Hewlett Packard Enterprise

Reference Configuration

HPE REFERENCE CONFIGURATION FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.10 ON HPE PROLIANT DL365 GEN10 PLUS SERVERS

Rapid deployment on HPE ProLiant DL365 Gen10 Plus Servers using Red Hat OpenShift Container Platform 4.10

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 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 4.10 (RHOCP) on the HPE ProLiant DL365 Gen10 Plus Servers to accelerate container application delivery.

This Reference Configuration provides architectural guidance for deploying Red Hat OpenShift Container Platform 4.10 and HPE ProLiant DL365 Gen10 Plus Servers for Compute. The compute requirements can easily be scaled by adding more HPE ProLiant DL365 Gen10 Plus 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 DL365 Gen10 Plus 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 the cost-effective yet reliable solution by leveraging the benefits of HPE ProLiant DL365 Gen10 Plus Servers for compute, storage, networking, and Red Hat OpenShift Container Platform 4.10.

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 DL365 Gen10 Plus Servers. This document assumes that the reader is familiar with HPE ProLiant DL365 Gen10 Plus Servers, Red Hat OpenShift Container Platform 4, 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 4.10 on HPE ProLiant DL365 Gen10 Plus 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/4.10-LTI/

INTRODUCTION

This Reference Configuration provides guidance for installing Red Hat OpenShift Container Platform 4.10 on HPE ProLiant DL365 Gen10 Plus Servers. The solution consists of six (6) HPE ProLiant DL365 Gen10 Plus Servers: three (3) HPE ProLiant DL365 Gen10 Plus Servers used for the Red Hat Enterprise Linux (RHEL) KVM-based Head Nodes and three (3) HPE ProLiant DL365 Gen10 Plus Servers used for the RHOCP worker nodes (out of which one node is used as a temporary RHOCP bootstrap node).

The persistent storage for this solution is provided by HPE Alletra 6000 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.

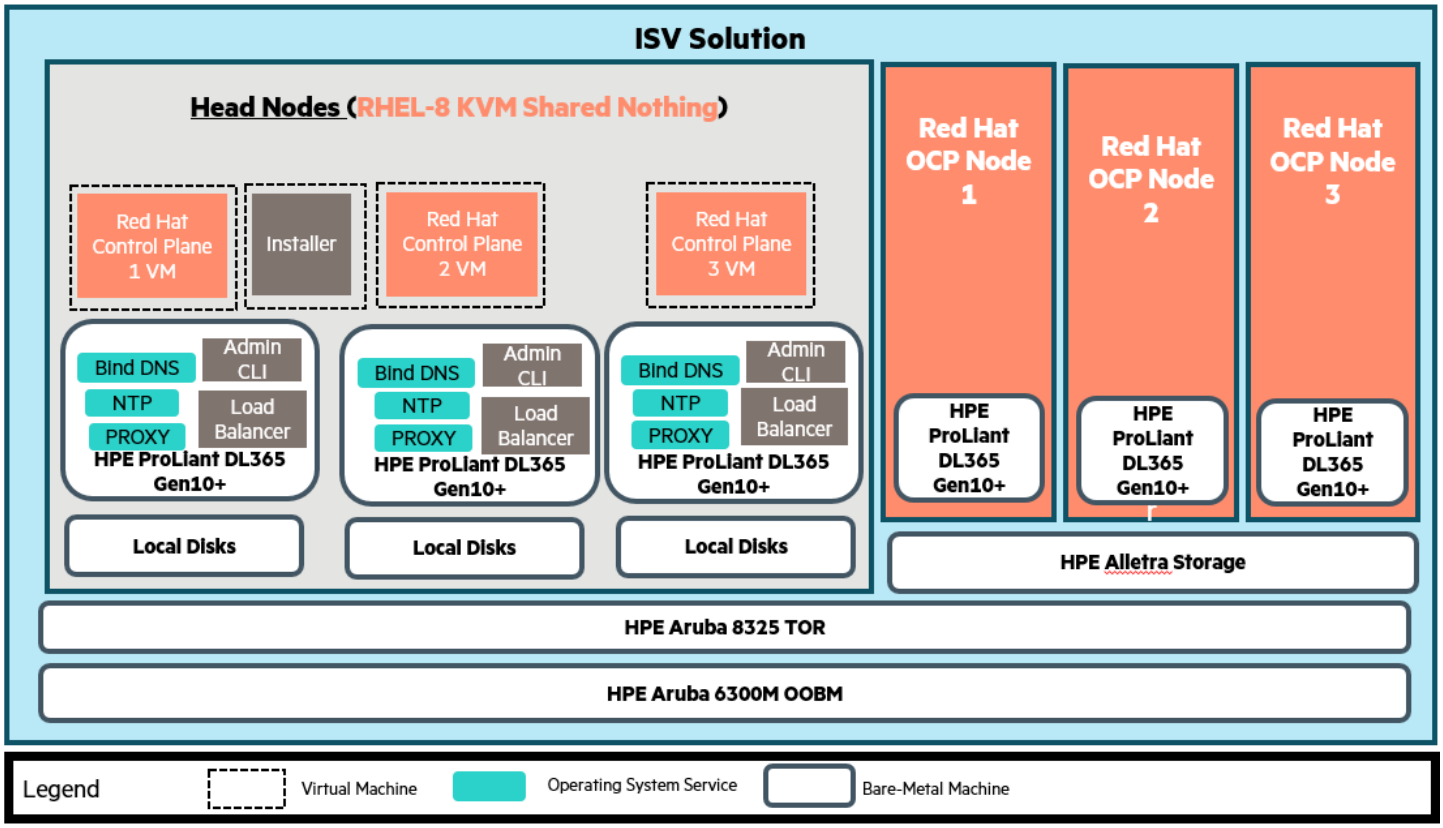

FIGURE 1. High-level architecture

This solution uses the Red Hat OpenShift User Provisioned Infrastructure method of installation to install Red Hat Linux CoreOS (RHCOS) and Red Hat Enterprise Linux® (RHEL) 8.6 on the HPE ProLiant DL365 Gen10 Plus 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 DL365 Gen10 Plus 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 RHOCP Control Plane nodes are deployed as KVM virtual machines running Red Hat CoreOS. These virtual machines are running Red Hat Enterprise Linux 8.6 and KVM on three (3) HPE ProLiant DL365 Gen10 Plus v2 Servers. Three (3) HPE ProLiant DL365 Gen10 Plus Servers are deployed as RHOCP worker nodes on bare metal. The temporary bootstrap node is deployed on one of the worker nodes and later configured as a worker node.

The solution uses the internal storage on the HPE ProLiant DL365 Gen10 Plus Servers for both the Operating System and RHOCP 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 OpenShift-installer tool is run to generate ignition files that contain information about the hosts that will be provisioned. The Red Hat Linux CoreOS for the nodes is then booted with the help of iPXE and the ignition files are passed with the OS image during installation. HPE ProLiant DL365 Gen10 Plus Servers use HPE Alletra 6070 via iSCSI to provide persistent container volume for the RHOCP application workload.

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

NOTE
Additional HPE ProLiant DL365 Gen10 Plus Servers can be added to this solution as per the customer's choice of configuration options.

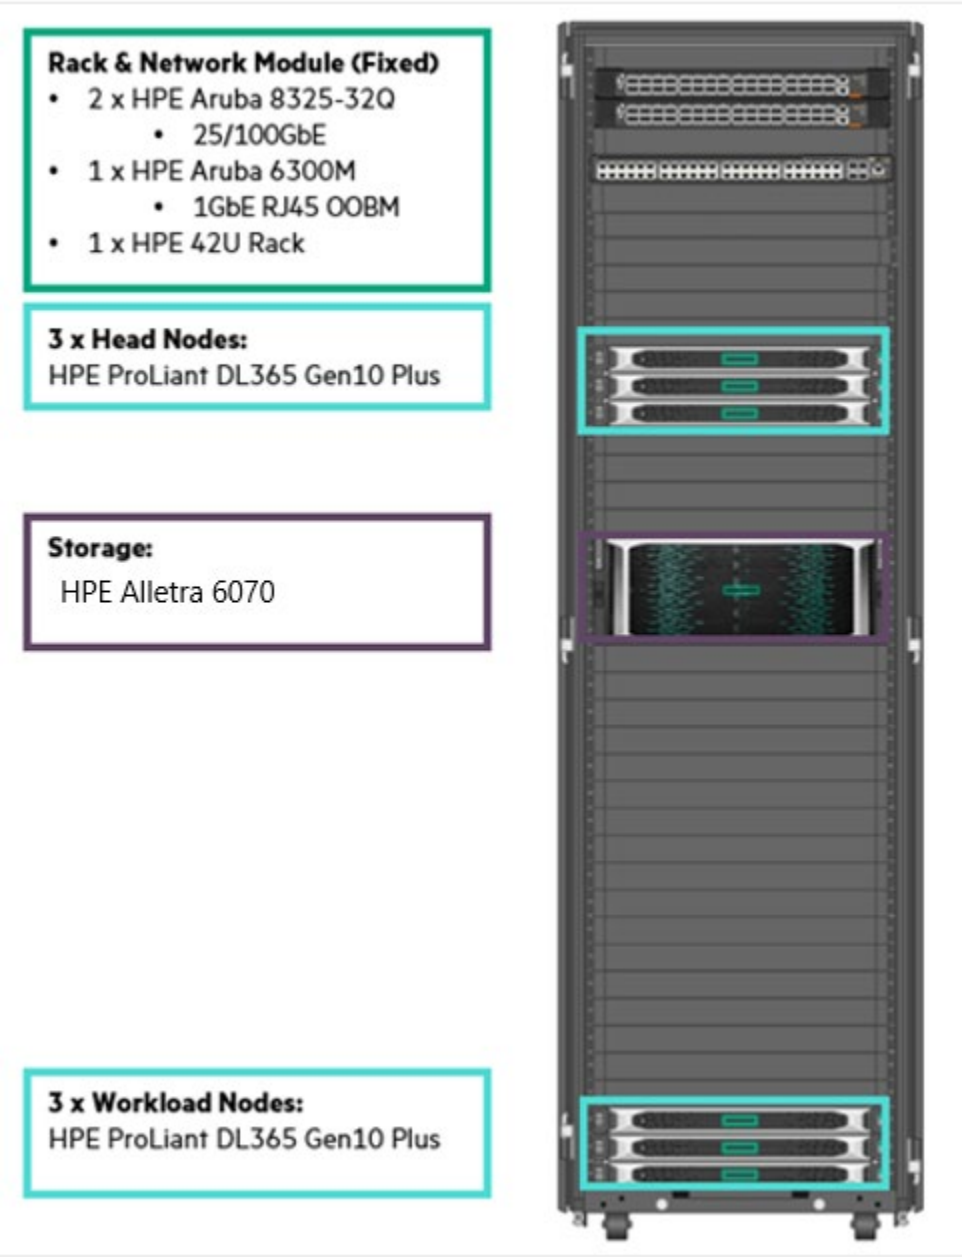

FIGURE 2. Solution components

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

SOLUTION COMPONENTS

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

Hardware

Table 1 shows the minimum resource requirements for the deployment of RHOCP without OpenShift Data Foundation.

TABLE 1. Minimum hardware requirement for RHOCP cluster

Node	Operating System	vCPU	Virtual RAM	Storage
Bootstrap node	RHCOS	4	16 GB	120 GB
RHOCP Control plane nodes	RHCOS	4	16 GB	120GB
Worker nodes	RHEL 8.6	2	8 GB	120GB

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 DL365 Gen10 Plus Servers	3	Provides compute capacity for head nodes
HPE ProLiant DL365 Gen10 Plus Servers	3	OpenShift worker nodes/bootstrap with RHEL8.6
HPE Alletra 6070	1	External iSCSI storage for Persistent Volumes
HPE Aruba 8325 switch	2	A network switch for datacenter network
HPE Aruba 6300M switch	1	A network switch for iLO Management network

Server roles and hardware 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	RAM	Storage
Bootstrap node	RHCOS	4	16 GB	Hard Disk 1 : 300 GB
RHOCP Control plane nodes	RHCOS	8	64 GB	Hard Disk 1 : 300 GB
Compute or Worker nodes	RHEL 8.6	64	512GB	Hard Disk 1 : 2x 800 GB Hard Disk 2 : 2x 1.6 TB
KVM Head Nodes	RHEL 8.6	64	512 GB	Hard Disk 1 : 2x 800 GB Hard Disk 2 : 2x 1.6 TB

NOTE

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

Figure 3 shows the HPE ProLiant DL365 Gen10 Plus Server.

FIGURE 3. HPE ProLiant DL365 Gen10 Plus Server

HPE ProLiant DL365 Gen10 Plus Server

The HPE ProLiant DL365 Gen10 Plus server is redefining price/performance with the new math for virtualized compute. Powered by the latest AMD® EPYC® 7003 Series Processors, the HPE ProLiant DL365 Gen10 Plus Server offers greater processing power, memory speeds up to 3200 MT/s. and data transfer rates with PCIe Gen4 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 DL365 Gen10 Plus Server is ideal for IT infrastructure, either physical, virtual, or containerized.

Table 4 lists the hardware configuration used in this solution.

TABLE 4. Hardware configuration in each of the HPE ProLiant DL365 Gen10 Plus Servers

Component	Description
Processor	2x AMD EPYC 7543 (2.8 GHz/32-core)
Memory	16x HPE 32GB Dual Rank x4 DDR4-3200 Smart Kit
Network	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56
Smart Array Controller	Broadcom MegaRAID MR216i-a x16 Lanes
Disks	• 2x HPE 800GB SAS 24G Mixed Use SFF BC Self-encrypting FIPS PM6 SSD • 2x HPE 1.6TB NVMe Gen4 High-Performance Mixed Use SFF BC Self-encrypting FIPS U.3 CM6 SSD

HPE iLO

HPE Integrated Lights Out (iLO) is embedded in HPE ProLiant DL365 Gen10 Plus platforms and provides server management that enables faster deployment, and simplified lifecycle operations while maintaining end-to-end security, thus increasing productivity.

Aruba 8325-32C BF switch

The 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 4 shows the Aruba 8325-32C BF switch.

FIGURE 4. Aruba 8325-32C BF switch

Aruba CX 6300M OOBM Switch

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 verifies that configurations are flawless and easy to deploy.

A powerful Aruba Gen7 ASICs architecture delivers fast, non-blocking performance, meaning your network is ready for tomorrow's unpredictable demands. 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 5 shows the Aruba 6300M OOBM Switch.

FIGURE 5. Aruba CX 6300M OOBM switch

HPE Alletra 6070

HPE Alletra 6070 powers your data from the edge to the core with the cloud experience for all your apps. For business-critical workloads, HPE Alletra 6070 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 delivers 6x9s1 availability guaranteed. Deliver recovery SLAs with fast, integrated app-aware backup and recovery, on-premises and in the cloud.

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 (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.

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 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 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 4.10 deployment on bare metal infrastructure, you must use RHCOS for all RHOCP control plane nodes, Bootstrap nodes, and RHEL 8.6 for worker nodes.

Table 5 lists the major software used in this solution.

TABLE 5. Software used in this solution

Component	Versions	Usage
Red Hat Enterprise Linux CoreOS	4.10	Red Hat OpenShift Control Plane VMs
Red Hat OpenShift Container Platform	4.10	Red Hat OpenShift Control Plane Nodes on KVM Virtual Machines and BareMetal worker nodes
Red Hat Enterprise Linux	8.6	KVM head node and Red Hat worker nodes BareMetal
Red Hat Enterprise Linux	7.9	Installer Machine required to execute automation scripts

CAPACITY AND SIZING

Sizing for a Red Hat OpenShift Container Platform 4 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 OpenShift Container Platform release. For more information, see RHOCP 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 RHOCP control plane 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 4.10 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.

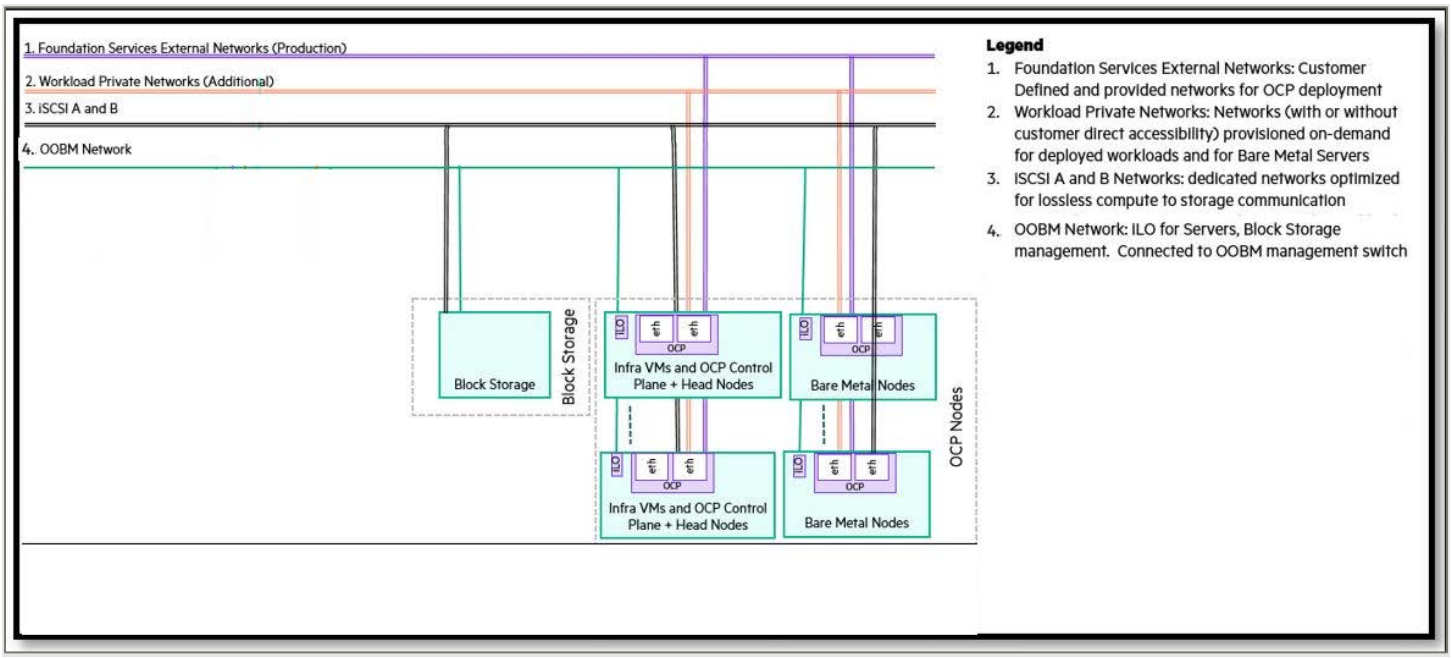

FIGURE 6. Networks for RHOCP solution

NOTE
For OOBM High Availability, use 2x Aruba 6300 switches.

DEPLOYMENT OVERVIEW

This section explains in detail the deployment of RHOCP 4.10 using external storage mode. In the external storage mode, HPE Alletra 6070 is connected via the iSCSI network to the RHOCP worker nodes.

Deploying the RHOCP 4.10 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 RHCOS ignition configuration files. These ignition files include the bootstrap ignition files, RHOCP control plane ignition files, and worker ignition files. The ignition files are used to configure RHCOS on each of the RHOCP control planes and worker nodes in the OpenShift cluster. For detailed installation and configuration information, see the Deployment guide. For detailed information on bare metal installation, see Installing on bare metal.

Figure 7 explains the RHOCP 4.10 deployment process.

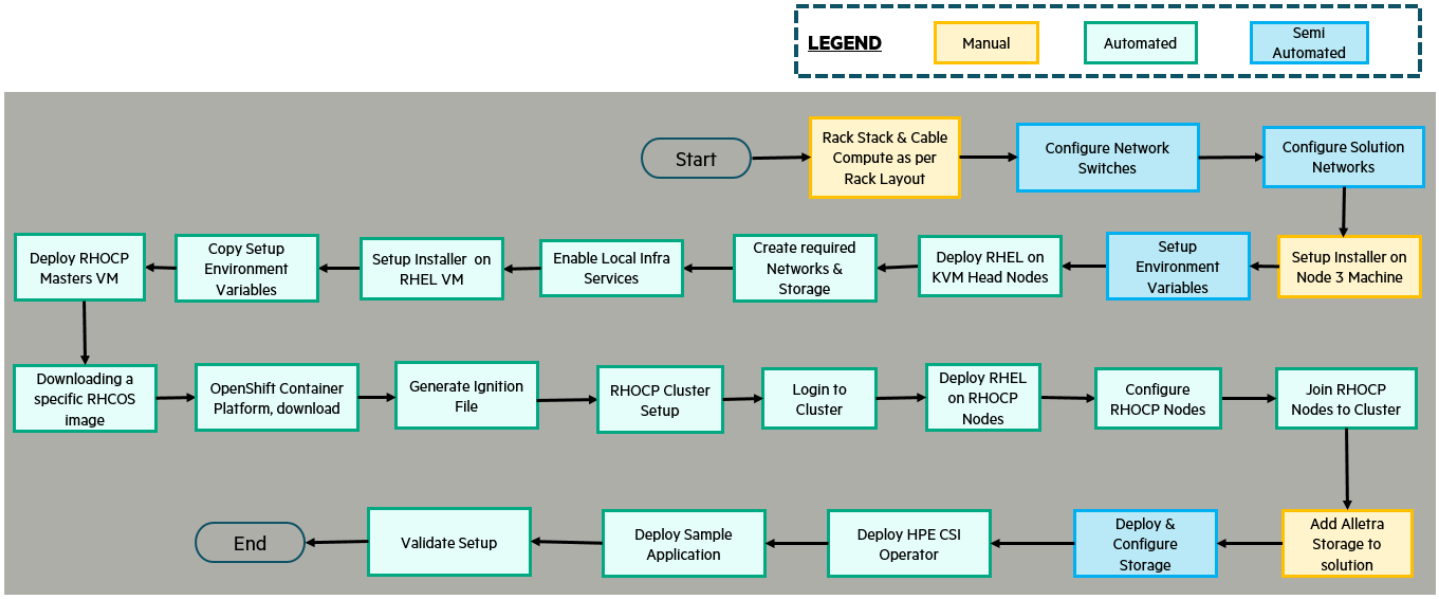

FIGURE 7. Deployment process for RHOCP 4.10 cluster using the UPI

NOTE
The load balancer described in this document is HAProxy.

Physical worker node labeling in OpenShift

We have discussed setting up the RHOCP cluster and then setting up the persistent volume so far. 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. A Red Hat OpenShift cluster can have many nodes. Each node can run multiple pods which, at scale, means that this process is both tedious and error-prone. With Red Hat 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 for 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, the compute in the HPE ProLiant DL365 Gen10 Plus Servers 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 thoroughly evaluated using the Intel TXT-enabled platform. For more information about node labeling configuration, see the HPE solutions for Red Hat OpenShift Platform documentation.

Post-validation deployment

After the successful deployment of RHOCP on the identified nodes and access to the OpenShift console is successful, the following services are deployed from the operator hub for monitoring the logs of the applications deployed on the cluster:

1. Install EFK operators in RHOCP.
2. Install the Elasticsearch and Cluster Logging operators (Fluentd and Kibana).
3. Install Cluster Logging operator.
4. Create a cluster logging Instance.
5. Launch Kibana.

Figure 8 shows the Kibana dashboard.

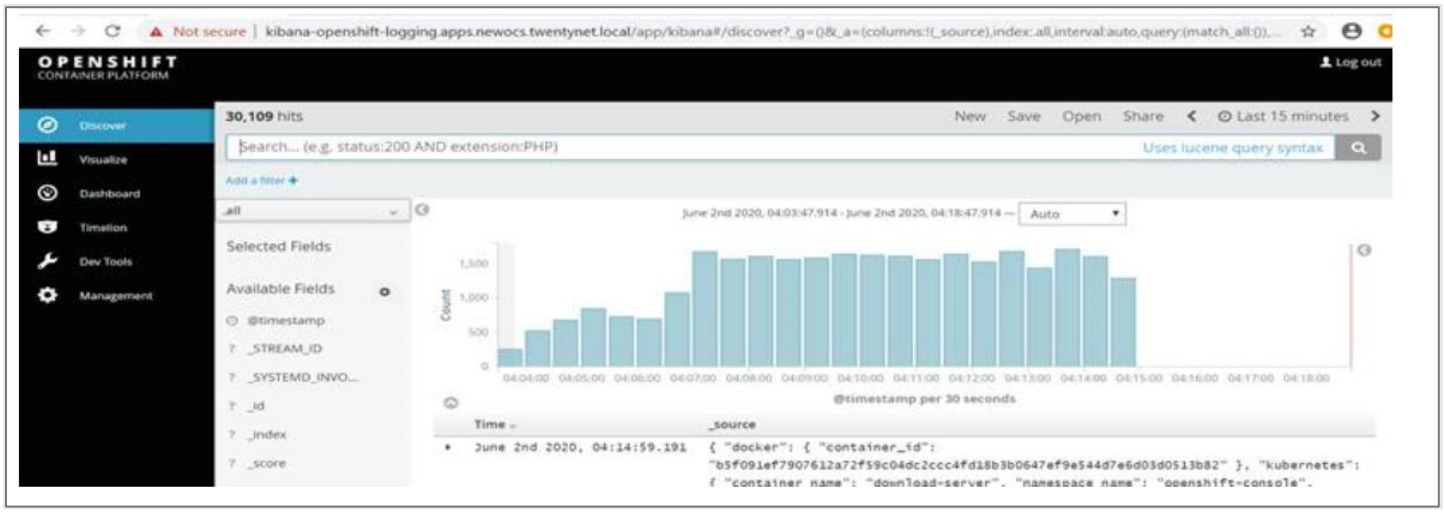

FIGURE 8. View of Kibana dashboard

For more information on the installation and configuration of EFK operators, cluster operators, Kibana, and so on, see the Deployment guide.

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.

Accelerating deployment

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 4.10. Automating the deployment 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 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 9 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.

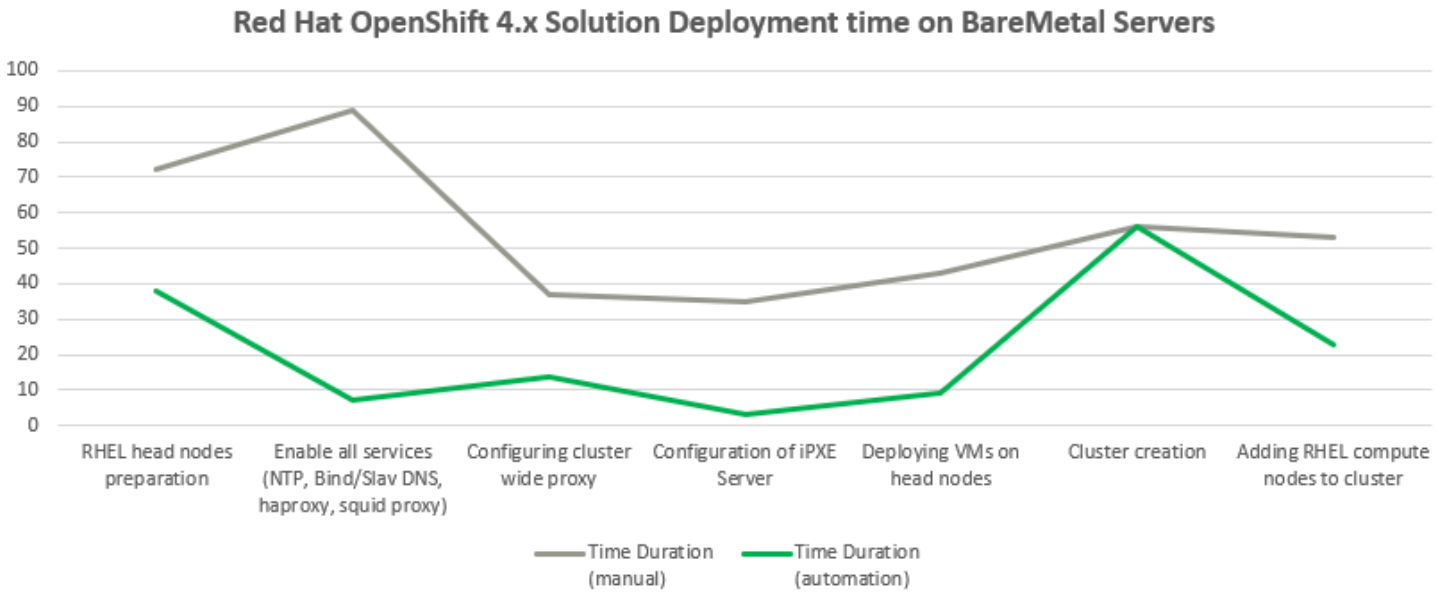

FIGURE 9. Red Hat OpenShift 4.x solution deployment manual and automation timelines on bare metal

Figure 10 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.

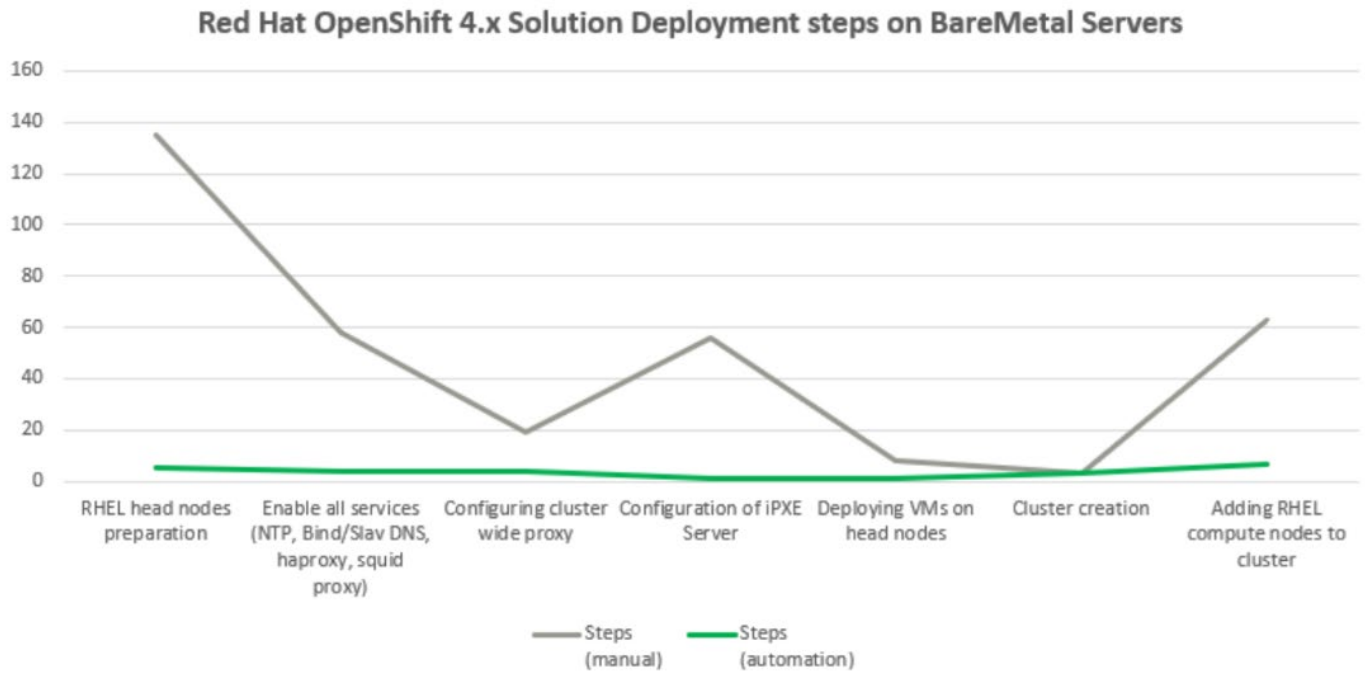

FIGURE 10. Red Hat OpenShift 4.x solution deployment manual and automation steps on bare metal

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 their 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, 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 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 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 4.10

Backup and restore is a day 2 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 4.10, Velero uses a controller model where it monitors custom resources and takes actions.

Figure 11 shows the overview of Velero backup and restore with RHOCP 4.10.

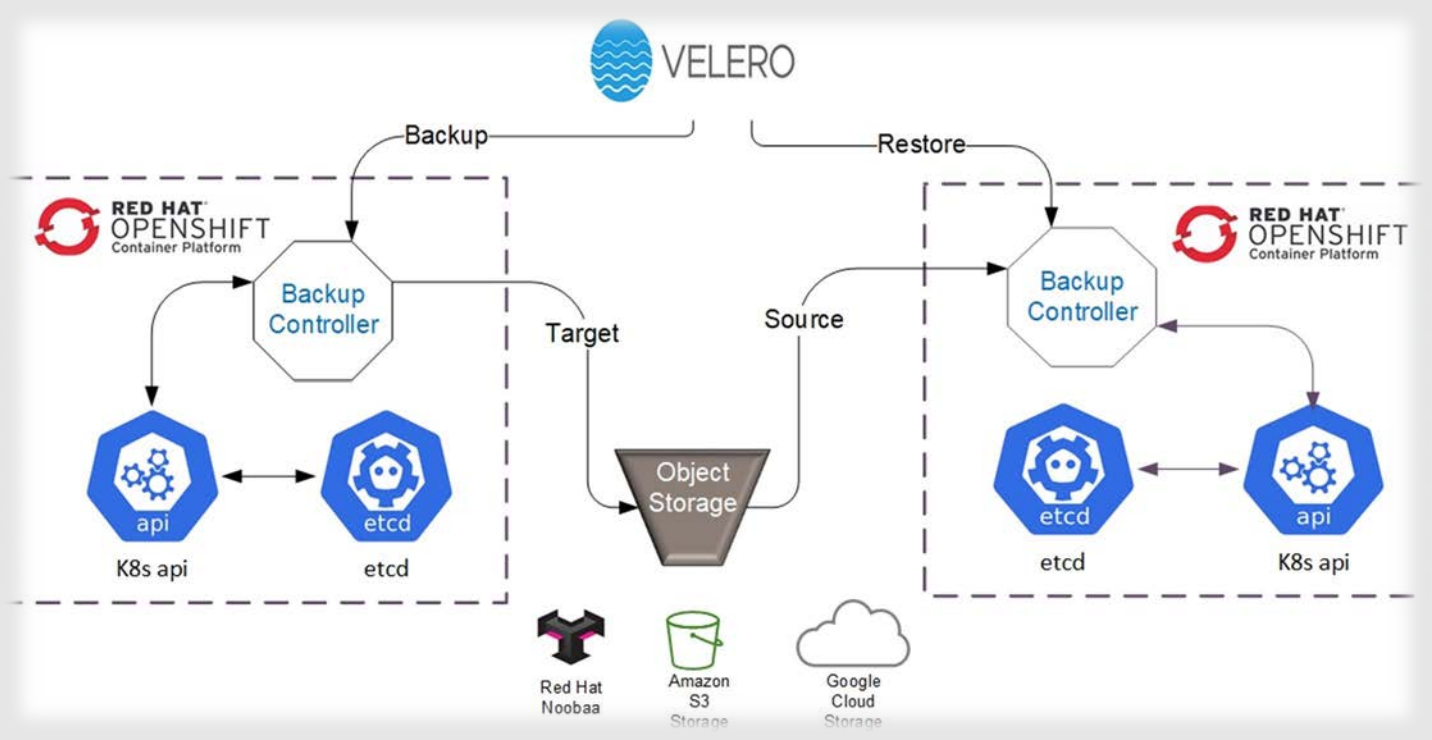

FIGURE 11. Velero backup and restore with RHOCP 4.10

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 OpenShift setup overview

Figure 12 shows the overview of Velero for the OpenShift setup.

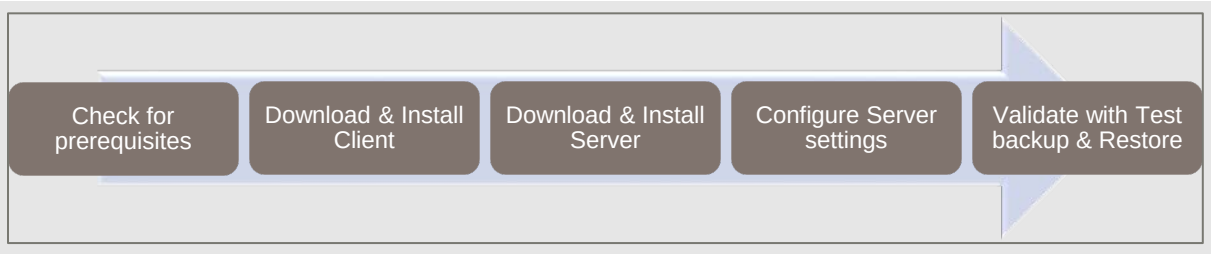

FIGURE 12. Velero for OpenShift setup overview

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/4.10-LTI/

BUSINESS CONTINUITY WITH DISASTER RECOVERY STRATEGIES FOR RED HAT OPENSHIFT CONTAINER PLATFORM 4.10

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 4.10 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 13 shows the comparison of the Red Hat OpenShift disaster recovery strategies using RTO and RPO objectives.

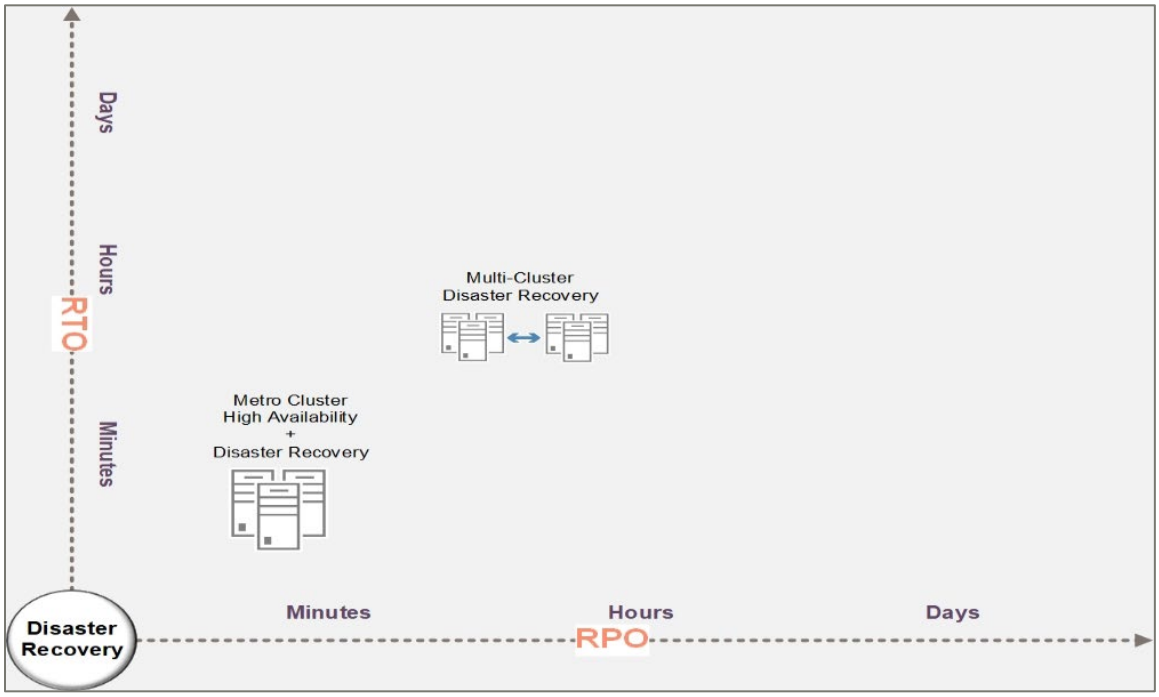

FIGURE 13. Red Hat OpenShift Disaster recovery strategies comparison using RTO and RPO objectives

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 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 HA like automatic recovery along with no data loss data mirroring:

1. Synchronous HA-DR for localized data center failures.
   - DR sites or 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 or fault domains are required for the cluster quorum.
   - Network latency between zones does not typically exceed 5 ms RTT.
2. RHOCP ensures pods and nodes get scheduled across zones during deployment.
3. ODF maintains consistent mirror copies across AZs resulting in less or no data loss.
4. Stretched RHOCP 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 14 shows an overview of the Red Hat OpenShift Metro Cluster design.

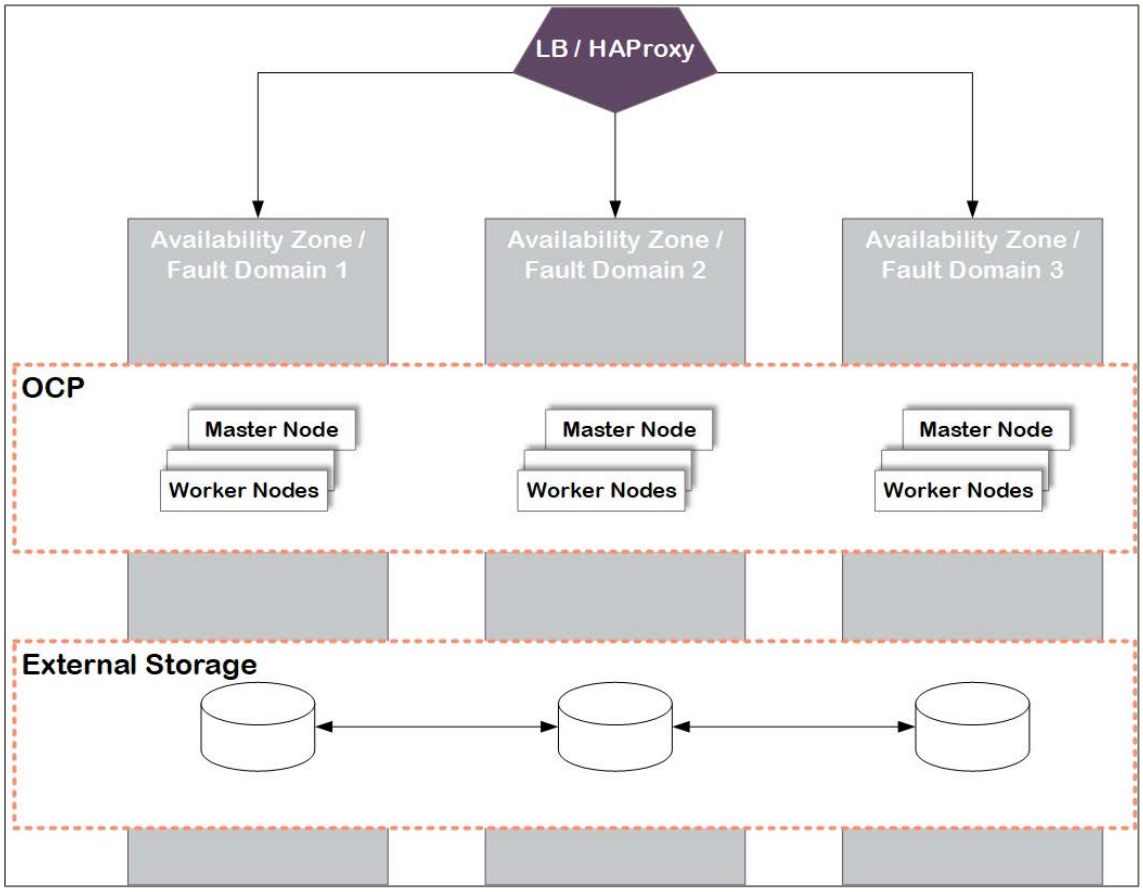

FIGURE 14. Red Hat OpenShift Metro Cluster design overview

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 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 15 shows the Red Hat OpenShift Multi-cluster disaster recovery approach.

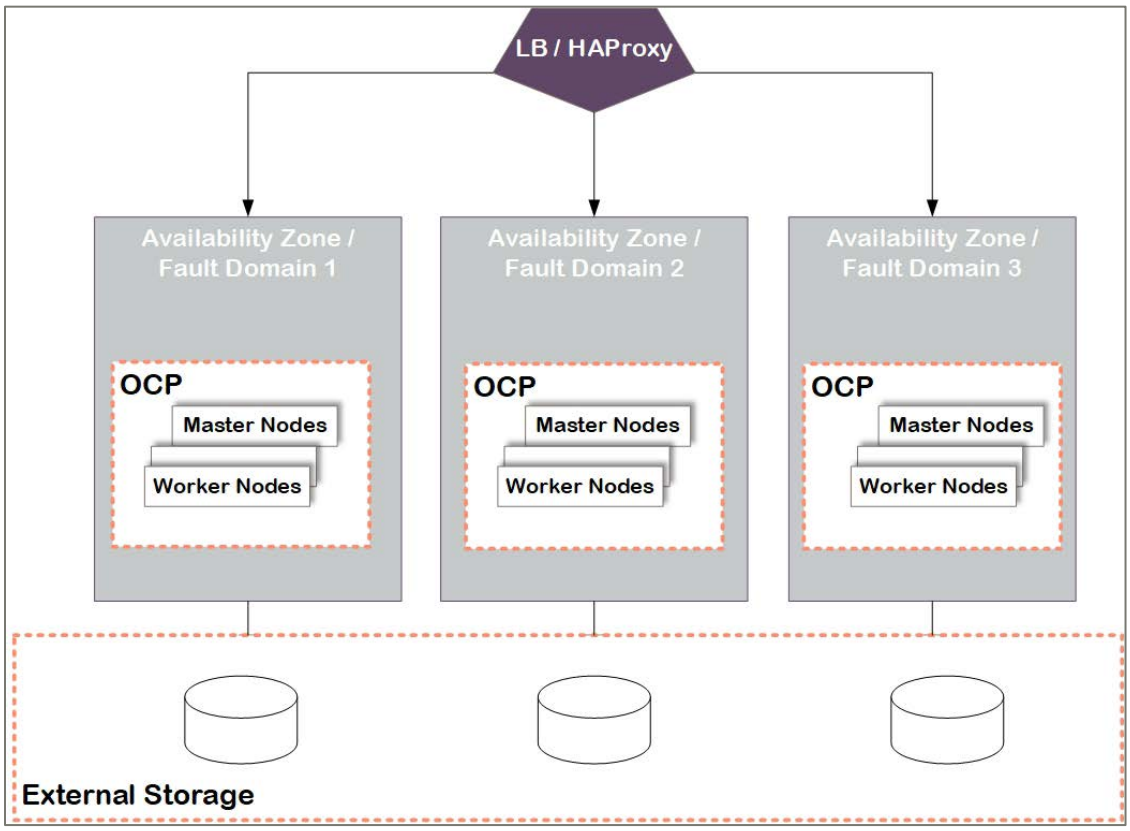

FIGURE 15. Red Hat OpenShift Multi-cluster disaster recovery approach

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

RED HAT 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, but they might not be good candidates for containerization. OpenShift Virtualization, also referred to as container-native virtualization, enables the developers to 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 within a single platform. In addition to this built-in feature of the Red Hat OpenShift Container Platform, OpenShift Virtualization combines the two technologies of virtualization and containerization into a single management platform. With this feature, organizations can 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 16 shows an overview of Red Hat OpenShift Virtualization.

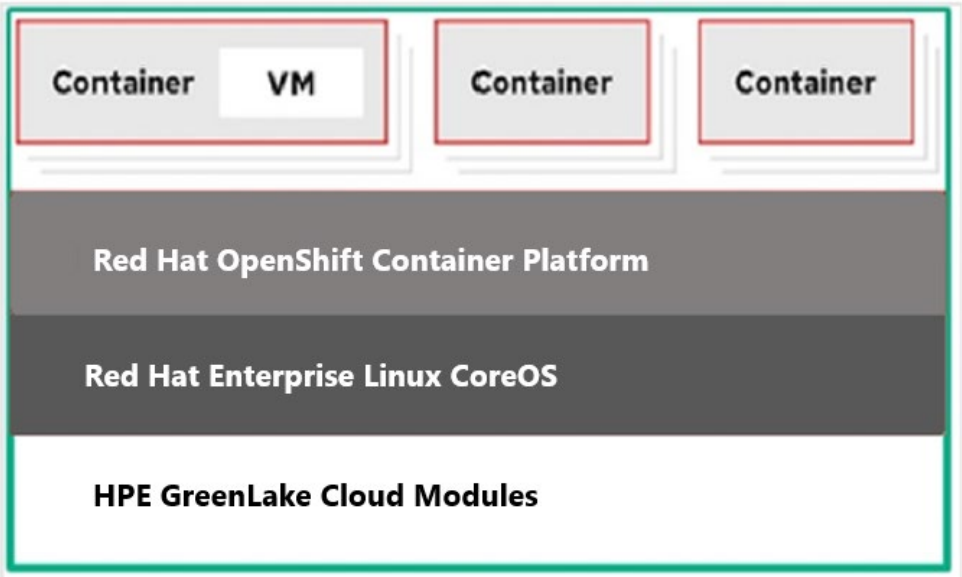

FIGURE 16. Red Hat OpenShift Virtualization overview layout

Red Hat OpenShift Virtualization adds new objects into your Red Hat 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 Red Hat OpenShift Container Platform cluster containers and infrastructure.
- OpenShift Virtualization is tested with Red Hat OpenShift Data Foundation (ODF) and designed to use ODF features for the best experience.
- OpenShift Virtualization allows the usage with either the OVN-Kubernetes or the OpenShift SDN default Container Network Interface (CNI) network provider.
- OpenShift Virtualization is supported for use on Red Hat OpenShift Container Platform 4.10 clusters.

Enabling Red Hat OpenShift Virtualization

Figure 17 shows an overview of Red Hat OpenShift Virtualization.

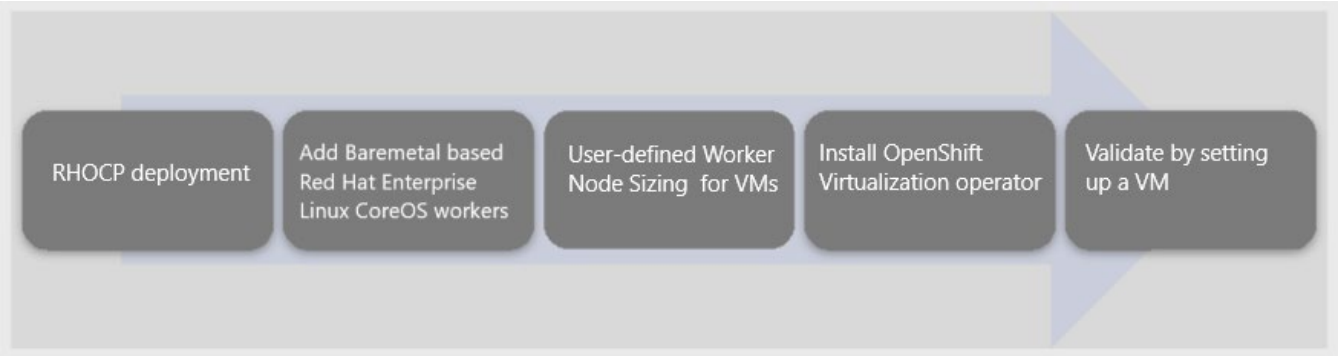

FIGURE 17. OpenShift virtualization setup

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 the OpenShift Virtualization documentation.

SUMMARY

The installation and configuration of the Red Hat OpenShift Container Platform Version 4.10 with HPE Alletra Storage Arrays on the HPE ProLiant DL365 Gen10 Plus Servers 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 eliminate the overhead associated with hypervisors and thus optimizes performance.
- Deploying RHOCP 4.10 on HPE ProLiant DL365 Gen10 Plus 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.

NOTE

Part numbers are at the time of publication/testing and are 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 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
P9K40A001	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
P38578-B21	6	HPE ProLiant DL365 Gen10 Plus 8SFF Configure-to-order Server
P38578-B21 0D1	6	Factory Integrated
P38578-B21 ABA	6	HPE DL365 Gen10 Plus 8SFF CTO Server
HA454A1-001	6	HPE FE ProLiant Svr Pkg 4 SVC
P38687-B21	12	AMD EPYC 7543 2.8GHz 32-core 225W Processor for HPE
P38687-B21 0D1	12	Factory Integrated
P07646-B21	96	HPE 32GB (1x32GB) Dual Rank x4 DDR4-3200 CAS-22-22-22 Registered Smart Memory Kit
P07646-B21 0D1	96	Factory Integrated
P26429-B21	6	HPE ProLiant DL360 Gen10 Plus 8SFF x4 Tri-Mode 24G U.3 BC Backplane Kit
P26429-B21 0D1	6	Factory Integrated
P41400-B21	12	HPE 800GB SAS 24G Mixed Use SFF BC Self-encrypting FIPS PM6 SSD
P41400-B21 0D1	12	Factory Integrated
P41404-B21	12	HPE 1.6TB NVMe Gen4 High-Performance Mixed Use SFF BC Self-encrypting FIPS U.3 CM6 SSD
P41404-B21 0D1	12	Factory Integrated
P38581-B21	6	HPE ProLiant DL365 Gen10 Plus 2SFF U.3 BC Premium Backplane Kit
P38581-B21 0D1	6	Factory Integrated
P31324-B21	6	HPE InfiniBand HDR/Ethernet 200Gb 2-port QSFP56 PCIe4 x16 MCX653106A-HDAT Adapter
P31324-B21 0D1	6	Factory Integrated
P26325-B21	6	Broadcom MegaRAID MR216i-a x16 Lanes without Cache NVMe/SAS 12G Controller for HPE Gen10 Plus
P26325-B21 0D1	6	Factory Integrated
P10097-B21	6	Broadcom BCM57416 Ethernet 10Gb 2-port BASE-T OCP3 Adapter for HPE
P10097-B21 0D1	6	Factory Integrated
P26477-B21	6	HPE ProLiant DL36X Gen10 Plus High-Performance Fan Kit

Component	Qty	Description
P26477-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
E5Y43A	6	HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU
P36657-B21	6	HPE ProLiant DL36X Gen10 Plus 2SFF Tri-Mode Cable Kit
P36657-B21 0D1	6	Factory Integrated
P13771-B21	6	HPE Trusted Platform Module 2.0 Gen10 Plus Black Rivets Kit
P13771-B21 0D1	6	Factory Integrated
P14604-B21	6	HPE Gen10 Plus Chassis Intrusion Detection Kit
P14604-B21 0D1	6	Factory Integrated
P07818-B21	6	HPE DDR4 DIMM Blank Kit
P07818-B21 0D1	6	Factory Integrated
P08040-B21	6	HPE iLO Common Password FIO Setting
P19368-B21	12	HPE DL385 Gen10 1U High-Performance Heat Sink Kit
P19368-B21 0D1	12	Factory Integrated
P26485-B21	6	HPE ProLiant DL300 Gen10 Plus 1U SFF Easy Install Rail Kit
P26485-B21 0D1	6	Factory Integrated
P49803-B21	6	HPE Server Platform LDevID FIO Setting
R9F63A	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE
R9F63A 0D1	2	Factory Integrated
R9F63A B2B	1	Aruba 6300M 48G Power to Port Airflow 2 Fans 1 Power Supply Unit Bundle for HPE PDU
HA454A1-021	2	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	2	Aruba 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable for HPE
R9F61A	1	HPE Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit
R9F61A B2B	1	Aruba 6300M 12VDC 250W 100-240VAC Power to Port Airflow Power Supply Unit for HPE PDU
R9F61A 0D1	2	Factory Integrated
R9F57A	1	Aruba 1U Universal 4-post Rack Mount Kit for HPE
R9F57A 0D1	2	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	1	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

Component	Qty	Description
C7535A	24	HPE RJ45 to RJ45 Cat5e Black M/M 7.6ft 1-pack Data Cable
C7535A 0D1	24	Factory Integrated
H6J85A	1	Factory Integrated
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
120672-B21	1	HPE Rack Ballast Kit
120672-B21 0D1	1	Factory Integrated
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack includes with Complete System Stabilizer Kit
R9G32AAE	1	Aruba Fabric Composer Device Management Service Tier 3 Switch 3-year Subscription E-STU for HPE
xHL251A1	8	HPE Compute Team Day SVC
R9G27AAE	2	Aruba Fabric Composer Device Management Service Tier 4 Switch 3-year Subscription E-STU for HPE
HU4A6A3	1	HPE 3Y Tech Care Essential Service
HU4A6A3 ZSE	6	HPE ProLiant DL365 Gen10 Plus Support
HU4A6A3 ZT0	1	HPE Std-16C ROK/FIO Support
HU4A6A3 ZT3	1	HPE Std-16C Add 16C Support
H1K92A3	1	HPE 3Y Proactive Care 24x7 SVC
H1K92A3 ZSG	1	HPE Aruba 6300M 48 SW Support
H1K92A3 ZND	2	HPE Aruba 8325-32 SW Support
HF385A1	10	HPE Training Credit Servers/HybridIT SVC

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 DL365 Gen10 Plus, hpe.com/servers

Red Hat OpenShift Container Platform, https://access.redhat.com/documentation/en-us/openshift_container_platform/4.10/

Red Hat OpenShift Container Storage, https://access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.10

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

a50007281enw, October 2022

Hewlett Packard Enterprise