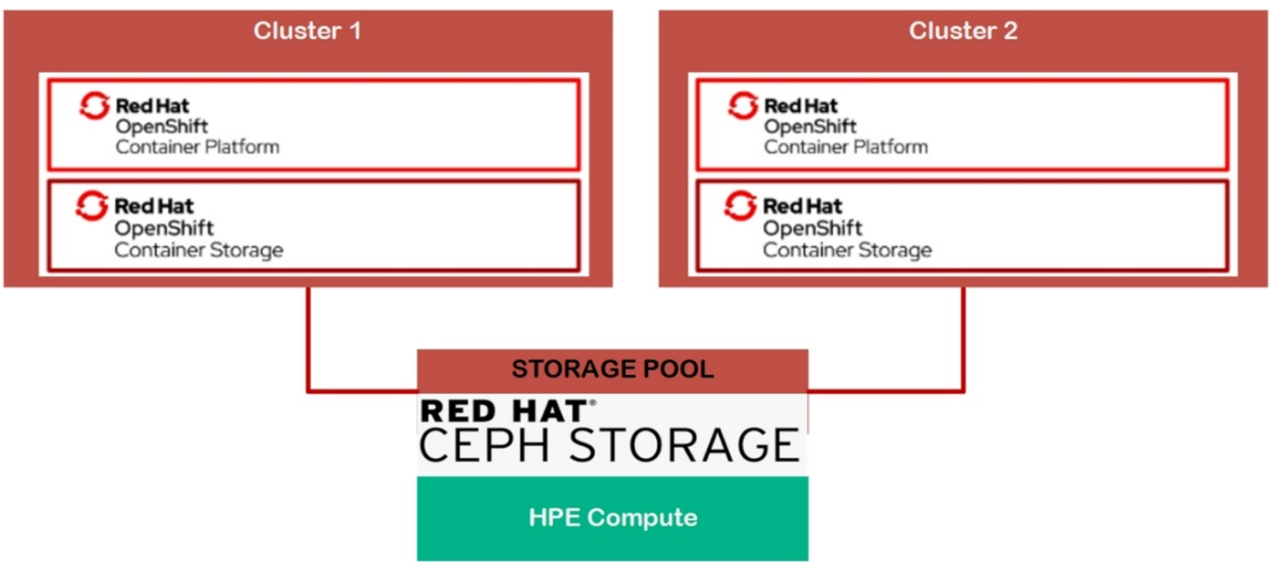

FIGURE 30. Red Hat OpenShift Data Foundation- Usage

RED HAT OPENSHIFT INSTALLER-PROVISIONED INSTALLATION FOR BARE-METAL

The OpenShift Container Platform installation program offers the flexibility of using the installation program to deploy a cluster on infrastructure that the installation program provisions and the cluster maintenance. With the installer-provisioned infrastructure clusters, OpenShift Container Platform manages all aspects of the cluster, including the operating system itself. Each machine boots with a configuration that references resources hosted in the cluster that it joins. This configuration allows the cluster to manage itself as updates are applied. When using the Red Hat OpenShift installer-provisioned installation, the installation program acts much as an installation wizard, requests input for values that it cannot determine on its own, and provides reasonable default values for the remaining parameters. In this section, we discuss the Installer-provisioned installation for installing OpenShift Container Platform on bare metal nodes.

Installer-provisioned installation of OpenShift Container Platform requires:

- One provisioner node with RHEL 8.1 installed
- Three Control Plane nodes

- At least two worker nodes
- Baseboard Management Controller (BMC) access to each node
- At least one network:
  - One required routable network

For more detailed information on Red Hat OpenShift Container Platform Installer-provisioned installation prerequisites, see https://dODF.openshift.com/container-platform/4.9/installing/installing_bare_metal_ipi/ipi-install-prerequisites.html.

Figure 31 shows the logical layout of Red Hat OpenShift Container Platform Installer-provisioned installation.

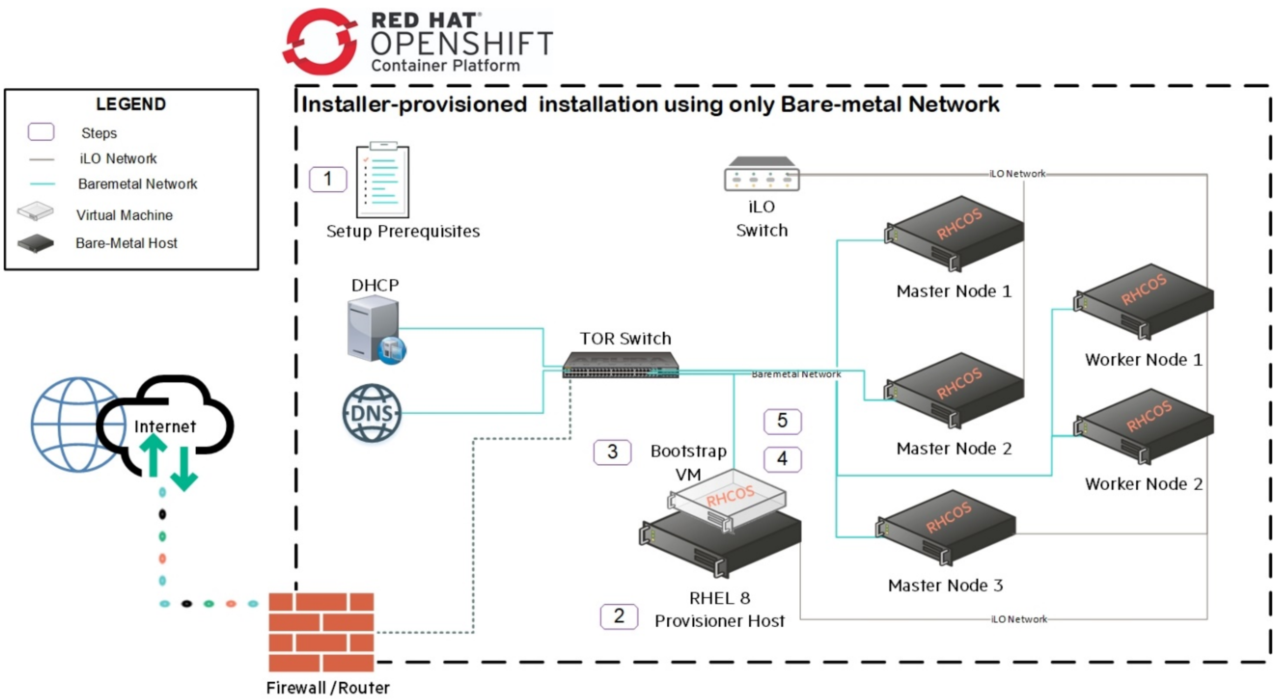

FIGURE 31. Red Hat OpenShift Container Platform Installer-provisioned installation – Logical layout

The high-level steps involved in setting up Red Hat OpenShift Container Platform Installer-provisioned infrastructure using three (3) master nodes, two (2) worker nodes, and using the bare metal network are as shown in Figure 32.

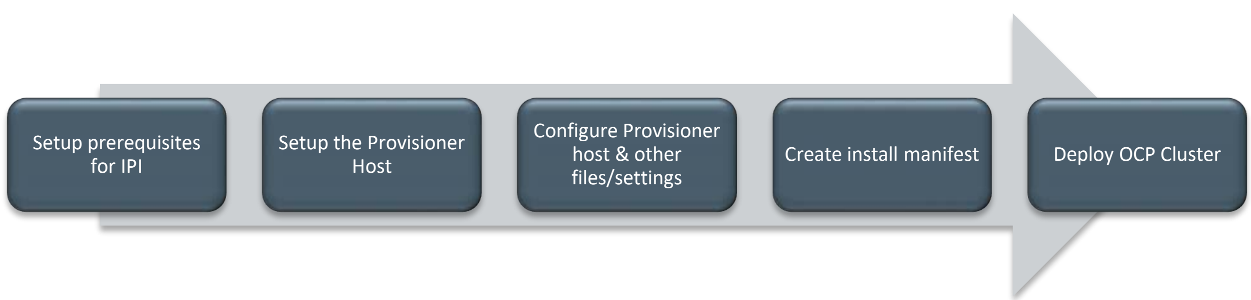

FIGURE 32. Red Hat OpenShift Container Platform Installer-provisioned installation – High-level setup flow

For more information, view the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/49-dl/.

SUMMARY

The installation and configuration of the Red Hat OpenShift Container Platform Version 4.9 with OpenShift Data Foundation on the bare-metal HPE ProLiant DL360 Gen10 and HPE ProLiant DL380 Gen10 servers solution provide customers with greater efficiency, higher utilization, and bare-metal performance by “collapsing the stack” and eliminating the need for virtualization. IT teams can manage multiple Kubernetes clusters with multitenant container isolation and data access, for any workload from edge to core to cloud. The benefits of containers beyond cloud-native microservices architected stateless applications can be extended by providing the ability to containerize monolithic stateful analytic applications with persistent data.

Benefits include:

- Deploying the management, etcd, and worker nodes on bare-metal eliminate the overhead associated with hypervisors and thus optimizes performance
- Deploying of OCP 4.9 on bare-metal HPE ProLiant DL servers using automation scripts saves significant efforts, resulting in quicker deployment
- The ODF storage can be configured in the internal or external storage mode based on the customer storage requirements

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.

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

TABLE A1. Bill of materials

Part Number	Qty	Description
		Rack
P9K10A	1	HPE 42U 600mmx1200mm G2 Kitted Advanced Shock Rack with Side Panels and Baying
P9K10A 001	1	HPE Factory Express Base Racking Service
		Master Nodes
P19766-B21	3	HPE ProLiant DL360 Gen10 8SFF NC Configure-to-order Server
P19766-B21 0D1	3	Factory Integrated
P19766-B21 ABA	3	HPE DL360 Gen10 8SFF ModX CTO
P02592-L21	3	Intel Xeon-Gold 5218 (2.3GHz/16-core/125W) FIO Processor Kit for HPE ProLiant DL360 Gen10
P02592-B21	3	Intel Xeon-Gold 5218 (2.3GHz/16-core/125W) Processor Kit for HPE ProLiant DL360 Gen10
P02592-B21 0D1	3	Factory Integrated
P00918-B21	12	HPE 8GB (1x8GB) Single Rank x8 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00918-B21 0D1	12	Factory Integrated
P18434-B21	6	HPE 960GB SATA 6G Mixed Use SFF (2.5in) SC 3yr Wty Multi-Vendor SSD
P18434-B21 0D1	6	Factory Integrated
P26966-B21	3	Pensando Distributed Services Platform DSC-25 Enterprise 10/25Gb 2-port SFP28 Card
P26966-B21 0D1	3	Factory Integrated

Part Number	Qty	Description
P02377-B21	3	HPE Smart Storage Hybrid Capacitor with 145mm Cable Kit
P02377-B21 0D1	3	Factory Integrated
804331-B21	3	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	3	Factory Integrated
879482-B21	3	HPE InfiniBand FDR/Ethernet 40/50Gb 2-port 547FLR-QSFP Adapter
879482-B21 0D1	3	Factory Integrated
865414-B21	6	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
865414-B21 0D1	6	Factory Integrated
BD505A	3	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	3	Factory Integrated
734811-B21	3	HPE 1U Cable Management Arm for Rail Kit
734811-B21 0D1	3	Factory Integrated
867998-B21	3	HPE 1U Gen10 Bezel Kit
867998-B21 0D1	3	Factory Integrated
874543-B21	3	HPE 1U Gen10 SFF Easy Install Rail Kit
874543-B21 0D1	3	Factory Integrated
P35876-B21	3	HPE CE Mark Removal FIO Enablement Kit
		Worker Nodes
868703-B21	3	HPE ProLiant DL380 Gen10 8SFF Configure-to-order server
868703-B21 0D1	3	Factory Integrated
868703-B21 ABA	3	HPE DL380 Gen10 8SFF CTO server
P02510-L21	3	Intel Xeon-Gold 6242 (2.8GHz/16-core/150W) FIO Processor Kit for HPE ProLiant DL380 Gen10
P02510-B21	3	Intel Xeon-Gold 6242 (2.8GHz/16-core/150W) Processor Kit for HPE ProLiant DL380 Gen10
P02510-B21 0D1	3	Factory Integrated
P00922-B21	36	HPE 16GB (1x16GB) Dual Rank x8 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00922-B21 0D1	36	Factory Integrated
P18434-B21	9	HPE 960GB SATA 6G Mixed Use SFF (2.5in) SC 3yr Wty Multi-Vendor SSD
P18434-B21 0D1	9	Factory Integrated
P01366-B21	3	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	3	Factory Integrated
804331-B21	3	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	3	Factory Integrated
879482-B21	3	HPE InfiniBand FDR/Ethernet 40/50Gb 2-port 547FLR-QSFP Adapter
879482-B21 0D1	3	Factory Integrated
867810-B21	3	HPE DL38X Gen10 High-Performance Temperature Fan Kit
867810-B21 0D1	3	Factory Integrated
830272-B21	6	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
830272-B21 0D1	6	Factory Integrated
BD505A	3	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	3	Factory Integrated
P8B31A	3	HPE OneView w/o iLO including 3yr 24x7 Support 1-server FIO LTU
733664-B21	3	HPE 2U Cable Management Arm for Easy Install Rail Kit

Part Number	Qty	Description
733664-B21 0D1	3	Factory Integrated
867809-B21	3	HPE Gen10 2U Bezel Kit
867809-B21 0D1	3	Factory Integrated
733660-B21	3	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	3	Factory Integrated
		Ceph Nodes
868703-B21	6	HPE ProLiant DL380 Gen10 8SFF Configure-to-order server
868703-B21 0D1	6	Factory Integrated
868703-B21 ABA	6	HPE DL380 Gen10 8SFF CTO server
P02510-L21	6	Intel Xeon-Gold 6242 (2.8GHz/16-core/150W) FIO Processor Kit for HPE ProLiant DL380 Gen10
P02510-B21	6	Intel Xeon-Gold 6242 (2.8GHz/16-core/150W) Processor Kit for HPE ProLiant DL380 Gen10
P02510-B21 0D1	6	Factory Integrated
P00922-B21	72	HPE 16GB (1x16GB) Dual Rank x8 DDR4-2933 CAS-21-21-21 Registered Smart Memory Kit
P00922-B21 0D1	72	Factory Integrated
P18434-B21	24	HPE 960GB SATA 6G Mixed Use SFF (2.5in) SC 3yr Wty Multi-Vendor SSD
P18434-B21 0D1	24	Factory Integrated
P01366-B21	6	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	6	Factory Integrated
804331-B21	6	HPE Smart Array P408i-a SR Gen10 (8 Internal Lanes/2GB Cache) 12G SAS Modular Controller
804331-B21 0D1	6	Factory Integrated
879482-B21	6	HPE InfiniBand FDR/Ethernet 40/50Gb 2-port 547FLR-QSFP Adapter
879482-B21 0D1	6	Factory Integrated
867810-B21	6	HPE DL38X Gen10 High High-Performance Temperature Fan Kit
867810-B21 0D1	6	Factory Integrated
830272-B21	12	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
830272-B21 0D1	12	Factory Integrated
BD505A	6	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	6	Factory Integrated
P8B31A	6	HPE OneView w/o iLO including 3yr 24x7 Support 1-server FIO LTU
733664-B21	6	HPE 2U Cable Management Arm for Easy Install Rail Kit
733664-B21 0D1	6	Factory Integrated
867809-B21	6	HPE Gen10 2U Bezel Kit
867809-B21 0D1	6	Factory Integrated
733660-B21	6	HPE 2U Small Form Factor Easy Install Rail Kit
733660-B21 0D1	6	Factory Integrated
		Networking
R9F67A	2	Aruba 8325-32C Power to Port Airflow 6 Fans 2 Power Supply Units Bundle for HPE
R9F67A 0D1	2	Factory Integrated
R9F67A ABA	2	Aruba 8325-32C 32-port 100G QSFP+/QSFP28 Back-to-Front 6 Fans and 2 Power Supply Bundle US en
R9F77A	4	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable for HPE
R9F77A B01	4	Factory Integrated
845416-B21	12	HPE 100Gb QSFP28 to 4x25Gb SFP28 3m Direct Attach Copper Cable

Part Number	Qty	Description
845416-B21 0D1	12	Factory Integrated
R9F59A	2	Aruba 4-post Rack Kit for HPE
R9F59A 0D1	2	Factory Integrated
R9F57A	2	Aruba 1U Universal 4-post Rack Mount Kit for HPE
R9F57A 0D1	2	Factory Integrated
R9F60A	2	Aruba Universal 4-post Duct Kit for HPE
R9F60A 0D1	2	Factory Integrated
R9F63A	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE
R9F63A B2B	1	Aruba 6300M 48G Pwr2Prt 2F 1PS Bdl PDU
R9F63A 0D1	1	Factory Integrated
R9F86A	12	Aruba 1G SFP LC SX 500m OM2 MMF Transceiver for HPE
R9F86A 0D1	12	Factory Integrated
R9F91A	2	Aruba 25G SFP28 to SFP28 0.65m Direct Attach Copper Cable for HPE
R9F91A B01	2	Aruba 25G SFP28 to SFP28 0.65m Direct Attach Copper Cable
R9F61A	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE
R9F61A B2B	1	Aruba X371 12VDC 250W 100-240VAC Power-to-Port Power Supply PDU
R9F61A 0D1	1	Factory Integrated
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack include with Complete System Stabilizer Kit
H8A01A3	1	HPE 3Y Foundation Care NBD Exchange
H8A01A3 X7P	2	HPE Aruba 8325-32 Support
H8A01A3 Z59	1	HPE Aruba 6300M 48 Support
HA113A1	1	HPE Installation SVC
HA113A1 5BY	1	HPE Rack and Rack Options Install SVC
HA114A1	1	HPE Installation and Startup Service
HA114A1 5A0	3	HPE Startup Entry 300 Series OS SVC
HA114A1 5A6	9	HPE Startup 300 Series OS SVC
R6A07AAE	3	Pensando Distributed Services Platform Enterprise 3-year Subscription 24x7 Support E-RTU
H8A01A3	1	HPE 3Y Foundation Care NBD Exchange
H8A01A3 X7P	2	HPE Aruba 8325-32 Support
H8A01A3 Z59	1	HPE Aruba 6300M 48 Support
Q9Y41AAE	1	HPE Network Orchestrator E-LTU
JL271A	6	HPE X240 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable
HU4A6A5	1	HPE 5Y Tech Care Essential Service
HU4A6A5 R2M	12	HPE iLO Advanced Non Blade Support
HU4A6A5 WAH	9	HPE DL38x Gen10 Support
HU4A6A5 SVP	9	HPE One View w/o Ilo Support
HU4A6A5 WAG	3	HPE DL360 Gen10 Support
HK771A1	5	HPE Training Credits for Linux SVC
		Optional Persistent Volume
Q8H41A	1	HPE Nimble Storage AF40 All-Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8C17B	1	HPE Nimble Storage 2x10GbE 4-port FIO Adapter Kit

Part Number	Qty	Description
Q8G27B	1	HPE Tier 1 Storage OS Default FIO Software
Q8H47A	1	HPE Nimble Storage AF40 All-Flash Array R2 11.52TB (24x480GB) FIO Flash Bundle
Q8J18A	2	HPE Nimble Storage NEMA 5-15P to C13 125V 10Amp 1.8m US FIO Power Cord
R3P91A	1	HPE Tier 1 Storage Array Standard Tracking
HU4A9A5	1	HPE 5Y Tech Care Essential Exchange Service
HU4A9A5 ZG0	1	HPE NS 2x10GbE 4p Adptr Supp
HU4A9A5 ZFF	1	HPE NS AF40 All Flash Base Array Supp
HU4A9A5 ZGL	1	HPE NS AF40/60/80 11TB Flash Upgr Supp
HA114A1	1	HPE Installation and Startup Service
HA114A1 5MR	1	HPE Tier 1 Storage Array Startup SVC
		Red Hat Enterprise Linux CoreOS (RHCOS) Server
J8J36AAE	12	Red Hat Enterprise Linux CoreOS (RHCOS) Server 2 Sockets 1 Guest 1 Year Subscription 24x7 Support
		Red Hat Enterprise Linux CoreOS (RHCOS) for Virtual data centers
G3J22AAE	12	Red Hat Enterprise Linux CoreOS (RHCOS) for Virtual Datacenters 2 Sockets 1 Year Subscription 24x7 Support

NOTE
For ODF – External approach, six (6) more HPE ProLiant DL360 Gen10 and/or HPE ProLiant DL380 Gen10 servers will be required. This is an optional requirement.

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 DL380 Gen10 server, https://h20195.www2.hpe.com/v2/default.aspx?cc=us&lc=en&oid=1010026818

HPE ProLiant DL360 Gen10 server, http://www.hpe.com/servers/dl360-gen10

Red Hat OpenShift Container Platform, https://access.redhat.com/documentation/en-us/openshift_container_platform/4.9/

Red Hat OpenShift Container Storage, https://access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.9

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

LEARN MORE AT

hpe.com/info/ra

Hewlett Packard Enterprise

a50005795enw, version 1.0, February 2022