Hewlett Packard Enterprise

HPE Reference Configuration for Red Hat OpenShift Container Platform 4.12 on HPE ProLiant DL360 & DL380 Gen11 servers

Rapid deployment on HPE ProLiant DL360 & DL380 Gen11 servers using Red Hat OpenShift Container Platform 4.12

CONTENTS

Executive summary 3
Introduction 4
Solution overview 4
Design objectives 4
Physical configuration 4
Solution components 5
Hardware Components 5
Software components 9
Capacity and sizing 10
Best practices and configuration guidance for the solution 10
Network Overview 10
Storage 11
Storage subscriptions 13
Deployment overview 13
Deploying the Red Hat OpenShift Container Platform 4.12 cluster using the User Provisioned Infrastructure 13
Red Hat OpenShift Data Foundation 13
HPE Alletra CSI 14
Accelerating deployment 14
Securing and monitoring Red Hat OpenShift with Sysdig SaaS 15
Business continuity with Data Protection for Red Hat OpenShift Container Platform 4.12 16
Business continuity with Disaster Recovery strategies for Red Hat OpenShift Container Platform 4.12 17
Red Hat OpenShift Data Foundation– Internal Mode 20
Summary 21
Appendix A: Bill of materials 21
Red Hat OpenShift Data Foundation subscription 21
Resources and additional links 26

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 Red Hat ® OpenShift® Container Platform 4.12 on the HPE ProLiant DL360 & DL380 Gen11 servers to accelerate container application delivery.

This Reference Configuration provides architectural guidance for deploying Red Hat OpenShift Container Platform 4.12 and HPE ProLiant DL360 Gen11 servers for Compute. The compute requirements can easily be scaled by adding more HPE ProLiant DL360 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 DL360 Gen11 servers.
2. To install and configure the Red Hat OpenShift Container Platform 4.x.
3. Validate the Red Hat OpenShift Container Platform installation.

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 DL360 & DL380 Gen11 servers for compute, storage, networking, and Red Hat OpenShift Container Platform 4.12.

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 DL360 Gen11 servers. This document assumes that the reader is familiar with HPE ProLiant DL360 Gen11 servers, Red Hat OpenShift Container Platform 4.12, 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 Container Platform 4.12 on HPE ProLiant DL360 Gen11 and HPE ProLiant DL380 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/4.12-INTEL-LTI/.

INTRODUCTION

This Reference Configuration provides guidance for installing Red Hat OpenShift Container Platform 4.12 (the solution), on HPE ProLiant DL360 & DL380 Gen11 servers. The solution consists of six (6) HPE ProLiant DL360 Gen11 servers: three (3) HPE ProLiant DL360 Gen11 servers used for the Red Hat Enterprise Linux (RHEL) KVM-based Head Nodes and three (3) HPE ProLiant DL360 Gen11 servers used for the solution worker nodes (out of which one node is used as a temporary Red Hat OpenShift Container Platform bootstrap node). HPE ProLiant DL380 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.

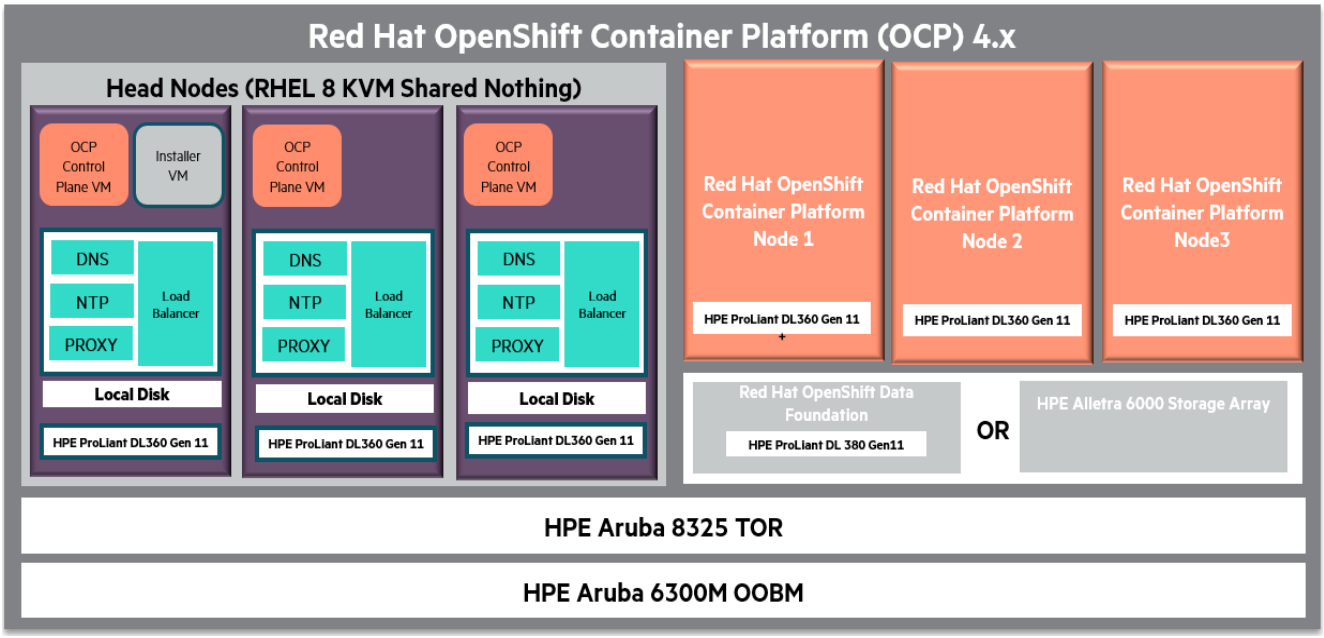

FIGURE 1. High-level architecture

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) 8.6 on the HPE ProLiant DL360 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 DL360 & DL380 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 RHEL8.6 and KVM on three (3) HPE ProLiant DL360 Gen11 servers. Three (3) HPE ProLiant DL360 Gen11 servers are deployed as solution worker nodes on bare metal. The temporary bootstrap node is deployed on one of the worker nodes and later configured as a worker node. HPE ProLiant DL380 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 DL360 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 DL360 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.

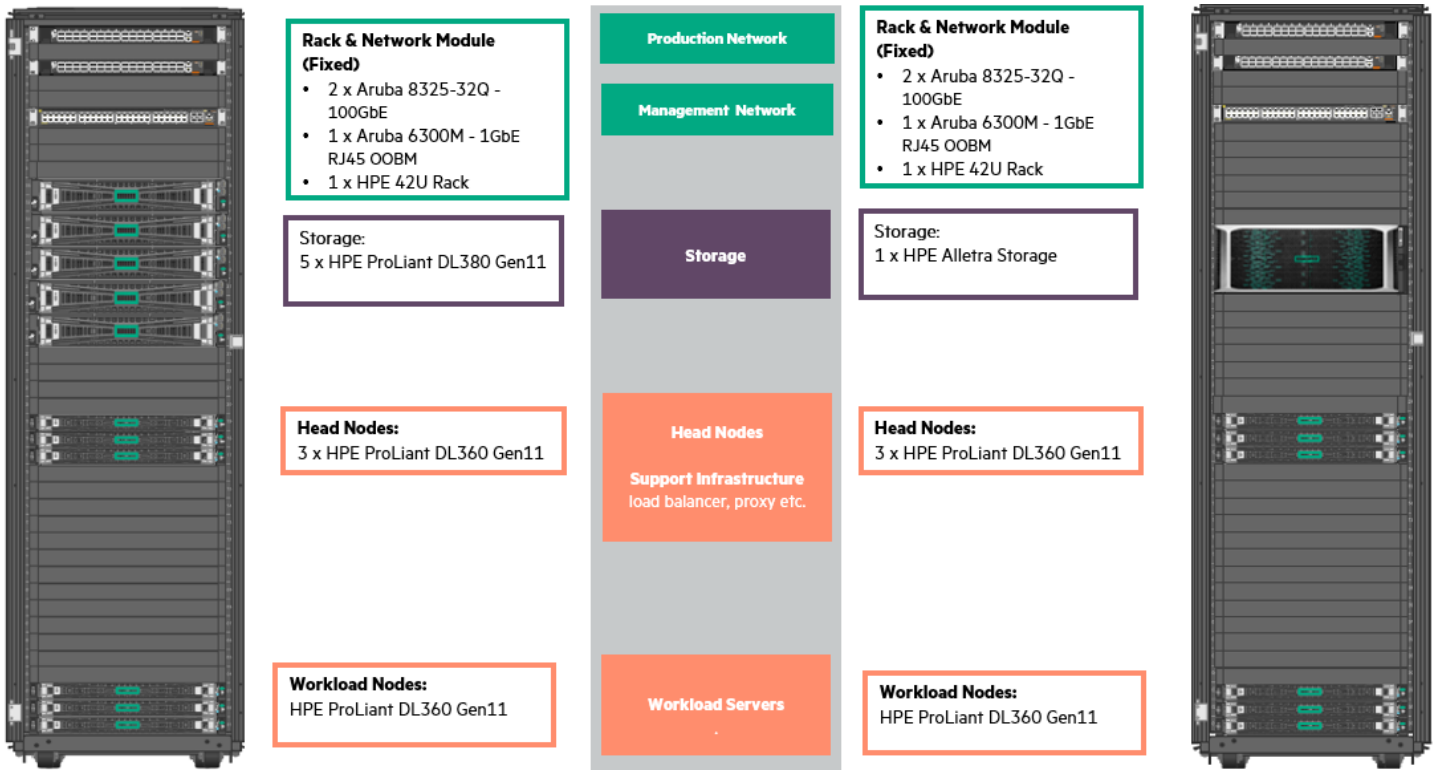

FIGURE 2. Solution components

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

Additional HPE ProLiant DL360 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 DL360 Gen11 servers	3	Provides capacity for head nodes with OpenShift master and bootstrap KVM VMS, HAProxy, DNS, Proxy
HPE ProLiant DL360 Gen11 servers	3	Provide OpenShift worker nodes based on KVM VM's
HPE Alletra 6070	1	External iSCSI storage for Persistent Volumes

Component	Qty	Description
HPE ProLiant DL380 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 8.6	64	512 GB	OS Disk: 2x 1.6 TB Data Disk: 2x 1.6 TB
Workload Nodes	RHEL 8.6	64	512 GB	OS Disk: 2x 1.6 TB Data Disk: 2x 1.6 TB
Red Hat Openshift Data Foundation Nodes	RHCOS	64	768 GB	OS Disk: 2x 1.2 TB Data Disk: 6x 3.2 TB

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 or Master nodes	RHCOS	8	16 GB	250GB
Worker nodes	RHCOS/RHEL	8	16 GB	250GB

NOTE

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

HPE ProLiant DL360 Gen11 server

The HPE ProLiant DL360 Gen11 server is redefining price/performance with the new math for virtualized compute. Powered by the latest Intel Xeon-Gold 6438Y Series Processors, the HPE ProLiant DL360 Gen11 server offers greater processing power, memory speeds up to 4800 MT/s. and data transfer rates with PCIe Gen5 capabilities. This 2P, 1U server has been designed with flexibility while delivering a high core count and large memory footprint. Choose this purpose-built platform for virtualization, High-Performance Compute, and memory-centric workloads. Designed for supreme versatility and resiliency while backed by a comprehensive warranty, the HPE ProLiant DL360 Gen11 server is ideal for IT infrastructure, either physical, virtual, or containerized.

Figure 3 shows the HPE ProLiant DL360 Gen11 server.

FIGURE 3. HPE ProLiant DL360 Gen11 server

Table 4 lists the hardware configuration used in this solution.

TABLE 4 Hardware configuration in each of the HPE ProLiant DL360 Gen11 servers

Component	Description
Processor	2x Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor
Memory	8x 32GB (1x32GB) Dual Rank x8 DDR5-4800
Network	InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter
Smart Array Controller	MR216i-o Gen11 x16 Lanes
Disks	4x HPE 1.6TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static Multi-Vendor SSD

HPE ProLiant DL380 Gen11 server

The HPE ProLiant DL380 Gen11 server is redefining price/performance with the new math for virtualized compute. Powered by the latest Intel Xeon-Gold 6430 Series Processors, the HPE ProLiant DL380 Gen11 servers offer greater processing power, memory speeds up to 4800 MT/s. and data transfer rates with PCIe Gen5 capabilities. This 2P, 2U server has been designed with flexibility while delivering a high core count and large memory footprint. Choose this purpose-built platform for virtualization, High-Performance Compute, and memory-centric workloads. One of the HPE ProLiant DL380 servers is initially used as the Red Hat OpenShift Data foundation storage cluster.

Figure 4 shows the HPE ProLiant DL380 Gen11 server.

FIGURE 4. HPE ProLiant DL380 Gen11 server

Table 5 lists the hardware configuration in each of the HPE ProLiant DL380 Gen11 servers used in this solution.

TABLE 5. Hardware configuration in each of the HPE ProLiant DL380 Gen11 servers

Component	Description
Processor	2 x Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor
Memory	5 x HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800
Network	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter
Array Controller	HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller
Disks	2XHPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD 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 DL360 & DL380 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 4.12 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 4.12

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 4.12 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 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	4.12	Red Hat OpenShift control plane nodes on KVM virtual machines and bare metal worker nodes
Red Hat Enterprise Linux	8.6	KVM head node and Red Hat worker nodes bare metal
Red Hat Enterprise Linux	8.6	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

CAPACITY AND SIZING

Sizing for a Red Hat OpenShift Container Platform 4.12 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 4.12 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 6 lists the various networks used for this solution. All the cluster nodes and iPXE servers are connected to the same network.

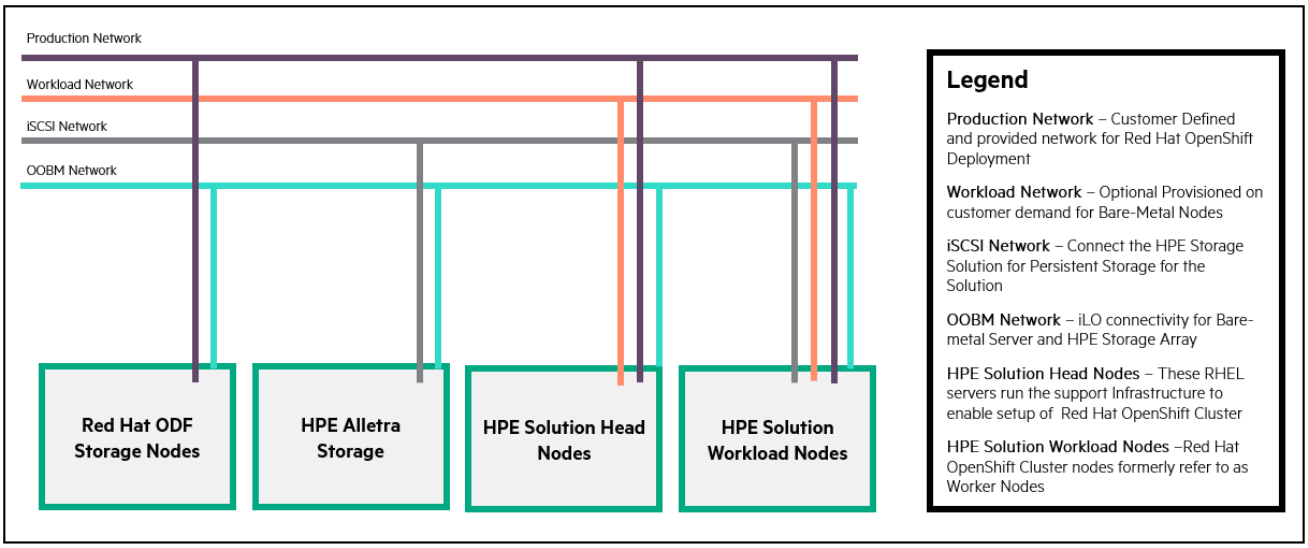

FIGURE 8. Networks for solution

NOTE
For OOBM High Availability, 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 DL360 and HPE ProLiant DL380 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	iSCSI Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated
HPE ProLiant DL380 Gen11 (Internal storage Red Hat OpenShift Data Foundation)	Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated

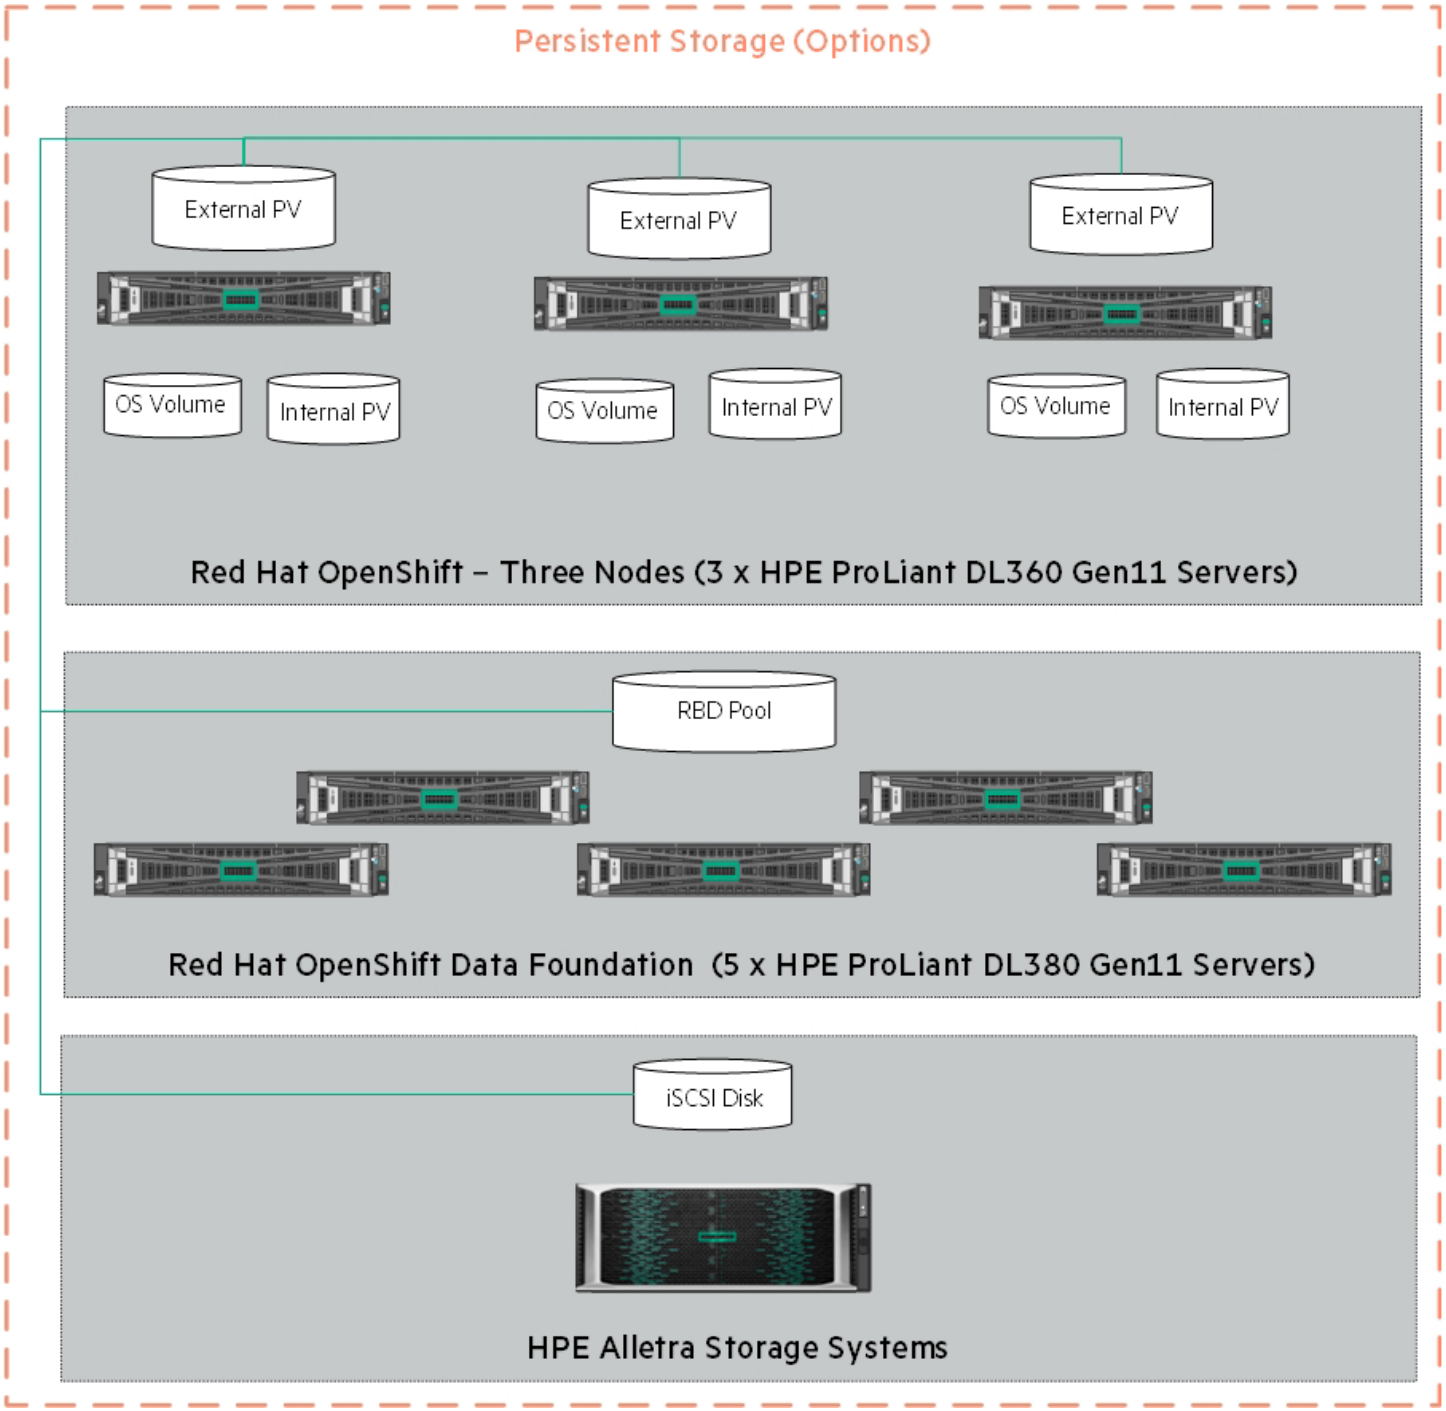

FIGURE 9. Logical storage layout for Red Hat OpenShift persistent volume options

NOTE

- HPE ProLiant DL380 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.

Storage subscriptions

In this solution, we have used Red Hat OpenShift Data Foundation internal and require a license for Red Hat OpenShift Container Platform Plus (OPP).

DEPLOYMENT OVERVIEW

This section explains in detail the deployment of Red Hat OpenShift Container Platform 4.12 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 4.12 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 OpenShift cluster. For detailed installation and configuration information, see the Deployment guide.

Figure 10 explains the Red Hat OpenShift Container Platform (RHOCP) 4.12 deployment process.

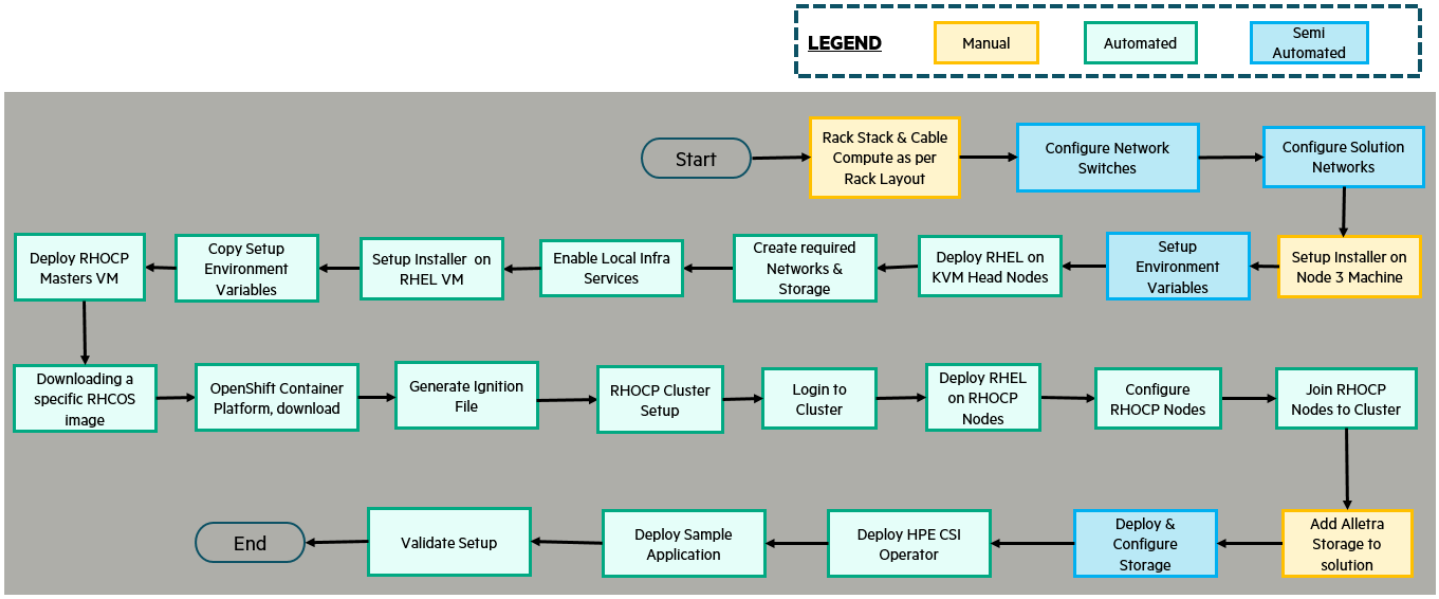

FIGURE 10. Deployment process for Red Hat OpenShift Container Platform 4.12 cluster using the UPI

NOTE

The load balancer described in this document is HAProxy.

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/4.12-INTEL-LTI/.

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/4.12-INTEL-LTI/.

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 4.12. 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 11 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.

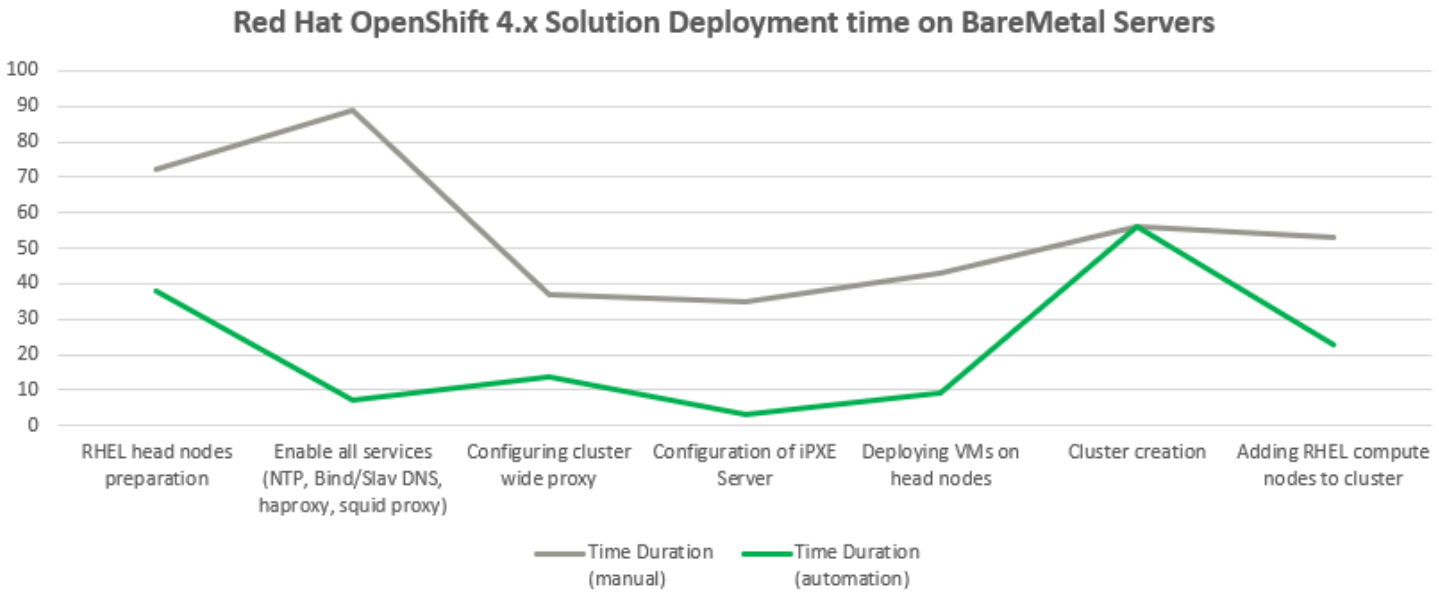

FIGURE 11. Red Hat OpenShift 4.x solution deployment manual and automation timelines on bare metal

Figure 12 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.

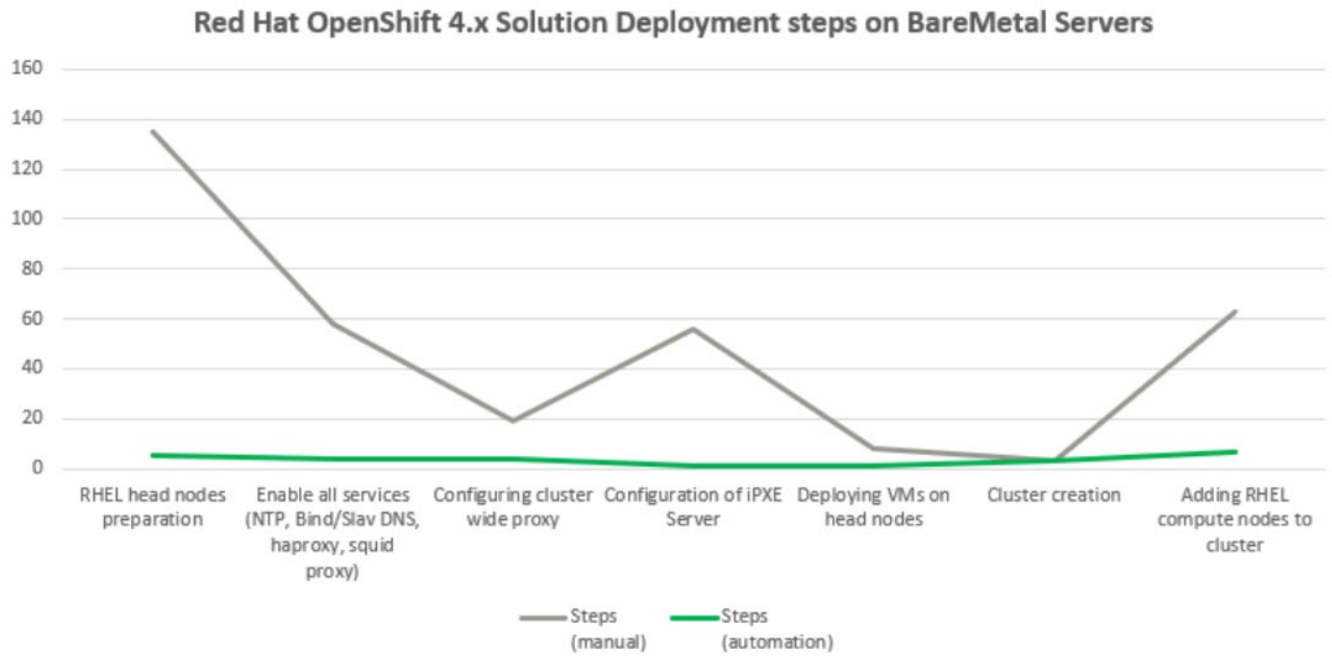

FIGURE 12. Red Hat OpenShift 4.x solution deployment manual and automation steps on bare metal

Securing and monitoring Red Hat 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 their choice of subscription. For security and monitoring of Red Hat OpenShift Containers, it is required to install the Sysdig Agent on the Red Hat OpenShift cluster. This means Sysdig Agents, which are lightweight entities, will be installed on each node in the Red Hat OpenShift cluster. These agents run as daemons 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 a Red Hat OpenShift cluster.

The key benefits are as follows:

- Faster incident resolution using Sysdig Monitor for Red Hat OpenShift cluster.
- Simplified compliance for the entire solution.
- Service-based access control for container security and monitoring.
- Less time is 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 Red Hat 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 on the Sysdig Agent deployment in the Red Hat OpenShift setup, see the HPE solutions for Red Hat OpenShift Platform documentation.

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 in the HPE solutions for Red Hat OpenShift Container Platform documentation.

BUSINESS CONTINUITY WITH DATA PROTECTION FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.12

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.

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 4.12, Velero uses a controller model where it monitors custom resources and takes actions.

Figure 13 shows the overview of Velero backup and restore with solution.

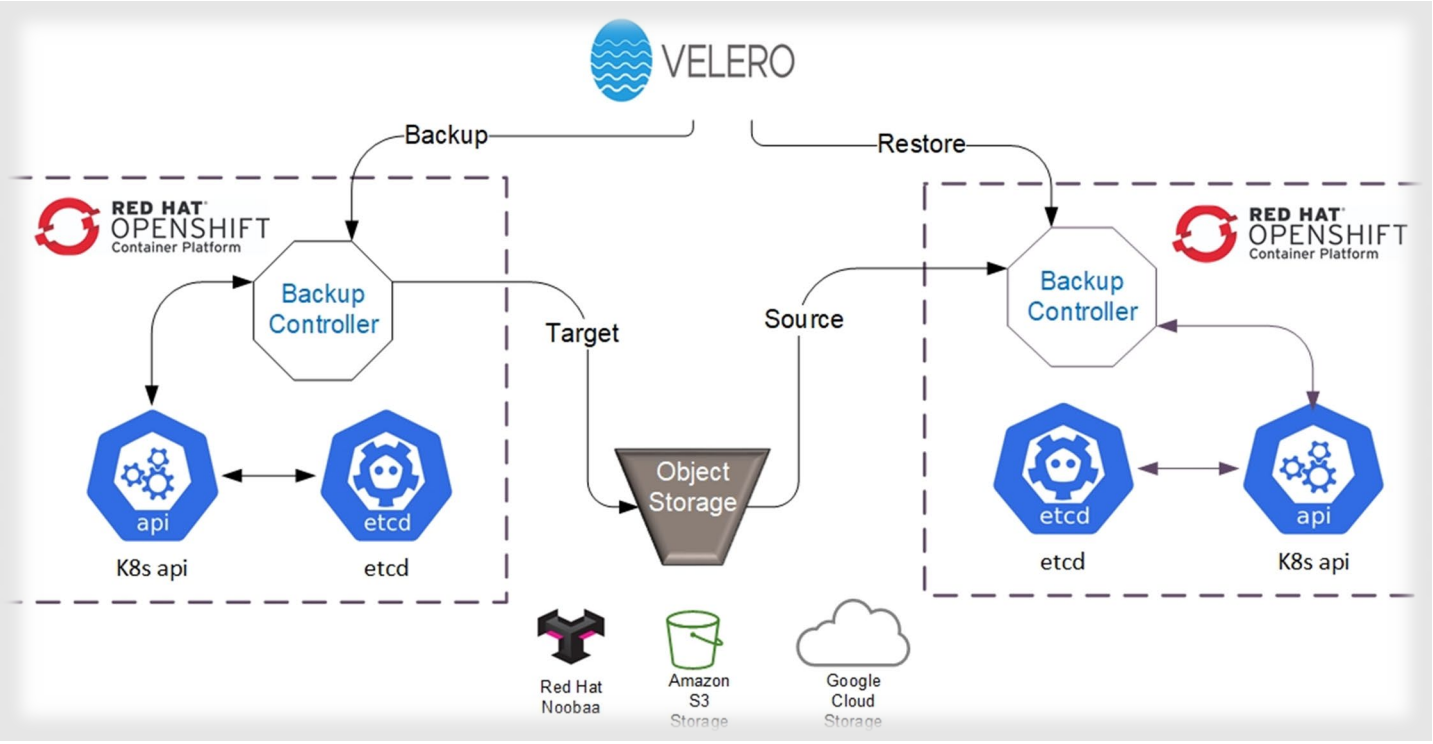

FIGURE 13. Velero backup and restore with Red Hat OpenShift Container Platform 4.12

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 14 shows the overview of Velero for the Red Hat OpenShift setup.

FIGURE 14. 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/4.12-INTEL-LTI/

BUSINESS CONTINUITY WITH DISASTER RECOVERY STRATEGIES FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.12

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 4.12 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 15 shows the comparison of the Red Hat OpenShift disaster recovery strategies using RTO and RPO objectives.

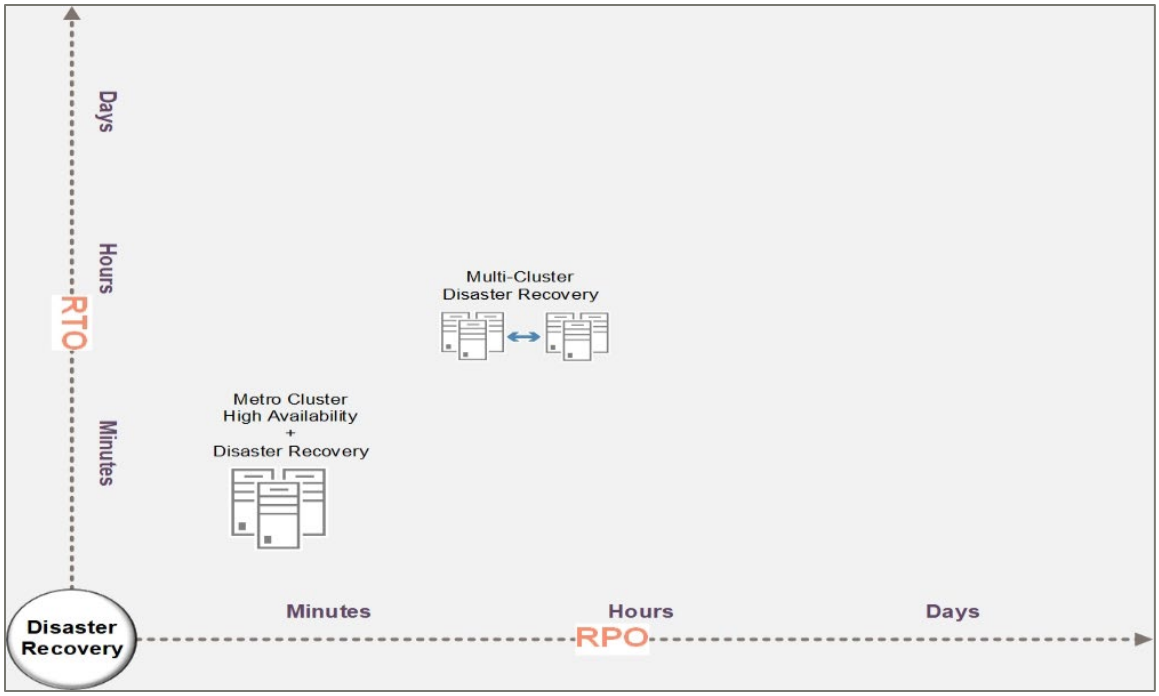

FIGURE 15. Red Hat OpenShift Disaster recovery strategies comparison using RTO and RPO objectives.

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.
   - DR sites or Availability Zones (AZs) connected by MAN or campus networks.
   - AZs are mapped to a fault domain (Heating,Ventilation and Air Conditioning, Power grids, etc.).
   - An odd number of AZ or fault domains are required for the cluster quorum.
   - 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 16 shows an overview of the Red Hat OpenShift (OCP) Metro Cluster design.

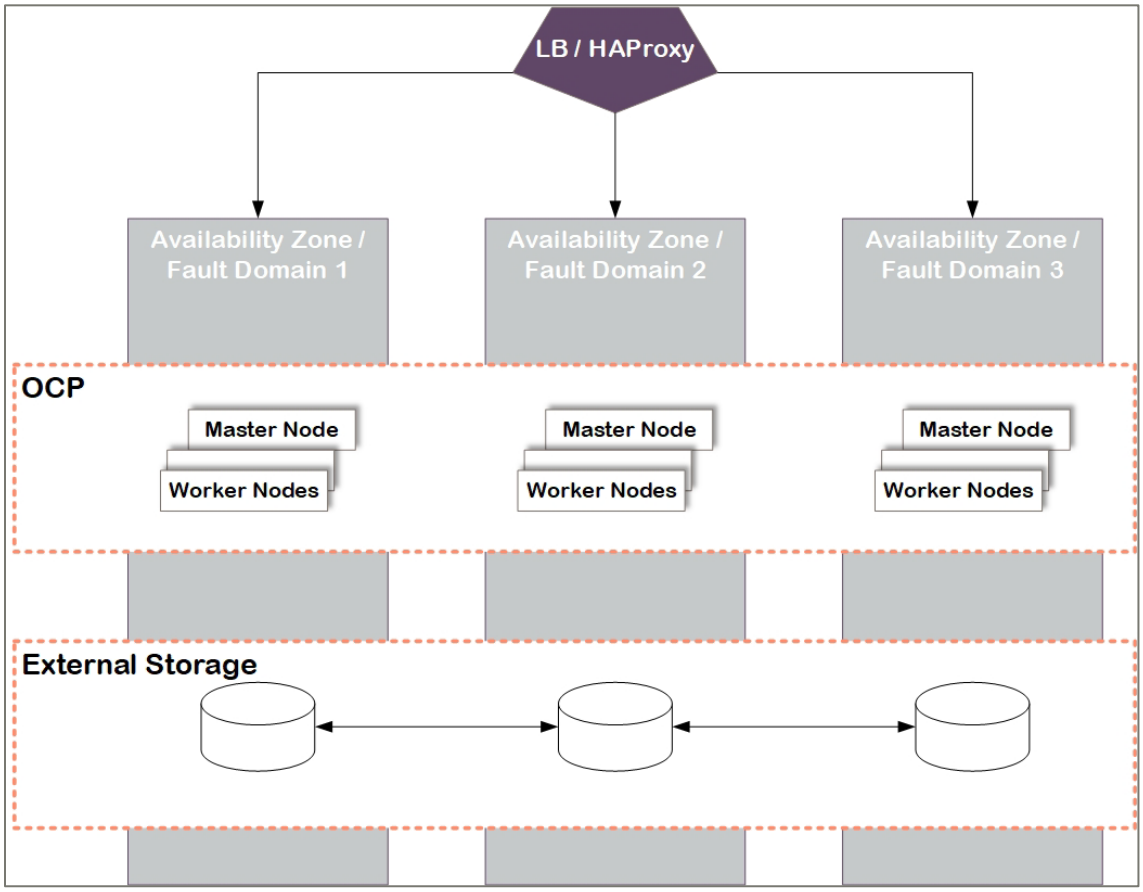

FIGURE 16. Red Hat OpenShift Metro Cluster design overview

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

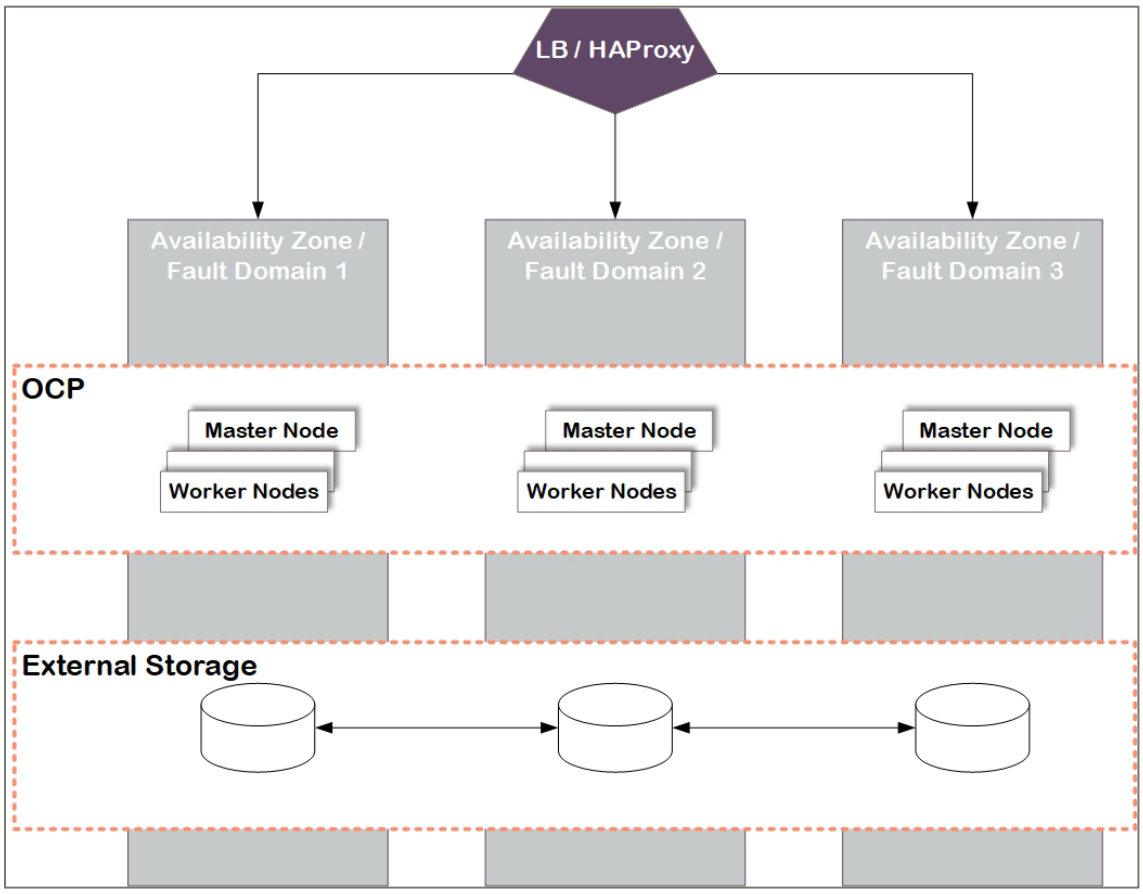

FIGURE 17. Red Hat OpenShift Multi-cluster disaster recovery approach

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 10GB for monitoring purposes and block storage of 500GB/2TB for OSD (Object Storage Daemon) 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.

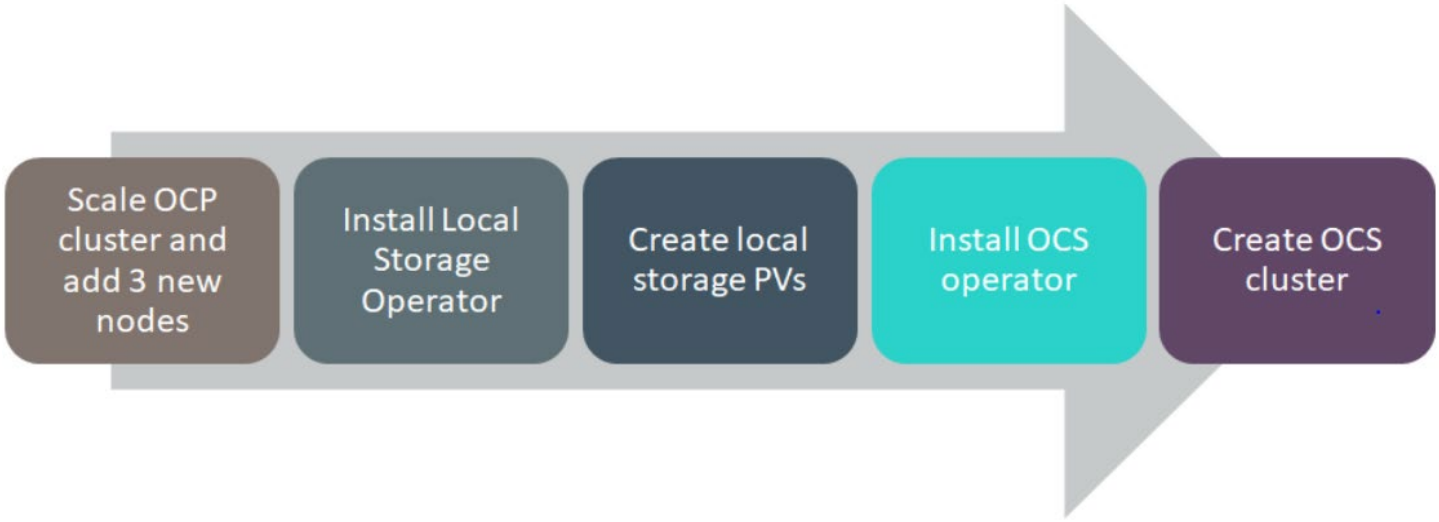

FIGURE 18. Red Hat OpenShift data foundation flow diagram

For more information, view the deployment guide at https://hewlettpackard.github.io/hpe-solutions-openshift/4.12-INTEL-LTI/.

SUMMARY

In this solution installation and configuration of the Red Hat OpenShift Container Platform Version 4.12 on the HPE ProLiant DL360 configured as master and worker nodes with storage like HPE Alletra Storage Arrays and DL380 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 virtualization. IT teams can manage multiple Kubernetes clusters with multitenant container isolation and data access, for any workload from edge to core and 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, etc., and worker nodes on bare metal eliminates the overhead associated with hypervisors and thus optimizes performance.
- Deploying Red Hat OpenShift Container Platform 4.12 on HPE ProLiant DL360 & DL380 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 Data Foundation subscription

Red Hat OpenShift Data Foundation requires Red Hat OpenShift Platform Plus license (includes Red Hat OpenShift Data Foundation Essentials) and is socket based subscription.

In this solution we have deployed Red Hat OpenShift Data Foundation on five (5) worker nodes. Each of these physical nodes have either one (1) or two (2) CPU sockets.

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
P9K40A	1	HPE 42U 600mmx1200mm G2 Enterprise Shock Rack
P9K40A 001	1	HPE Factory Express Base Racking Service
H4F42A1	15	HPE Factory Express Complex Unit of SVC
HA454A1-000	1	HPE FE Solution Package 4 SVC
P52499-B21	6	HPE ProLiant DL360 Gen11 8SFF NC Configure-to-order Server
P52499-B21 0D1	6	Factory Integrated
P52499-B21 ABA	6	HPE DL360 G11 SFF CTO Server
HA454A1-001	6	HPE FE Proliant Svr Pkg 4 SVC
P43328-B21	96	HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P43328-B21 0D1	96	Factory Integrated
P48895-B21	6	HPE ProLiant DL360 Gen11 8SFF x1 U.3 Tri-Mode Backplane Kit
P48895-B21 0D1	6	Factory Integrated
P47838-B21	24	HPE 1.6TB NVMe Gen4 Mainstream Performance Mixed Use SFF BC U.3 Static Multi Vendor SSD
P47838-B21 0D1	24	Factory Integrated
P47789-B21	6	HPE MR216i-o Gen11 x16 Lanes without Cache OCP SPDM Storage Controller
P47789-B21 0D1	6	Factory Integrated
P31348-B21	6	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter
P31348-B21 0D1	6	Factory Integrated
P48908-B21	6	HPE ProLiant DL3X0 Gen11 1U High Performance Fan Kit
P48908-B21 0D1	6	Factory Integrated
P38997-B21	12	HPE 1600W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
P38997-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
P48828-B21	6	HPE ProLiant DL3XX Gen11 OCP2 x16 Enablement Kit
P48828-B21 0D1	6	Factory Integrated
P52416-B21	6	HPE ProLiant DL360 Gen11 OROC Tri-Mode Cable Kit
P52416-B21 0D1	6	Factory Integrated
P49615-B21	12	Intel Xeon-Gold 6438Y+ 2.0GHz 32-core 205W Processor for HPE
P49615-B21 0D1	12	Factory Integrated
P48922-B21	6	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21 0D1	6	Factory Integrated
P08040-B21	6	HPE iLO Common Password FIO Setting
P48905-B21	12	HPE ProLiant DL3XX Gen11 High Performance Heat Sink Kit
P48905-B21 0D1	12	Factory Integrated
P52341-B21	6	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit

Component	Qty	Description
P52341-B21 0D1	6	Factory Integrated
P52534-B21	5	HPE ProLiant DL380 Gen11 8SFF NC Configure-to-order Server
P52534-B21 0D1	5	Factory Integrated
P52534-B21 ABA	5	HPE DL380 Gen11 8SFF NC CTO Svr
P49614-B21	10	Intel Xeon-Gold 6430 2.1GHz 32-core 270W Processor for HPE
P49614-B21 0D1	10	Factory Integrated
P43328-B21	60	HPE 32GB (1x32GB) Dual Rank x8 DDR5-4800 CAS-40-39-39 EC8 Registered Smart Memory Kit
P43328-B21 0D1	60	Factory Integrated
P48810-B21	5	HPE DL380 G11 2SFF U.3 Prim/Sec Cage Kit
P48810-B21 0D1	5	Factory Integrated
P28586-B21	10	HPE 1.2TB SAS 12G Mission Critical 10K SFF BC 3-year Warranty Multi Vendor HDD
P28586-B21 0D1	10	Factory Integrated
P48813-B21	5	HPE DL380 Gen11 2U 8SFF x1 TM Kit
P48813-B21 0D1	5	Factory Integrated
P50230-B21	40	HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD
P50230-B21 0D1	40	Factory Integrated
P48802-B21	5	HPE DL380 G11 2U x8/x16/x8 Sec Riser Kit
P48802-B21 0D1	5	Factory Integrated
P47777-B21	5	HPE MR416i-p Gen11 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller
P47777-B21 0D1	5	Factory Integrated
P01366-B21	5	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	5	Factory Integrated
P48918-B21	5	HPE ProLiant DL360 Gen11 Storage Controller Enablement Cable Kit
P48918-B21 0D1	5	Factory Integrated
P31348-B21	5	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 OCP3 MCX653436A-HDAI Adapter
P31348-B21 0D1	5	Factory Integrated
P38995-B21	10	HPE 800W Flex Slot Platinum Hot Plug Low Halogen Power Supply Kit
P38995-B21 0D1	10	Factory Integrated
P48828-B21	5	HPE ProLiant DL3XX Gen11 OCP2 x16 Enablement Kit
P48828-B21 0D1	5	Factory Integrated
P48820-B21	5	HPE DL380/DL560 G11 2U High Perf Fan Kit
P48820-B21 0D1	5	Factory Integrated
P48818-B21	10	HPE DL380/DL560 G11 High Perf 2U HS Kit
P48818-B21 0D1	10	Factory Integrated
P52341-B21	5	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit
P52341-B21 0D1	5	Factory Integrated
R9F63A	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE
R9F63A 0D1	1	Factory Integrated
R9F63A B2B	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE PDU
HA454A1-021	1	HPE FE Strg and Ntwking Pkg 4 SVC
R9G06A	1	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9G06A B01	1	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE

Component	Qty	Description
R9F61A	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE
R9F61A B2B	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE PDU
R9F61A 0D1	1	Factory Integrated
R9F57A	1	Aruba 1U Universal 4-post Rack Mount Kit for HPE
R9F57A 0D1	1	Factory Integrated
R9F59A	2	Aruba 4-post Rack Kit for HPE
R9F59A 0D1	2	Factory Integrated
R9F67A	2	Aruba 8325-32C Power to Port Airflow 6 Fans 2 Power Supply Units Bundle for HPE
R9F67A 0D1	2	Factory Integrated
R9F67A B2B	2	Aruba 8325-32C Power to Port Airflow 6 Fans 2 Power Supply Units Bundle for HPE PDU
HA454A1-021	2	HPE FE Strg and Ntwking Pkg 4 SVC
R9F77A	2	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable for HPE
R9F77A B01	2	Aruba 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable for HPE
C7535A	24	HPE RJ45 to RJ45 Cat5e Black M/M 7.6ft 1-pack Data Cable
C7535A 0D1	24	Factory Integrated
HA113A1	1	HPE Installation SVC
HA113A1 5MW	3	HPE Aruba 6xxxN8xxx Install Swt SVC
H6J85A	1	HPE Rack Hardware Kit
H6J85A 0D1	1	Factory Integrated
P9L11A	1	HPE G2 Rack Grounding Kit
P9L11A 0D1	1	Factory Integrated
P9L12A	1	HPE G2 Rack Baying Kit
P9L12A B01	1	HPE G2 Rack Baying Kit
P9L16A	1	HPE G2 Rack 42U 1200mm Side Panel Kit
P9L16A 0D1	1	Factory Integrated
P9S23A	2	HPE G2 Metered/Switched 3Ph 17.3kVA/60309 4-wire 48A/208V Out (36) C13 (12) C19/Vertical NA/JP PDU
P9S23A 0D1	2	Factory Integrated
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack include with Complete System Stabilizer Kit
P8B26AAE	6	HPE OneView w/o iLO including 3yr 24x7 Support Flexible Quantity E-LTU
R9G32AAE	1	Aruba Fabric Composer Device Management Service Tier 3 Switch 3 year Subscription E-STU for HPE
R9G27AAE	2	Aruba Fabric Composer Device Management Service Tier 4 Switch 3 year Subscription E-STU for HPE
HL263A1	1	HPE Applied Network CI SVC 40H onsite
HU4A6A5	1	HPE 5Y Tech Care Essential Service
HU4A6A5 R2M	6	HPE iLO Advanced Non Bltade Support
HU4A6A5 SVP	6	HPE One View w/o iLO Support
HU4A6A500DJ	6	HPE DL360 Gen11 Support
H7J34A3	1	HPE 3Y Foundation Care 24x7 SVC
H7J34A3 ZSG	1	HPE Aruba 6300M 48 SW Support
H7J34A3 ZND	2	HPE Aruba 8325-32 SW Support
HA124A1	1	HPE Technical Installation Startup SVC
HA124A1 5MR	1	HPE Tier 1 Storage Array Startup SVC

Component	Qty	Description
HV9V8A1	1	HPE Digital Learner - SMB Ed 1Yr Sub SVC
H33XSA1	100	HPE Edu 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
R7S86A	1	HPE Alletra 6000 184TB (24x7.68TB) 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
S1E76AAE	1	HPE Alletra 6000 Software and Support SaaS
S1E76AAE CTF	603	3yr Subscription
S1F94AAE	1	HPE GreenLake for Alletra Storage Tracking
HU4A6A3	1	HPE 3Y Tech Care Essential Service
HU4A6A3 ZUN	2	HPE Alletra 6000 2x10/25GbE 2p Kit Supp
HU4A6A3 ZUH	1	HPE Alletra 6070 Base Array Supp
HU4A6A3 ZV8	1	HPE Alletra 6000 AF184TB 7.68 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 DL360 Gen11, hpe.com/servers

Red Hat OpenShift Container Platform, https://access.redhat.com/documentation/en-us/openshift_container_platform/4.12/

Red Hat OpenShift Container Storage, https://access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.12

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

a50009527enw, October 2024

Hewlett Packard Enterprise