HPE Reference Architecture for Red Hat OpenShift Virtualization on HPE ProLiant Compute DL320 and DL380 Gen12 servers using Red Hat OpenShift Container Platform 4.21

Rapid deployment of Red Hat OpenShift Virtualization on HPE ProLiant Compute DL320 and DL380 Gen12 servers using Red Hat OpenShift Container Platform 4.21 automated approach

Contents

Executive summary

Enterprises across industries are accelerating hybrid-cloud initiatives to drive digital transformation, foster innovation, and grow more quickly. Their priorities include speed, agility, simplicity, consistency, and cost efficiency however traditional IT practices and fragmented deployment environments create persistent obstacles:

— Modernizing legacy applications for cloud-native delivery remains complex and slow.

— Operating workloads across multiple clouds is difficult and error prone.

— Provisioning new environments can be slow, delaying development and innovation.

— Vendor lock-in concerns limit strategic flexibility.

— Siloed infrastructure raises both capital and operational costs.

— Rapid deployment of a secure, disconnected end-to-end container platform remains challenging.

This reference architecture explains the design, implementation, and operational considerations for deploying Red Hat OpenShift Container Platform 4.x on HPE ProLiant Compute Gen12 infrastructure. It describes an end-to-end solution, covering physical configuration, hardware and software components, networking and storage design, virtualization via KubeVirt, security hardening, backup and disaster-recovery patterns, and validated capacity/sizing, for CIOs, CTOs, data center managers, enterprise architects, and implementation teams who already hold valid Red Hat subscriptions and are familiar with core networking and HPE server platforms.

The architecture targets production-grade clusters running Red Hat OpenShift Container Platform 4.20/4.21 on HPE ProLiant Compute DL320 Gen12 and HPE ProLiant Compute DL380 Gen12 servers with HPE Alletra Storage MP B10000 and Aruba switching. It consolidates component-level guidance (server roles, iLO management, switch OOBM), platform software (Red Hat OpenShift Container Platform, Red Hat Enterprise CoreOS, Red Hat OpenShift Data Foundation), and virtualization/VM lifecycle features (KubeVirt) into a single operational playbook. Capacity and sizing guidance is included to help translate business requirements into node counts, storage pools, and performance targets.

Operational resilience and security are core themes: the document prescribes platform baselines, Red Hat Advanced Cluster Security controls, software supply-chain protections, compliance and audit readiness across mixed container/VM workloads, and concrete disaster-recovery strategies (metro-cluster HA and multi-cluster DR). It also covers platform-specific caveats such as device, CPU, memory, networking, and storage exceptions when running virtual machines in Red Hat OpenShift and offers best-practice configuration steps to avoid operational pitfalls.

Finally, the reference architecture contains practical, implementable artifacts for deployment and operations: a UPI deployment walkthrough for Red Hat OpenShift Container Platform 4.21, instructions for installing Red Hat OpenShift Data Foundation and the HPE CSI Operator, guidance for enabling and managing Red Hat OpenShift Virtualization, migration modes (cold/warm), GPU and SR-IOV/DPDK considerations, and runbook items for backup/restore and incident response. Use this document to align executive stakeholders on value and risk, and to give implementers a prescriptive, assessor-friendly blueprint for delivering a performant, secure OpenShift platform on HPE Gen12 infrastructure provided by Hewlett Packard Enterprise.

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 Compute DL320 Gen12 servers. This document assumes that the reader is familiar with HPE ProLiant Compute DL320 Gen12 servers, Red Hat OpenShift Container Platform 4.21, core networking, and has a valid Red Hat OpenShift Container Platform and Red Hat Enterprise Linux Subscriptions.

Document purpose: This document describes the benefits and technical details of deploying Red Hat OpenShift Container Platform 4.20/4.21 on HPE ProLiant Compute DL320 Gen12 and HPE ProLiant Compute DL380 Gen12 servers, implementation details, and the processes. This guide is accompanied by a Deployment Guide which can be found here.

Introduction

What's changed since earlier releases

Since OpenShift Virtualization was first introduced on OpenShift 4.5, the solution has continued to evolve through 4.21 with focused improvements that strengthen production readiness, performance, and platform integration:

- Better platform integration and lifecycle management: The virtualization operator and console integrations have matured, simplifying deployment, upgrades and day-2 operations through the standard OpenShift operator and cluster lifecycle workflows.
- Improved migration and availability: Migration tooling and live-migration robustness have been enhanced to reduce downtime and simplify lift-and-shift migrations. Migration Toolkit for Virtualization (MTV) remains central to importing workloads from other hypervisors and cloud targets.
- Higher performance and resource efficiency: Enhancements in scheduling, NUMA awareness and CPU/memory handling improve VM performance and densification on modern servers.
- Accelerator and device support: Expanded support for GPU passthrough, SR-IOV and PCIe device assignment improves support for ML/AI, accelerated compute and high-performance networking workloads.
- Storage and snapshot improvements: Better integration with CSI-based storage, snapshot, and backup tooling simplifies persistent VM storage management within OpenShift storage ecosystems.
- Security and compliance: Strengthened defaults, SELinux and other platform hardening efforts, plus tighter operator-managed policies, improve the security posture for mixed VM/container deployments.
- Windows and guest OS support: Continued improvements for Windows and other guest operating systems enable smoother modernization paths for legacy enterprise applications.
- Developer and observability experience: Console UX improvements, enhanced metrics/telemetry integration, and tighter DevOps toolchain alignment make it easier to build, monitor and operate mixed workload applications.

How OpenShift Virtualization delivers value

Red Hat OpenShift Virtualization (an integrated capability of Red Hat OpenShift) (RHOV) provides a unified, enterprise-grade platform for running both new containerized applications and existing virtual machine workloads. Key benefits include:

- Easier migration: Built-in migration tooling and Red Hat services options help move VMs from hypervisors or to cloud targets while preserving business continuity.
- Faster time to production: A single platform for containers and VMs streamlines pipelines and reduces friction between development and operations.
- Operational consistency: Unified tooling and the OpenShift operator model let teams manage containers, VMs and serverless functions with the same enterprise processes and toolchain.

— Investment protection and modernization: Run legacy virtualized workloads alongside cloud-native services to incrementally modernize applications without disruptive forklift migrations.

— Windows application modernization: Full support for Linux and Windows guests enables lift-and-shift as well as gradual modernization of Windows-based workloads into containerized microservices where appropriate.

— Virtualization management automation: The full end-to-end lifecycle of VMs – from provisioning, to patching, to enforcing configuration standards – as well as quickly and consistently migration to Red Hat OpenShift Virtualization can be achieved by Red Hat Ansible Automation Platform.

HPE Gen12 compute for modern OpenShift Virtualization deployments

HPE ProLiant Compute Gen12 servers (for example, DL360 Gen12 / DL380 Gen12) bring modern compute and I/O platforms that are well matched to the improvements in OpenShift Virtualization through version 4.21.

— New CPU and memory technologies: Support for latest server-class processors and DDR5 memory increases single-socket and cluster performance for both VMs and containers.

— Higher I/O and storage throughput: PCIe Gen5 and enhanced NVMe support accelerate data-intensive workloads and improve persistent VM storage performance.

— Improved accelerator and networking support: Enhanced capability for GPUs, DPUs and SR-IOV networking aligns with RHOV's expanded device passthrough and acceleration features.

— Greater density and efficiency: Higher core counts, memory capacity and platform efficiency allow higher consolidation ratios and lower total cost of ownership (TCO).

— Scale without disruption: Validated architectures and HPE management integrations enable adding compute nodes non-disruptively to scale capacity.

Organizations must adopt cloud-native principles and next-generation infrastructure patterns that let them run and manage virtual machines and containers together. Red Hat and Hewlett Packard Enterprise provide a validated path to do that by combining Red Hat OpenShift and OpenShift Virtualization with HPE ProLiant Compute.

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 Compute DL320 Gen12 servers.
2. To install and configure the Red Hat OpenShift Container Platform 4.x Cluster and Red Hat OpenShift Virtualization feature.
3. Validate the Red Hat OpenShift Virtualization and Red Hat OpenShift Container Platform installation.
4. Migrate Virtual machines from VMware to Red Hat OpenShift Virtualization.

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

The Reference architecture demonstrates a cost-effective yet reliable solution by leveraging the benefits of HPE ProLiant Compute DL360 and DL380 Gen12 servers for compute, storage, networking, and Red Hat OpenShift Container Platform 4.21. This document provides guidance for installing Red Hat OpenShift Container Platform 4.21 on HPE ProLiant Compute DL320 and DL380 Gen12 servers. The solution consists of three (3) HPE ProLiant Compute DL320 Gen12 servers: three (3) HPE ProLiant Compute DL320 Gen12 servers used for the Red Hat Enterprise Linux (RHEL) Kernel-based Virtual Machine(KVM)head nodes and three (3) HPE ProLiant Compute DL360 Gen12 servers used for the solution worker nodes (out of which one node is used as a temporary Red Hat OpenShift Container Platform installer node). HPE ProLiant Compute DL380 Gen12 servers can be added as a cluster for Red Hat OpenShift Data Foundation (ODF).

Persistent storage for this solution is provided by HPE Alletra Storage MP B10000. It delivers mission-critical storage at midrange economics with the only disaggregated, scale-out block storage with a 100% data availability guarantee.

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.

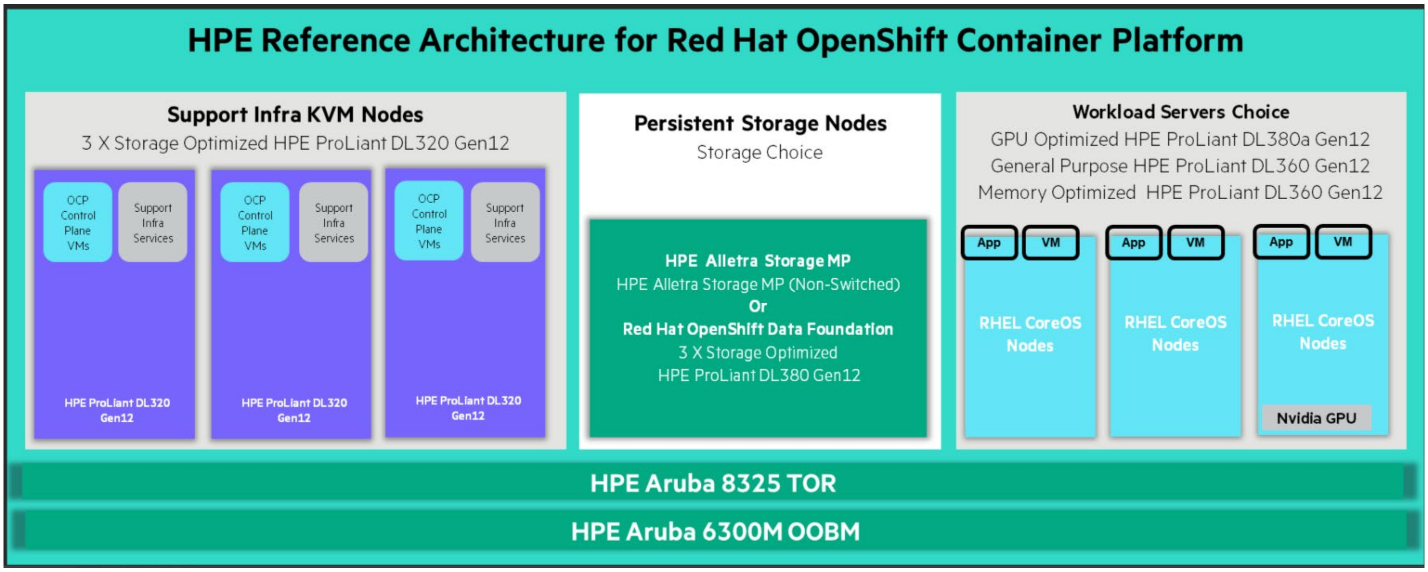

Figure 1. High-level architecture

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) 9.6 on the HPE ProLiant Compute DL320 Gen12 servers and configure the Red Hat OpenShift Container Platform cluster.

Design objectives

The objective of this reference architecture 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 Compute DL320 and DL380 Gen12 servers and HPE Alletra Storage MP B10000 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 RHCOS. These virtual machines are running on RHEL 9.6 and KVM on three (3) HPE ProLiant Compute DL320 Gen12 servers. Three (3) HPE ProLiant Compute DL320 Gen12 servers are deployed as solution worker nodes on bare metal. The temporary installer node is deployed on one of the worker nodes and later configured as a worker node. HPE ProLiant Compute DL380 Gen12 servers can be added as a cluster for Red Hat OpenShift Data Foundation (ODF) external storage or HPE Alletra Storage MP (Non-switched) could be leveraged as persistent volume storage for the solution.

The solution uses internal storage on the HPE ProLiant Compute DL320 Gen12 servers for both the operating system and solution applications. The environmental infrastructure support components (Installer machine, 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 RHCOS 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 Compute DL320 or DL360 Gen12 servers use HPE Alletra Storage MP (Non-switched) 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.

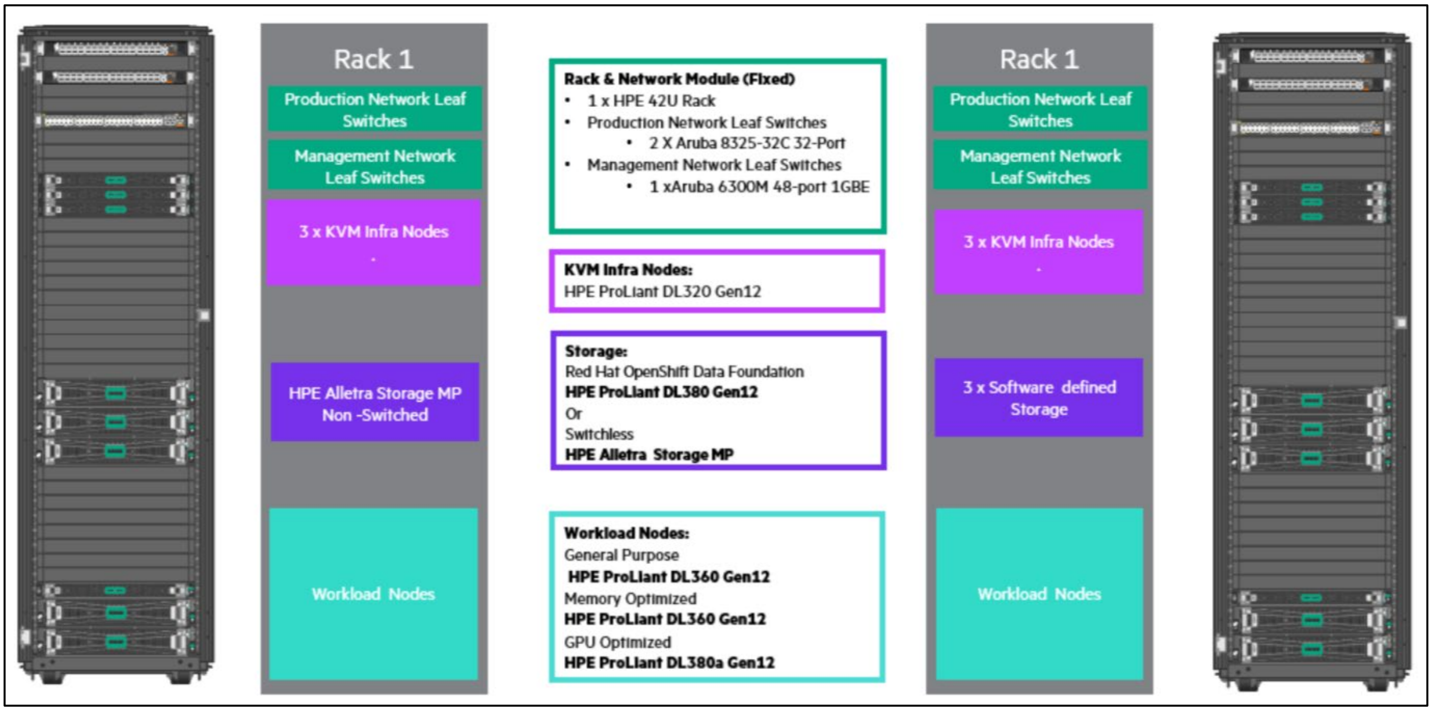

Figure 2. Solution components

Note

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

Additional HPE ProLiant Compute DL360 Gen12 or HPE ProLiant Compute DL380a Gen12 can be added to this solution as per the customer's choice of configuration workload options.

Solution components

This section provides 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 are utilized in this solution

Component	Qty	Description
HPE ProLiant Compute DL320 Gen12 servers	3	Provides capacity for head nodes with OpenShift Master and Bootstrap KVM, VMS, HAProxy, DNS, Proxy

Component	Qty	Description
HPE ProLiant Compute DL320 Gen12 servers Or HPE ProLiant Compute DL360 Gen12 servers Or HPE ProLiant Compute DL380a Gen12 servers	3	Provide Red Hat OpenShift worker nodes
HPE Alletra Storage MP B10000(Optional)	1	External iSCSI storage for persistent volumes (non-switched with one controller two data nodes)
HPE ProLiant Compute DL380 Gen12 servers	3	Red Hat OpenShift Data Foundation nodes – Internal or external 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	CPU Core	RAM	Storage
KVM head nodes	RHEL 9.6	64	512GB	OS Disk: HPE NS204i-u v2 480GB NVMe Data Disk: 8x 3.84 TB
Workload nodes option 1 General purpose	RHCOS	64	512GB	OS Disk: HPE NS204i-u v2 480GB NVMe
Workload nodes option 2 memory optimized	RHCOS	64	1024 GB	OS Disk: HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device
Workload nodes option 3 GPU optimized	RHCOS	172	3072 GB	OS Disk: HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device Data Disk: 4x 3.2 TB
Red Hat ODF external nodes	RHEL 9.6	32	1024 GB	OS Disk: 2x 1.6 TB Data Disk: 8x 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	250 GB

Note

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

HPE ProLiant Compute DL320 Gen12 server

The HPE ProLiant Compute DL320 Gen12 maximizes your rack utilization while mitigating virtualization risks in power-constrained environments. Power your workloads with a server providing greater expansion capabilities compared to previous generations. The latest Intel® Xeon® 6 processors with up to 144 cores, increased memory capability (up to 4 TB), and high-speed PCIe Gen5 help provide a high-performance solution with better datacenter efficiency.

Figure 3 shows the HPE ProLiant Compute DL320 Gen12 server.

Figure 3. HPE ProLiant Compute DL320 Gen12 server

Table 4 lists the hardware configuration used in this solution.

Table 4. Hardware configuration in each of the HPE ProLiant Compute DL320 Gen12 servers

Component	Description
Processor	1x Intel Xeon 6767P 2.4GHz 64-core 350W Processor for HPE
Memory	16x 32GB (1x32GB) Dual Rank x8 DDR5-4800
Network	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T Adapter for HPE
Smart Array Controller	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
Disks	8x HPE 3.84TB SFF SATA SSD

HPE ProLiant Compute DL360 Gen12 server

The HPE ProLiant Compute DL360 Gen12 is a compact 1U 2P server that delivers exceptional compute performance, memory density with scalability and high-speed data transfer rate to run your most demanding applications. Powered by Intel® Xeon® 6 processors with up to 144 cores and up to 8 TB of memory as well as 20 EDSFF E3.S NVMe drives, the HPE ProLiant Compute DL360 Gen12 is an ideal hybrid cloud platform for enterprise applications and workloads. This server also provides extreme scalability through a hybrid front storage cage that supports not only SFF and E3.S drives, but also an OS Boot device and two front OCP NICs as optional.

Figure 4 shows the HPE ProLiant Compute DL360 Gen12 server.

Figure 4. HPE ProLiant Compute DL360 Gen12 server

Table 5 lists the hardware configuration used in this solution.

Table 5. Hardware configuration in each of the HPE ProLiant Compute DL360 Gen12 servers

Component	Description
Processor	Intel Xeon 6530P 2.3GHz 32-core 225W Processor for HPE
Memory	HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
Network	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
Disks	OS Disk: HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device

HPE ProLiant Compute DL380 Gen12 server

The HPE ProLiant Compute DL380 Gen12 is a scalable 2U 2P server that delivers exceptional compute performance, memory density with scalability and high-speed data transfer rate to run your most demanding applications. It is powered by Intel® Xeon® 6 processors with up to 144 cores and up to 8 TB of memory as well as 36 EDSFF E3.S drives, the HPE ProLiant Compute DL380 Gen12 is an ideal hybrid cloud platform for enterprise applications and workloads.

Figure 5 shows the HPE ProLiant Compute DL380 Gen12 server. The HPE ProLiant Compute DL380 Gen12 is used for the ODF nodes.

Figure 5. HPE ProLiant Compute DL380 Gen12 server

Table 6 lists the hardware configuration in each of the HPE ProLiant Compute DL380 Gen12 servers used in this solution.

Table 6. Hardware configuration in each of the HPE ProLiant Compute DL380 Gen12 servers

Component	Description
Processor	2 x Intel Xeon 6530P 2.3GHz 32-core 225W Processor for HPE
Memory	32 x HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400
Network	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
Array controller	HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device
Disks	4x HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD

HPE ProLiant Compute DL380a Gen12 server

The HPE ProLiant Compute DL380a Gen12 server is a rack-optimized, 4U 2P solution that delivers breakthrough performance with advanced GPU accelerators on an ultra-scalable architecture. It is powered by Intel® Xeon® 6 processors with up to 144 cores and eight double-wide GPUs as well as increased memory bandwidth and high-speed PCIe Gen5 I/O, the HPE ProLiant Compute DL380a Gen12 server is a perfect solution for AI inference workloads.

Figure 6 shows the HPE ProLiant Compute DL380a Gen12 server.

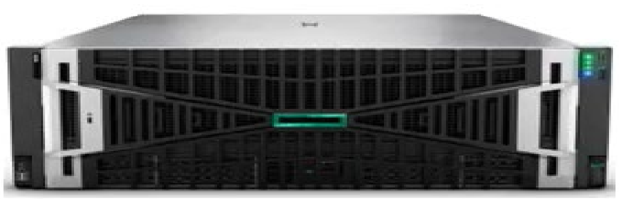

Figure 6. HPE ProLiant Compute DL380a Gen12 server

Table 7 lists the hardware configuration in each of the HPE ProLiant Compute DL380a Gen12 servers used in this solution.

Table 7. Hardware configuration in each of the HPE ProLiant Compute DL380 Gen12 servers

Component	Description
Processor	2 x Intel Xeon-6527P 3.0GHz 24-core 225W Processor for HPE
Memory	12 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 Gen12 x16 Lanes 8GB Cache PCI SPDM Plug-in Storage Controller
GPU	NVIDIA L40S 48GB PCIe Accelerator OS Disk: HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device
Disks	6x HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD

HPE Alletra Storage MP B10000 (Ethernet)

HPE GreenLake for Block Storage built on HPE Alletra Storage MP provides mission-critical storage at midrange economics with the industry's first disaggregated, scale-out block storage with a 100% data availability guarantee. Built on new HPE Alletra Storage MP hardware and managed via GreenLake, this unique block storage offer brings an intuitive cloud experience, efficient scale, and extreme resilience and performance to mission-critical apps — from large databases to modern cloud native apps to consolidated mixed workloads - at a midrange price point.

Figure 7. HPE Alletra storage MP (Ethernet iSCSI)

HPE iLO

HPE Integrated Lights Out (iLO) is embedded in HPE ProLiant Compute DL320 and DL380 Gen12 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 8 shows the HPE Aruba 8325-32C BF switch.

Figure 8. HPE Aruba 8325-32C BF switch

HPE Aruba CX 6300M OOBM switch

The HPE Aruba CX 6300M 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 verify 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 future 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 9 shows the HPE Aruba 6300M OOBM Switch.

Figure 9. 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. 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 4.20 deployment on bare metal infrastructure, you must use RHCOS for all Red Hat OpenShift Container Platform control plane nodes, Bootstrap nodes, and RHCOS for worker nodes.

HPE CSI operator for Kubernetes is an operator for the HPE CSI Driver for Kubernetes to provision and manage storage resources for HPE Alletra Storage MP Block, Alletra 5000/6000/9000.The HPE Container Storage Interface (CSI) Driver is a multi-vendor and multi-backend driver where each implementation has a Container Storage Provider (CSP). It allows any vendor or project to develop its own Container Storage Provider (CSP) by using the CSP specification. This makes it easier 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 4.21

Red Hat OpenShift Data Foundation is software-defined storage that is optimized for container environments. It runs as an operator on Red Hat OpenShift Container Platform 4.21 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 8 lists the major software used in this solution.

Table 8. Software used in this solution.

Component	Versions	Usage
Red Hat Enterprise Linux CoreOS	9.6	Red Hat OpenShift control plane VMs and Red Hat worker nodes bare metal
Red Hat OpenShift Container Platform	4.21	Red Hat OpenShift control plane nodes on KVM virtual machines and bare metal worker nodes
Red Hat Enterprise Linux	9.6	KVM head node and Red Hat worker nodes bare metal
Red Hat Enterprise Linux	9.6	Installer Virtual Machine required to execute automation scripts
HPE Alletra Storage MP (Ethernet)	32c 2N	External storage
Red Hat OpenShift Data Foundation	4.21	External Storage

Capacity and sizing for a Red Hat OpenShift Container Platform 4.21

Size for a Red Hat OpenShift Container Platform 4.21 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 documentation.

KubeVirt

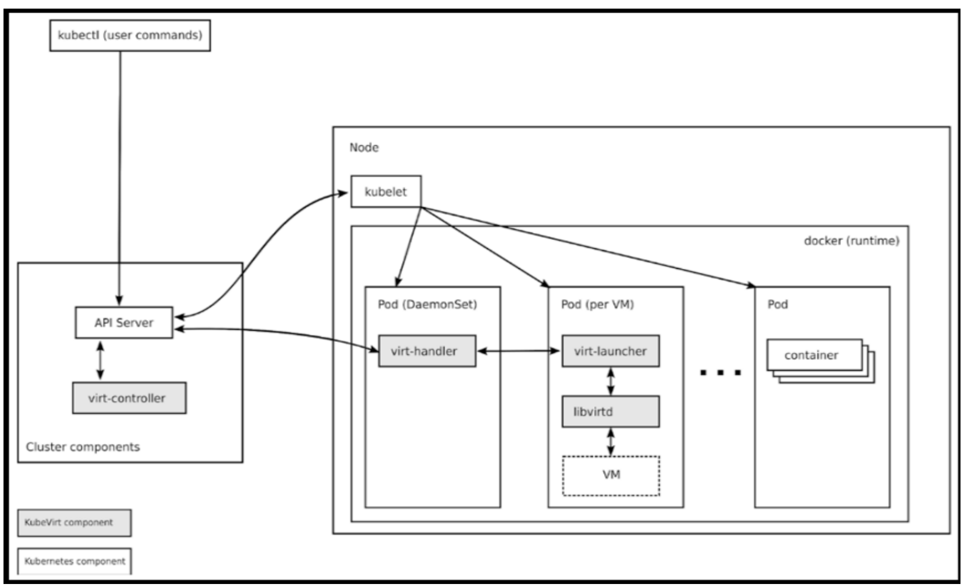

Figure 8. KubeVirt architecture

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

Red Hat OpenShift Virtualization 4.21 is powered by the open-source KubeVirt project, which integrates traditional virtual machines (VMs) with Kubernetes by treating them as native Kubernetes objects. The architecture works by extending the Kubernetes API with Custom Resource Definitions (CRDs) and leveraging the KVM hypervisor for VM execution within specialized pods.

Core architectural components

The main components of the KubeVirt architecture in OpenShift Virtualization are:

Virtual Machine (VM) and Virtual Machine Instance (VMI) CRDs: These custom resources define the specifications and lifecycle of VMs as Kubernetes-native objects. The Virtual Machine CRD represents the desired state of a VM, while the Virtual Machine Instance (VMI) represents a running instance of that VM on a specific node.

virt-controller: A cluster-level controller that watches the VMI CRDs and ensures the cluster's actual state matches the desired state. It manages VM operations such as creation, updates, and deletion, and orchestrates actions like live migration.

virt-handler: This agent runs as a DaemonSet on each worker node in the cluster. It communicates with the virt-controller and manages node-specific tasks, such as preparing the environment for a VM and monitoring its status.

virt-launcher: A pod created for each VMI. This pod is the container for the VM process itself, providing the necessary isolation and integrating the VM with Kubernetes networking and storage. It embeds an instance of libvirtd to manage the VM's lifecycle and interface with the underlying hypervisor.

KVM and libvirt: The underlying virtualization technologies used for executing the VM. KVM is a module in the Linux kernel that acts as a hypervisor, and libvirt is the virtualization API used by virt-launcher to interact with KVM.

Key integrations

KubeVirt integrates seamlessly with existing Kubernetes infrastructure:

- Networking: It uses Kubernetes networking solutions, including standard OVN-Kubernetes, Multus (for multiple network interfaces), and SR-IOV (for high-performance networking).
- Storage: VMs use Kubernetes Persistent Volume Claims (PVCs) and Storage Classes to access reliable storage, which can include solutions like Ceph RBD or NFS. The Containerized Data Importer (CDI) is used for managing data volumes and importing existing VM images.
- Management: Users can manage VMs through the Red Hat OpenShift web console's Virtualization perspective or using standard CLI tools like oc and the virtualization-specific virtctl command.
- Value adds: HPE Morpheus Enterprise provides deep, enterprise-grade integration with Red Hat OpenShift Virtualization, enabling organizations to manage virtual machines (VMs) alongside containerized workloads from a single, unified control plane.

Capacity sizing Red Hat OpenShift Virtualization

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

Table 9. Scale limits for OpenShift Virtualization 4.x, per cluster

Objective	4.x tested maximum	Theoretical limit
Virtual Machine Maximums	The following maximums apply to the OpenShift VM run on OpenShift Virtualization	
Maximum virtual CPUs per virtual machine	216	255
Maximum memory per virtual machine	6 TB	16 TB
Minimum memory per virtual machine	512MB	
Maximum single disk size per virtual machine	20 TB	100 TB
Maximum number of hot-pluggable disks per virtual machine	255	
Host Maximums	The following maximums apply to the OpenShift hosts used for OpenShift Virtualization	
Logical CPU cores or threads	Same as RHEL with KVM	
RAM	Same as RHEL with RHEL as KVM	
Simultaneous live migrations	Default to two outbound migrations per node, and five concurrent migrations per cluster	Depends on NIC bandwidth
Live migration bandwidth	There is no default bandwidth limit for live migrations	Depends on NIC bandwidth
Cluster Maximums	The following maximums apply to objects defined in OpenShift Virtualization	
Maximum number of attached PVs per node	Not specified (CSI-dependent	CSI storage provider dependent
Maximum PV size	Not specified (CSI-dependent	CSI storage provider dependent
Maximum number of hosts	500 (<100 recommended)2	Same as OpenShift release
Maximum number of defined VMs	10,000(per cluster)	Same as OpenShift release

For more information, see the OpenShift Virtualization - Tuning and Scaling Guide which is a supplemental document for OpenShift Virtualization. It focuses on fine-tuning cluster scalability and VM performance for several use cases and environments, for a variety of workloads and cluster sizes. This document provides guidance for scaling up and scaling out an environment, up to 100 OpenShift nodes. Additional OpenShift Virtualization cluster guidance can also be found in the Reference Implementation Guide.

For sizing a Red Hat OpenShift Container Platform 4.21 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 4.20 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 10 lists the various networks used for this solution. All the cluster nodes and iPXE servers are connected to the same network.

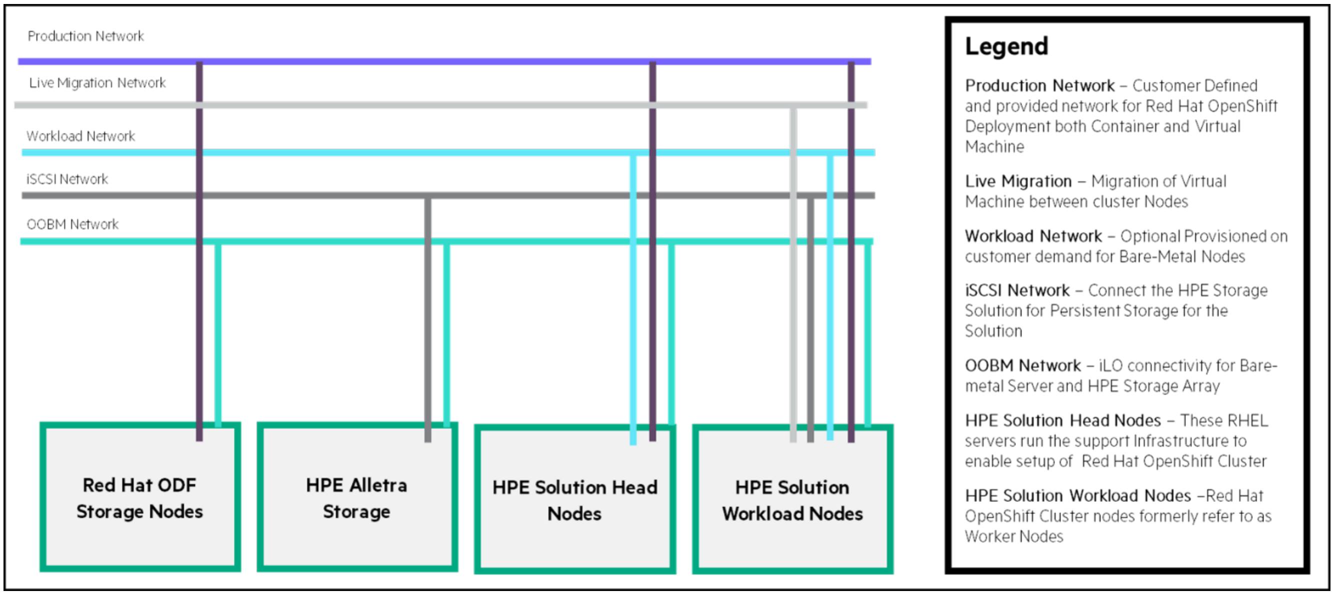

Figure 10. Networks for solution

Note

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

Storage

In the internal storage mode, the storage for the OS and internal persistent volume is provided by the local storage disks (SSD) on the HPE ProLiant Compute DL320 and HPE ProLiant Compute DL380 Gen12 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 OpenShift Data Foundation (ODF) operator installation will be using the Local Storage operator. ODF provides persistent storage services including OpenShift, monitoring, logging and registry, and other container-based applications that require persistent storage.

Table 10 lists all volumes used within the solution for the storage systems and highlights what storage provides the capacity and performance for each function.

Table 10. Details of the volume

Source	Volume/Disk Function	Hosts	Shared/Dedicated
Local storage on the servers	Red Hat OpenShift Container Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated
	Operating System	All Nodes	Dedicated
HPE Alletra Storage MP B10000	iSCSI Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated
HPE ProLiant Compute DL380 Gen12 (Internal storage ODF)	Persistent Volume	Red Hat OpenShift Container Platform worker nodes	Dedicated

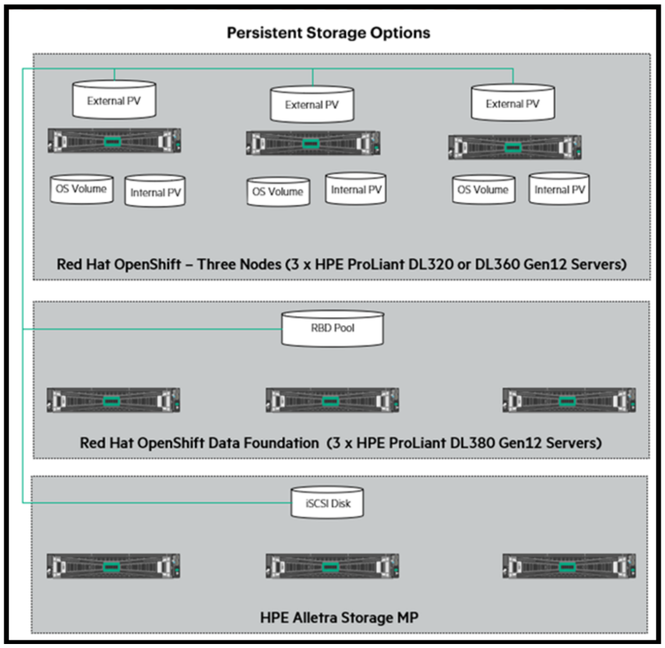

Figure 11. Logical storage layout for Red Hat OpenShift persistent volume options

Note

HPE ProLiant Compute DL380 Gen12 servers can be added to this solution as per the customer's choice of configuration storage options.

HPE Alletra Storage MP B10000 provides external storage to solution.

Storage subscriptions

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

If you want Advance replication features use ODF External and require a license for Red Hat OpenShift Container Platform Plus (OPP) and Advance Red Hat OpenShift Data Foundation Advance subscription

Deployment overview

This section explains in detail the deployment of Red Hat OpenShift Container Platform 4.20 using internal and external storage mode. In the external storage mode, HPE Alletra Storage MP is connected via the iSCSI network to the solution worker nodes.

Deploying the Red Hat OpenShift Container Platform 4.21 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, the solution control plane ignition files, and Workload ignition files. The ignition files are used to configure RHCOS 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 12 explains the Red Hat OpenShift Container platform 4.21 deployment process.

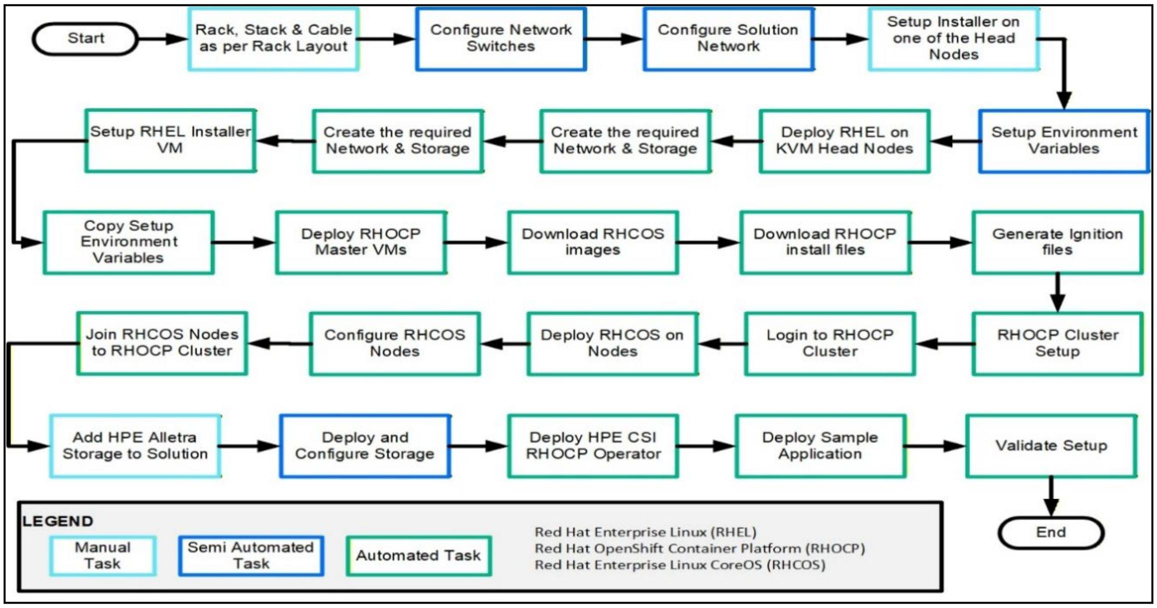

Figure 12. Deployment process for Red Hat OpenShift Container Platform 4.21 cluster using the UPI

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 external storage mode with a minimal cluster of three (3) ODF 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 ODF configuration and procedure on storage sizing are described in the Deployment guide at here.

HPE CSI Operator for Kubernetes

An operator for the HPE CSI Driver for Kubernetes, to provision and manage storage resources for HPE Alletra Storage MP B10000, HPE Alletra 5000/6000/9000, Nimble Storage, Primera and 3PAR.

Deploying Red Hat OpenShift Virtualization

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

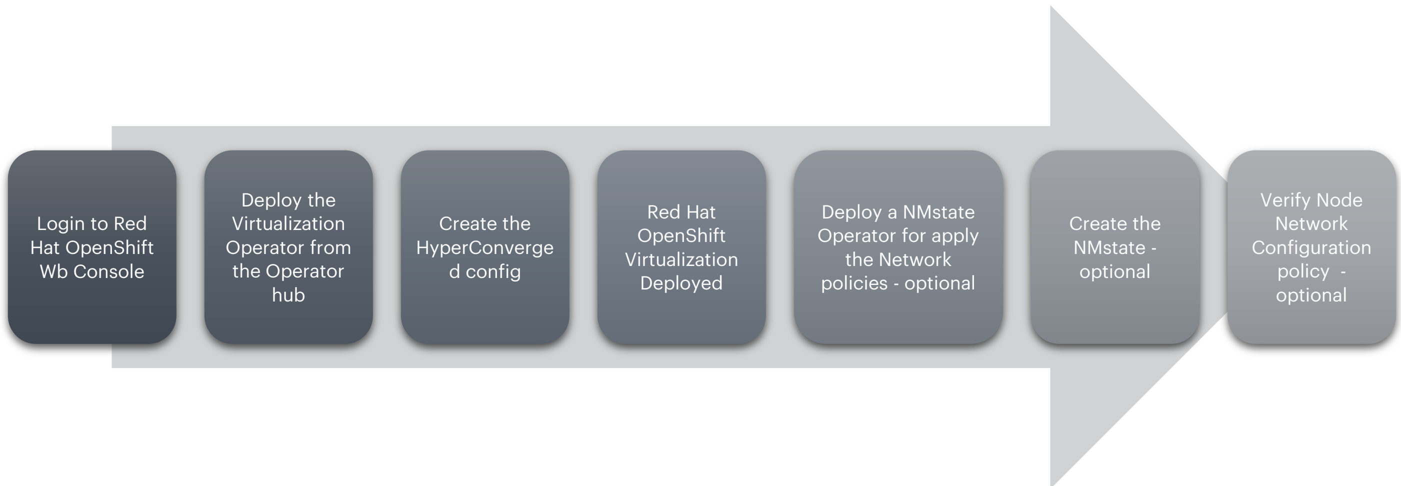

Figure 13. Red Hat OpenShift Virtualization deployment flow

Red Hat documentation states that NMState is used for post-installation networking tasks, such as:

— Creating Linux bridges

— Configuring a dedicated network for live migration

— Managing secondary NICs

Therefore, NMState is optional to install OpenShift Virtualization 4.21 however is needed for advanced networking.

Managing common tasks with Red Hat OpenShift Virtualization

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

— Importing the image as a container disk from a registry.

— Importing the image from a web page.

— Uploading the image from a local machine.

— Cloning a persistent volume claim (PVC) that contains the image.

— Managing virtual machine.

Virtual machine wizard

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

— General: Fill in essential details such as the virtual machine name, namespace, and template. Required fields are marked with an asterisk (*). You can also select a template to autofill some fields.

— Networking: Configure Network Interface Cards (NICs) for your virtual machine. By default, a nic0 NIC is attached, but you can add more if needed. NICs can also be created after the virtual machine is created.

— Storage: Add disks to your virtual machine. You can create additional disks and customize their settings. If you choose URL or container as the source in the general step, a root disk is automatically created and attached as the bootable disk. You can modify the root disk but cannot remove it. A bootable disk is not required for virtual machines provisioned from a PXE source if no disks are attached.

— Review and Create: Review your configuration and create the virtual machine. The results screen displays the JSON configuration file for the virtual machine.

Note

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

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

Connecting to virtual machines through a variety of consoles and CLI tools

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

Accessing virtual machine consoles in the OpenShift Container Platform web console:

— Connecting to the serial console

— Connecting to the VNC console

— Connecting to a Windows virtual machine with RDP

Accessing virtual machine consoles by using CLI commands

— Accessing a virtual machine via SSH by using `virtctl`.

Importing and cloning existing virtual machines

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

— Importing the image from a web page.

— Importing the image as a container disk from a registry

— Uploading the image from a local machine.

— Cloning a persistent volume claim (PVC) that contains the image.

Managing network interface controllers and storage disks attached to virtual machines

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

— OpenShift Virtualization networking overview

— Storage configuration overview

Live migrating virtual machines between nodes

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

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

Data plane development kit and SR-IOV (Single Root I/O Virtualization)

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

DPDK and SR-IOV: DPDK provides a set of libraries and drivers for fast packet processing in user space. By leveraging DPDK, you can achieve lower latency and higher throughput for packet handling.

SR-IOV allows you to share physical network interfaces among multiple virtual machines (VMs) while maintaining hardware-level isolation. It is particularly useful for high-performance workloads.

Graphics processing unit acceleration in OpenShift Virtualization

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

Cross-cluster migration with User-Defined Networks (UDN)

This allows virtual machines to be migrated between OpenShift clusters while maintaining consistent, isolated L2 or L3 network connectivity across those environments. UDNs implemented using User Defined Network and Cluster User Defined Network custom resources—create dedicated tenant networks that operate independently of the default pod network, enabling VMs to communicate over private, cluster-spanning network segments. When combined with the Migration Toolkit for Virtualization (MTV), which supports cold migration between remote OpenShift Virtualization clusters, UDN ensures that migrated VMs land in an identical network topology on the target cluster, preserving IP addressing, isolation boundaries, and multi-namespace connectivity. This makes UDN especially valuable for cross-cluster relocation, data center consolidation, and multi-cluster lifecycle operations, because it decouples workload mobility from physical network dependencies while maintaining consistent L2 adjacency for applications that require it. UDN thereby enhances the reliability and predictability of cross-cluster migrations by ensuring the VM's virtual network environment remains intact across both source and destination clusters.

CPU Load-Aware VM rebalancing

Red Hat OpenShift Virtualization CPU Load-Aware VM rebalancing is an enhancement introduced in OpenShift 4.20 that intelligently distributes virtual machines across cluster nodes based on actual CPU utilization, rather than relying solely on requested CPU resources. In many real-world environments, VMs with identical resource requests may generate very different CPU loads, causing hotspots where heavily loaded VMs accumulate on the same node and compete for CPU time. CPU load-aware rebalancing detects these imbalances using utilization metrics and automatically triggers live migrations to move busy VMs onto less-loaded nodes, ensuring more even

CPU distribution and preventing performance degradation caused by CPU contention. This feature improves overall VM performance, maximizes hardware efficiency, and provides a more autonomous, self-optimizing virtualization environment within OpenShift.

Scheduler intelligence that optimizes placement

Red Hat OpenShift Virtualization introduces scheduler intelligence that optimizes placement and performance for mixed VM and container workloads by combining Kubernetes' native scheduling with virtualization-aware insights. By default, OpenShift schedules virtual machines based on their declared vCPU and memory requests; however, real-world environments often experience imbalances when heavily utilized VMs accumulate on a subset of nodes. To address this, OpenShift Virtualization incorporates CPU load-aware rebalancing, which monitors actual CPU utilization and detects hotspots where busy VMs cause contention. When such imbalances occur, the platform leverages the descheduler and live migration to proactively redistribute VMs across the cluster, ensuring that CPU-intensive VMs are spread out and that underutilized nodes are efficiently used. This enhances overall cluster fairness and ensures containers and VMs can coexist without one starving the other of compute resources—especially important given Kubernetes' mixed-workload scheduling model. As a result, OpenShift provides a more intelligent, self-optimizing environment that preserves performance, prevents node overload, and maintains consistency across diverse workload types.

Migration Toolkit for Virtualization

The Migration Toolkit for Virtualization (MTV) enables you to migrate virtual machines from VMware vSphere, Red Hat Virtualization, or OpenStack to OpenShift Virtualization running on Red Hat OpenShift Virtualization platform. The Migration Toolkit for Virtualization (MTV) streamlines and automates the process of migrating virtual machine workloads into Red Hat OpenShift Virtualization, enabling a consistent, efficient, and low-risk transition from traditional virtualization platforms.

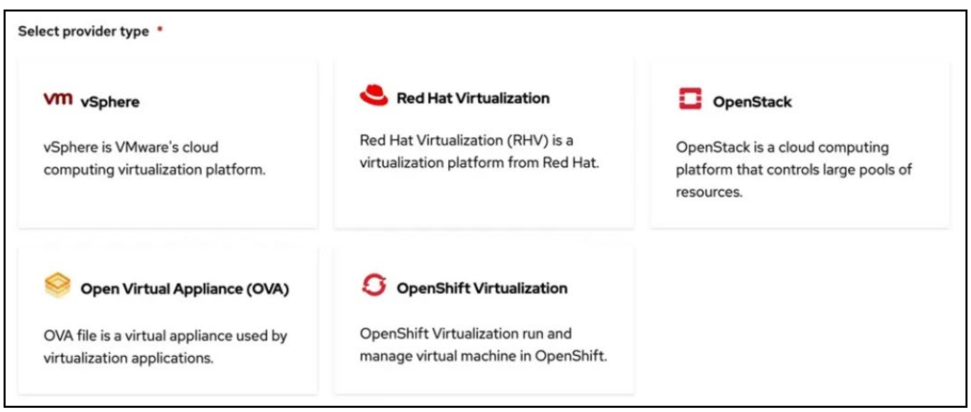

Figure 14. Migration Toolkit for Virtualization supported providers

Cold migration (default migration mode)

MTV supports cold migration as the default approach, where the source virtual machines are powered off while their data is copied to the OpenShift Virtualization environment. Cold migration is supported from:

— VMware vSphere

— Red Hat Virtualization (RHV)

— OpenStack

— Remote OpenShift Virtualization clusters

Cold migration is the simplest and safest method because the VM remains shut down during the data transfer, ensuring data consistency throughout the migration.

Warm migration (pre-copy migration)

MTV also supports warm migration from:

— VMware vSphere

— Red Hat Virtualization (RHV)

In warm migration, MTV first performs a pre-copy stage where most VM data is transferred while the source VM remains running. During the final cutover stage, the VM is powered off and only the remaining changed blocks are copied, minimizing downtime. Warm migration is designed for environments where lowering service interruption is essential, offering a near-seamless transition with shorter application downtime than cold migration.

Accelerating deployment

Automating 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 4.21. 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 15 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.

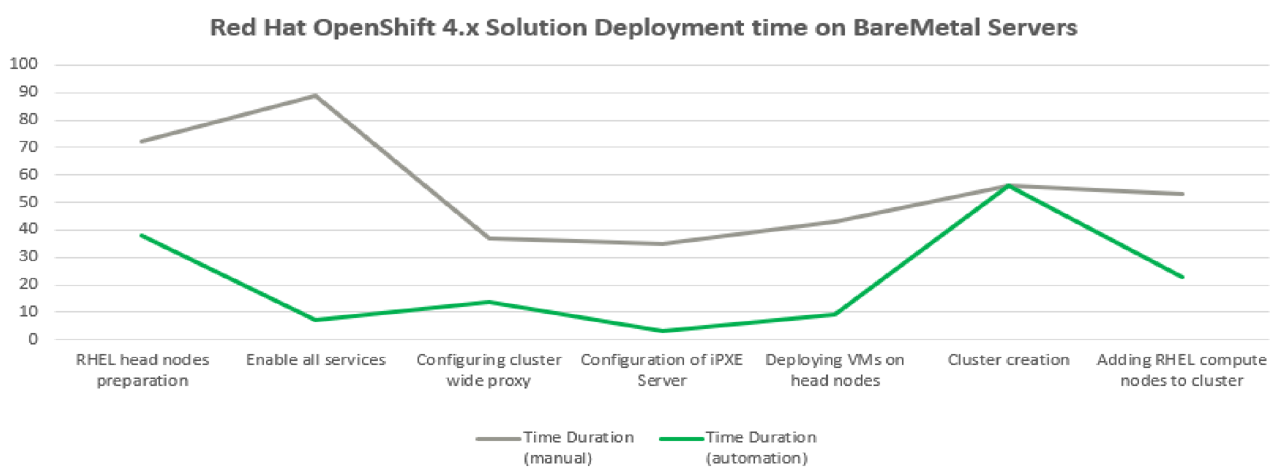

Figure 15. Red Hat OpenShift 4.x solution deployment manual and automation timelines on bare metal

Figure 16 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.

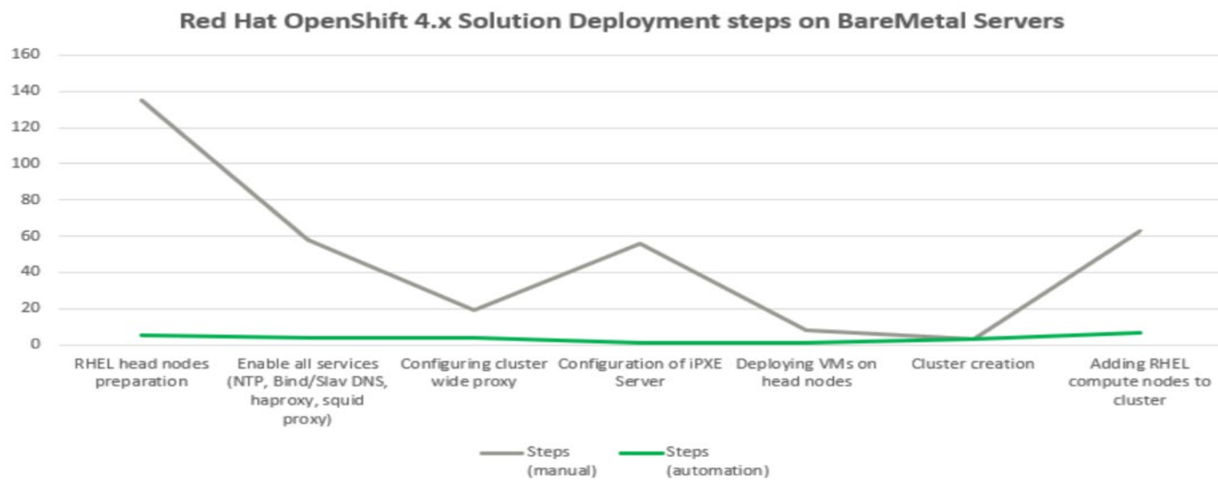

Figure 16. Red Hat OpenShift 4.x solution deployment manual and automation steps on bare metal

Securing Red Hat OpenShift with Red Hat Advanced Cluster Security

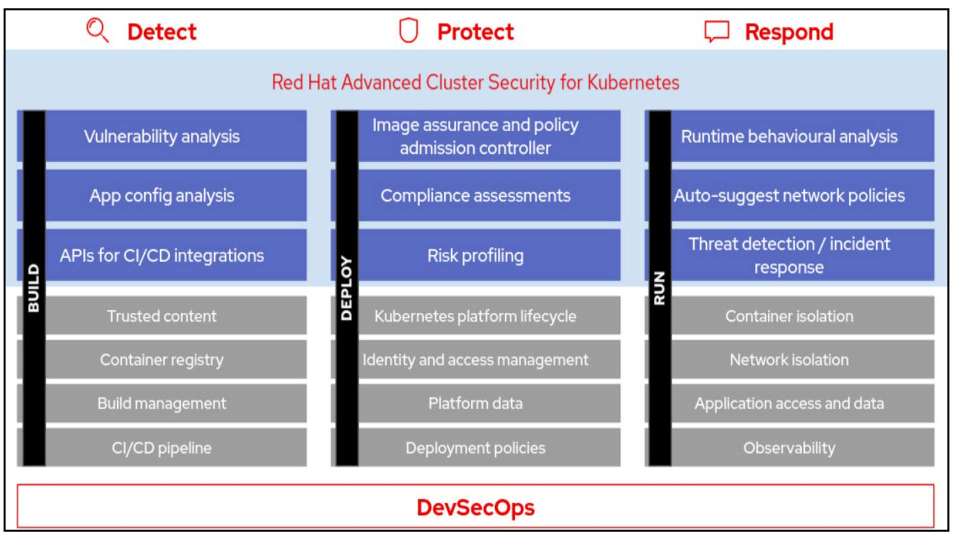

Figure 17. Red Hat Advanced Cluster security

Red Hat Advanced Cluster Security for Kubernetes is a comprehensive security solution designed to protect containerized applications and Kubernetes environments. Red Hat Advanced Cluster Security enhances the security posture of a Kubernetes environment, reducing the risk of breaches and ensuring business continuity. This translates to protecting brand reputation and trust. By leveraging a comprehensive security solution, you can potentially lower the costs associated with security incidents and compliance fines. The managed service option also reduces the need for in-house security expertise, allowing you to allocate resources more effectively. Implementing robust security measures can be a key differentiator in the market, showcasing your commitment to protecting customer data and adhering to industry standards.

— Enhanced Security: It provides robust security features to protect Kubernetes' workloads across various environments, including on-premises, cloud, and hybrid platforms.

— Integration with DevOps: The platform integrates seamlessly with DevOps tools, enabling automated security checks and policies that help mitigate risks during the development and deployment processes.

— Vulnerability Management: Continuous scanning and vulnerability management are core features, ensuring that container images and Kubernetes configurations are secure throughout the software development lifecycle.

— Compliance and Auditing: It supports compliance with industry standards like CIS Creating and managing Linux and Windows virtual machines (VMs).

Red Hat Advanced Cluster Security for Kubernetes (RHACS/ACS) adds tangible security value to Red Hat OpenShift Virtualization 4.21 by extending Kubernetes-native security controls across the platform your VMs run on (namespaces, pods, networks, images, and nodes). While ACS is container/Kubernetes-centric, OpenShift Virtualization runs each VM inside a Kubernetes pod (virt-launcher) as a non-root, restricted workload, so ACS's prevention, detection, and compliance features directly strengthen the VM hosting surface, its network posture, and the surrounding supply chain. OpenShift Virtualization already runs VMs as unprivileged pods under the restricted profile and uses SCCs for the Kubevirt-controller. ACS complements this with preventive (admission, policy), detective (sensor/collector insight), and reporting (compliance) controls, closing gaps that aren't covered by SCCs alone.

Harden the platform your VMs run on (KSPM and policy guardrails)

Kubernetes Security Posture Management (KSPM): ACS continuously evaluates cluster configuration (RBAC, namespaces, pod security, node settings), helping you harden the substrate that runs virt-launcher pods and KubeVirt controllers. This reduces blast radius for VM workloads.

Admission control and enforce-before-run: ACS policies can block risky deployments (e.g., privileged containers, hostPath mounts, outdated base images) that would otherwise weaken nodes hosting your VMs. This aligns well with OpenShift Virtualization model where VMs run in unprivileged pods under restricted SCCs.

Why it matters for VMs: VMs inherit the security of the pod/node they run on. Strengthening cluster posture and preventing unsafe DaemonSets/sidecars or misconfigurations reduces the chance that a compromised container undermines VM hosts or migration controllers.

Network exposure control for VM traffic (network graph to policies)

Network visibility and policy generation: ACS provides a network graph and can help generate/refine Kubernetes Network Policies, giving you visibility into east-west flows and a way to micro-segment traffic to and from VMs (via their VMIs/virt-launcher pods and attached networks)

In OpenShift Virtualization, networking is Kubernetes-managed (pod network, Multus, SR-IOV). Using ACS to visualize and codify least-privilege policies helps reduce attack surface around live-migration networks, bridge networks, and service exposure paths that VMs depend on.

Why it matters for VMs: Even if a guest OS is patched, over-permissive cluster networking can expose the VM's services. ACS turns observed flows into policy candidates, making it far easier to lock down communication to only what VM applications truly need.

Compliance, reporting, and audit readiness across mixed VMs and containers

Built-in compliance profiles and one-click reports (e.g., CIS, NIST, PCI, HIPAA) help document control effectiveness across your OpenShift clusters that host VMs, useful for auditors who want evidence that the platform and its security policies meet standards.

Because OpenShift Virtualization treats VMs as first-class Kubernetes resources, these compliance checks cover the same cluster control plane that schedules and isolates your VM workloads.

Why it matters for VMs: Many regulatory controls focus on platform hardening and network segmentation. ACS lets you prove these are enforced uniformly for both containers and VMs.

Software supply chain security for images that touch your VMs

ACS integrates with registries and CI/CD to scan images, verify signatures, and enforce deployment policies. While the VM guest disks aren't container images, your virt-launcher, operators, and sidecars are. Securing these artifacts reduces the risk of hosting-surface compromise that could impact VMs.

You can embed scans and policy checks in OpenShift Pipelines/Tekton so only compliant images (including those used by virtualization operators or supporting services) reach production.

Why it matters for VMs: A compromised operator or base image on a virtualization node can endanger guest VMs. ACS helps ensure only verified, low-risk images power platform components around your VMs.

Unified visibility and incident response

Centralized dashboards aggregate risks, violations, CVEs, and network findings across clusters, including those that host OpenShift Virtualization. Security teams get a single place to triage and respond, instead of stitching together hypervisor-only and Kubernetes-only views.

ACS works across multiple clusters (hub-and-spoke), aligning with common OpenShift footprints where one or more clusters run virtualization at scale.

Application backup and restore OpenShift API for Data Protection

Red Hat OpenShift API for Data Protection (OADP) product safeguards customer applications on OpenShift Container Platform. It offers comprehensive disaster recovery protection, covering OpenShift Container Platform applications, application-related cluster resources, persistent volumes, and internal images. OADP is also capable of backing up both containerized applications and virtual machines (VMs). However, OADP **does not serve** as a disaster recovery solution for etcd or Red Hat OpenShift Operators.

Red Hat OpenShift API for Data Protection (OADP) is the Red Hat OpenShift tool used to back up and restore Kubernetes/Red Hat OpenShift cluster resources and persistent volumes. It is based on the Velero project.

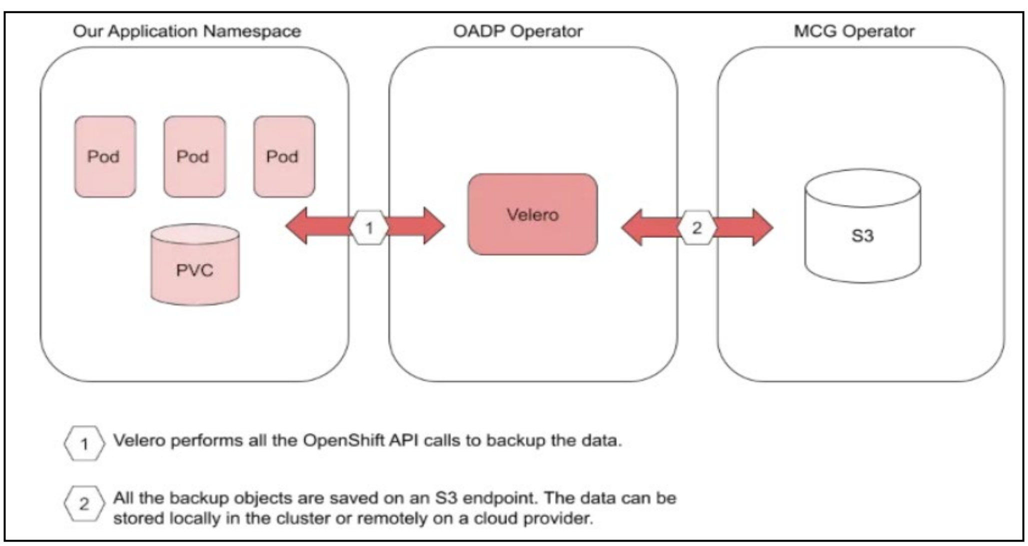

Figure 18. Red Hat OpenShift API for Data Protection

Business continuity with Disaster Recovery strategies for Red Hat OpenShift Container Platform 4.21 using Red Hat OpenShift Data Foundation Storage

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 4.20 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 19 shows the comparison of the Red Hat OpenShift disaster recovery strategies using RTO and RPO objectives.

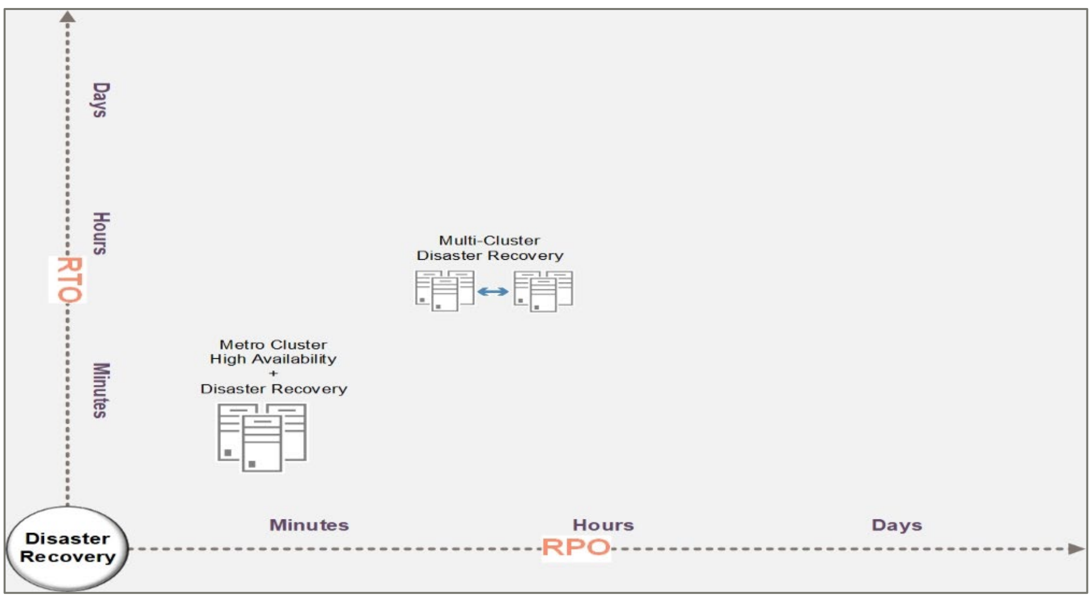

Figure 19. Red Hat OpenShift Disaster recovery strategies comparison using RTO and RPO objectives.

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

— 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 5-10 ms RTT.

— The solution ensures pods and nodes get scheduled across zones during deployment.

— ODF External maintains consistent mirror copies across AZs resulting in less or no data loss.

— Stretched solution cluster provides automatic and non-disruptive recovery for apps across AZs.

— An application with a consensus protocol that allows it to determine which instances of the cluster are active and healthy.

Figure 20 shows an overview of the Red Hat OpenShift Metro Cluster design.

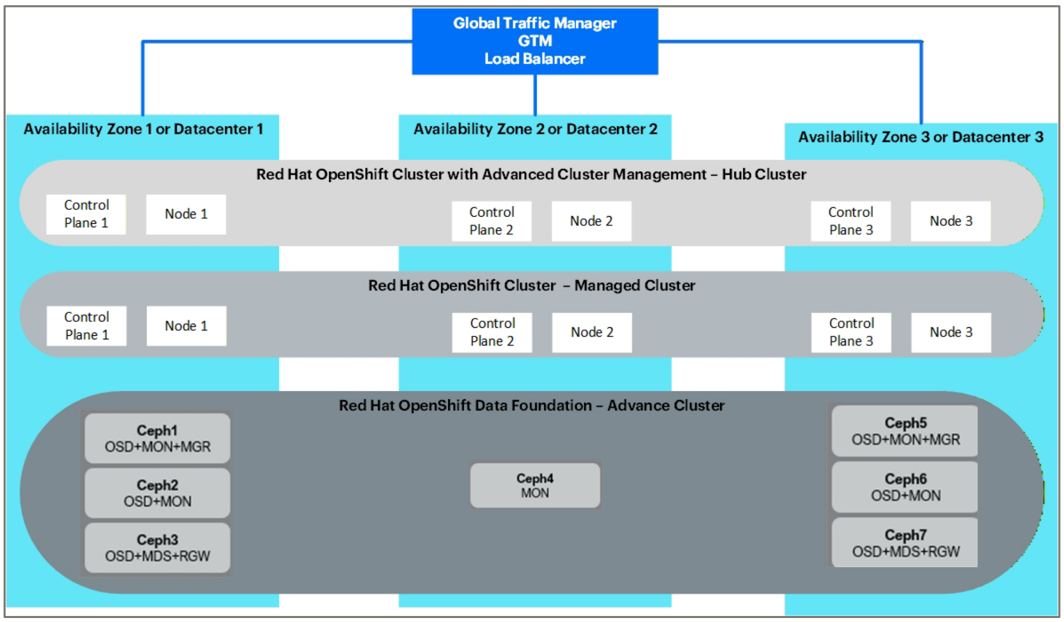

Figure 20. Red Hat OpenShift Metro Cluster design overview

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 computer clusters. 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 be reorganized to manage the loss of a cluster member.

Figure 21 shows the Red Hat OpenShift Multi-cluster disaster recovery approach.

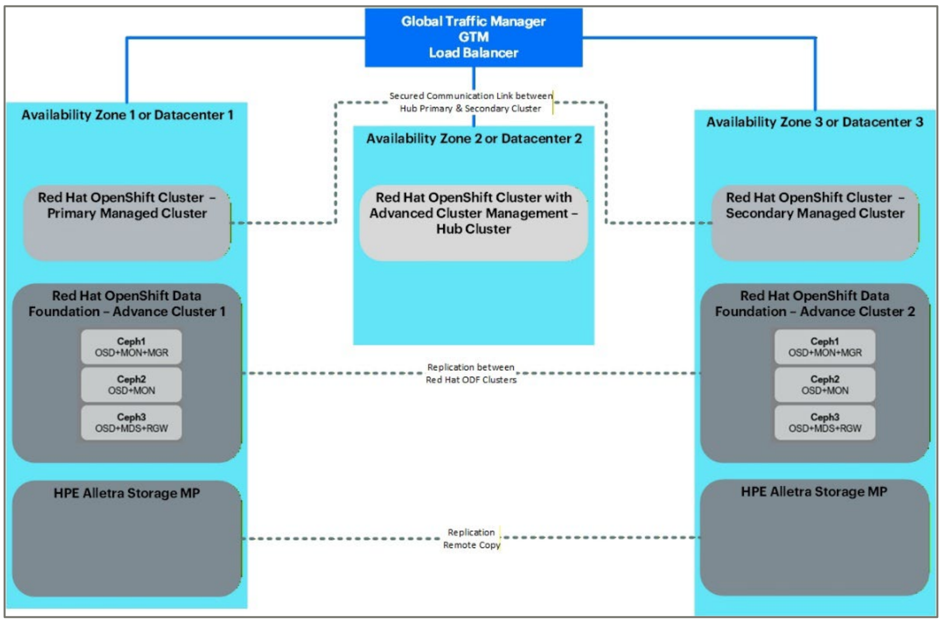

Figure 21. Red Hat OpenShift Multi-cluster disaster recovery approach

Regional-DR is composed of Red Hat Advanced Cluster Management for Kubernetes and Red Hat OpenShift Data Foundation components to provide application and data mobility across Red Hat OpenShift Container Platform clusters. OpenShift DR. Red Hat OpenShift Disaster Recovery (DR) is a set of orchestrators to configure and manage stateful applications across a set of peers OpenShift clusters which are managed using RHACM and provide cloud-native interfaces to orchestrate the life cycle of an application's state on Persistent Volumes. These include:

— Protecting an application and its state relationship across Red Hat OpenShift clusters

— Failing over an application and its state to a peer cluster

— Relocate an application and its state to the previously deployed cluster

Red Hat OpenShift DR is split into three components:

— ODF Multi-cluster Orchestrator: Installed on the multi-cluster control plane (RHACM Hub), it orchestrates configuration and peering of Red Hat OpenShift Data Foundation clusters for Metro and Regional DR relationships

— Red Hat OpenShift DR Hub Operator: Automatically installed as part of ODF Multi-cluster Orchestrator installation on the hub cluster to orchestrate failover or relocation of DR enabled applications.

— Red Hat OpenShift DR Cluster Operator: Automatically installed on each managed cluster that is part of a Metro and Regional DR relationship to manage the lifecycle of all PVCs of an application.

For more information, see Disaster Recovery Strategies for Red Hat OpenShift Data Foundation.

Summary

This reference architecture demonstrates how organizations can confidently standardize on Red Hat OpenShift Container Platform 4.x running on HPE ProLiant Compute Gen12 infrastructure as a unified foundation for modern hybrid workloads spanning both containers and virtual machines. By leveraging validated hardware platforms, including HPE ProLiant Compute DL Gen12 servers, alongside integrated storage, networking, and lifecycle management components, the solution delivers a scalable, secure, and enterprise-ready environment for cloud-native and traditional applications. Through its detailed treatment of physical configuration, the software stack, Red Hat OpenShift Virtualization (KubeVirt), capacity planning, and operational design patterns, this reference architecture connects strategy to execution and equips teams with the guidance needed to deploy Red Hat OpenShift Container Platform with confidence and predictability.

Security, compliance, and resilience are embedded throughout the architecture, reflecting their importance as foundational design principles rather than optional add-ons. Integrated controls for platform hardening, workload isolation, backup and restore, metro-cluster high availability, and multi-cluster disaster recovery ensure that organizations can meet regulatory requirements while minimizing operational risk. Combined with capabilities for VM migration, workload portability, and unified management of mixed VM and container estates, the architecture accelerates modernization journeys without compromising stability. Ultimately, this document serves as a strategic blueprint for building a resilient, future-ready application platform on HPE infrastructure, enabling enterprises to reduce design uncertainty, accelerate time to value, and establish a standardized, secure foundation for digital transformation powered by Red Hat OpenShift Container Platform 4.x.

Appendix A: Red Hat OpenShift solution backup restore using Veeam Kasten

Modernizing and securing applications is more critical now than ever. Red Hat, in collaboration with Veeam, offers a comprehensive solution that leverages the power of Kubernetes and OpenShift, ensuring that your applications are both modernized and protected. By combining Red Hat OpenShift with Veeam Kasten, businesses can gain a powerful solution to address the challenges listed above. Together, these solutions offer:

— Business continuity: Assurance that your applications and their data remain available and resilient.

— Data governance: Control your data even across various environments.

— Application and data protection: Protection against data loss and cyber threats.

Red Hat OpenShift, enhanced by Veeam Kasten, supports application modernization through backup and disaster recovery (DR) solutions, application mobility, migration capabilities and ransomware protection. This includes integration with various cloud platforms such as AWS, Microsoft Azure, or multi-cloud environments, helping ensure that your applications can be easily managed and scaled while maintaining high availability and security. Red Hat OpenShift and Veeam Kasten work together to manage and protect both VMs and cloud-native workloads. Key features include:

— Cloud-native integration: Fully integrated with cloud-native technologies.

— Compliance: FIPS 140-3 compliant for security.

— Encryption: Integrated end-to-end encryption ensures data is transferred and stored securely.

— Immutability: Support for immutable repositories to ensure backup data can't be held ransom or accidentally deleted.

— Advanced management: Leveraging Advanced Cluster Manager (ACM) and CephRBD integration for efficient backup and snapshot management.

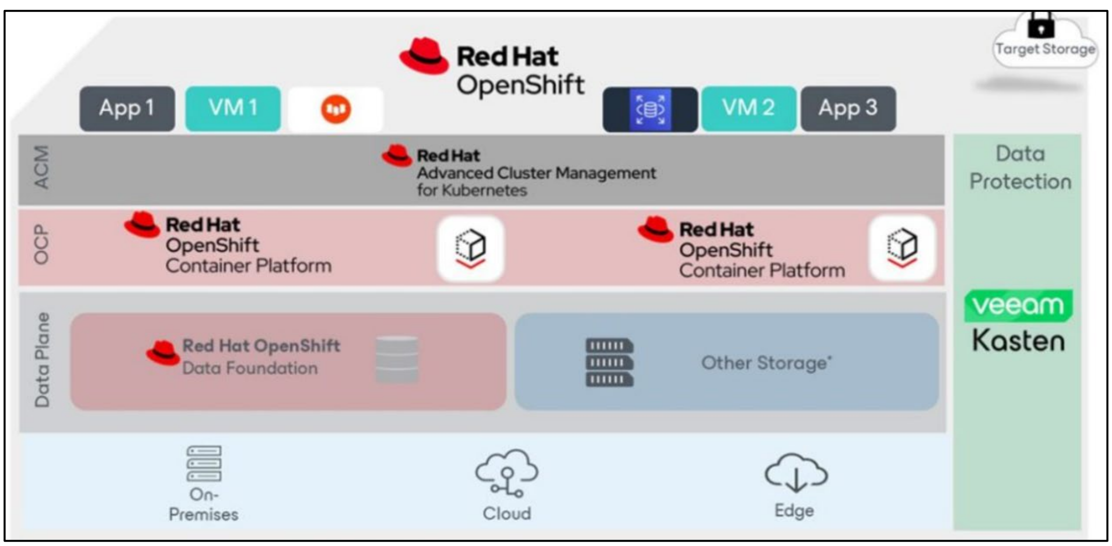

Figure A1. Red Hat OpenShift Solution Backup Restore using Veeam Kasten

Red Hat and Veeam have partnered to deliver a robust, cloud-native solution that empowers enterprises to modernize and secure applications across hybrid and multi-cloud environments. By integrating Red Hat OpenShift with Veeam Kasten, organizations gain comprehensive capabilities for business continuity, data governance, and protection against cyber threats. This joint platform supports backup, disaster recovery, application mobility, and ransomware resilience for both virtual machines and containerized workloads. Key features include FIPS 140-3 compliance, end-to-end encryption, immutable backups, and advanced management via OpenShift ACM and CephRBD—ensuring scalable, secure, and compliant application modernization.

Appendix B: Red Hat OpenShift disconnected installation

An air-gapped environment is a computer network or system that is isolated from other networks, such as the internet, to prevent unauthorized access and cyber threats. The name "air-gapped" refers to the air that separates the system from potential threats. Air-gapped environments are used in situations where there is a high risk of cyber threats, such as:

Military and government: Computer networks and systems for the military and government.

Financial: Computer systems for financial institutions, such as stock exchanges.

Industrial control systems: Systems for industrial control, such as SCADA in oil and gas fields.

Life-critical systems: Systems that are critical for life, such as nuclear power plant controls, air traffic control systems, and computerized medical equipment.

Deploying Red Hat OpenShift in disconnected or air-gapped environment overview

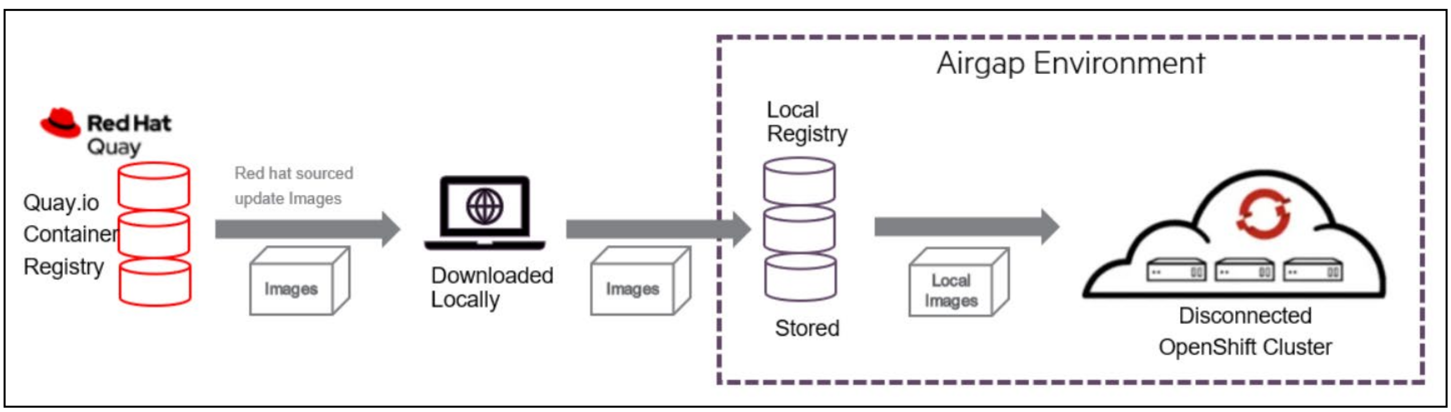

Figure B1. Red Hat OpenShift disconnected installation overview

The reference architecture describes the UPI or user provisioned infrastructure approach to deploying Red Hat OpenShift - this enables us to deploy Red Hat OpenShift Cluster with maximum customized, or flexibility required for Optimal solution on HPE Hardware.

To the setup the infrastructure we need the following:

— Repository server

- For RHEL OS images and RPM files
- For support tool installation files

To set up Red Hat OpenShift nodes:

— Mirror Registry server

- To store OpenShift base images
- OpenShift Operator Images

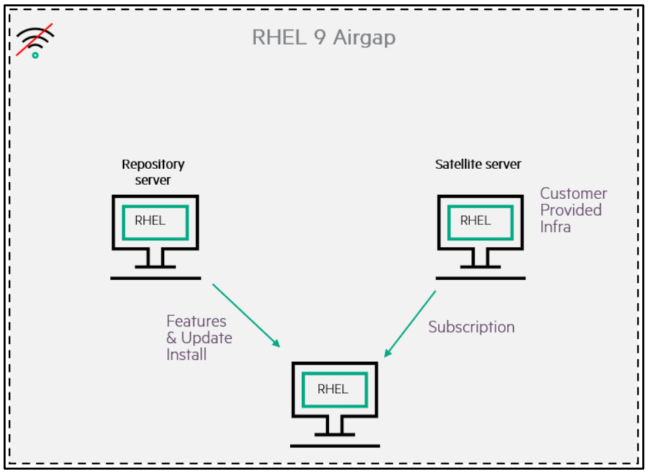

Figure B2. Red Hat OpenShift disconnected installation components

Before you install a cluster on infrastructure that you provision in a restricted network, you must mirror the required container images into that environment. To mirror container images, you must have a registry for mirroring. Both these roles can be deployed on a dedicated HPE ProLiant Compute GL325 Gen12 server running Red Hat Enterprise Linux (RHEL) 9 prior to installing Red Hat OpenShift Container Platform components or KVM Infra Node as mentioned in this reference architecture.

For more details on Red Hat OpenShift disconnected Installation, refer to the deployment guide.

Appendix C: HPE Morpheus Enterprise integration

The Red Hat OpenShift integration adds a new Cluster type, which can be added to HPE Morpheus Enterprise from the Clusters list page (**Infrastructure > Clusters**). Once a pre-existing OpenShift cluster is integrated, HPE Morpheus Enterprise can onboard any currently running virtualized workloads along with the cluster hosts, network objects, storage objects, and virtual images needed to provision additional virtualized workloads. In addition to a new Cluster type, the OpenShift integration adds a new Instance Type to the provisioning wizard, which is used to provision new Instances to any integrated Clusters.

The OpenShift integration is developed as a standalone plugin for HPE Morpheus Enterprise and is not part of the product by default. Information on accessing the plugin and adding it to the HPE Morpheus Enterprise appliance is included in a subsequent.

The Red Hat OpenShift plugin enables the integration of OpenShift cluster with HPE Morpheus Enterprise, allowing users to manage and provision virtualized workloads on OpenShift from the Morpheus UI. With this integration, users can onboard existing clusters, expose cluster networking, storage, and images to Morpheus, and use Morpheus provisioning and automation to deploy and operate OpenShift workloads.

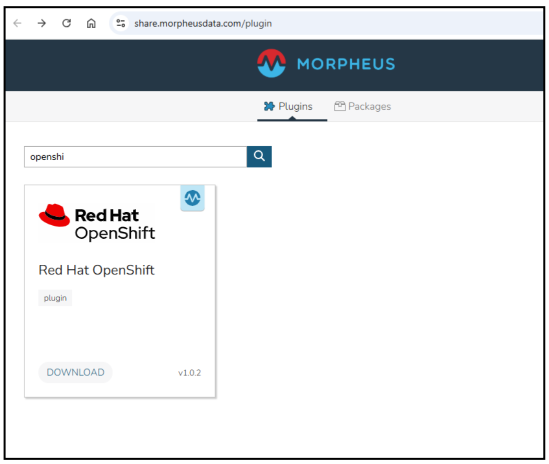

Figure C1. HPE Morpheus Enterprise Red Hat OpenShift plugin

Key features:

— Onboard existing OpenShift clusters into HPE Morpheus Enterprise quickly and securely.

— Import and expose OpenShift network objects, storage objects, and virtual images so they are available in the Morpheus provisioning wizard.

— Adds a dedicated Instance Type in the Morpheus provisioning wizard for simple, guided provisioning of OpenShift virtualized workloads.

— Leverages the HPE Morpheus automation engine for one-time provisioning and ongoing (day-2) operations across the full instance lifecycle, including scaling, updates, and decommissioning.

Note

The OpenShift integration with HPE Morpheus Enterprise is currently limited to support of virtualized workloads only. There is not currently support for container-based workloads.

Prerequisites

— HPE Morpheus Enterprise appliance running version 8.0.13 or higher

— At least one pre-existing Red Hat OpenShift 4.x cluster, which can communicate back to the HPE Morpheus Enterprise appliance on the HTTPS port

— Access to a full administrator account for each cluster which will be integrated

— Ability for the HPE Morpheus Enterprise appliance to communicate with the OpenShift API on port 6443

— Ability for the HPE Morpheus Enterprise appliance to communicate with each provisioned VM on port 22 for SSH and on port 3389 for Windows RDP

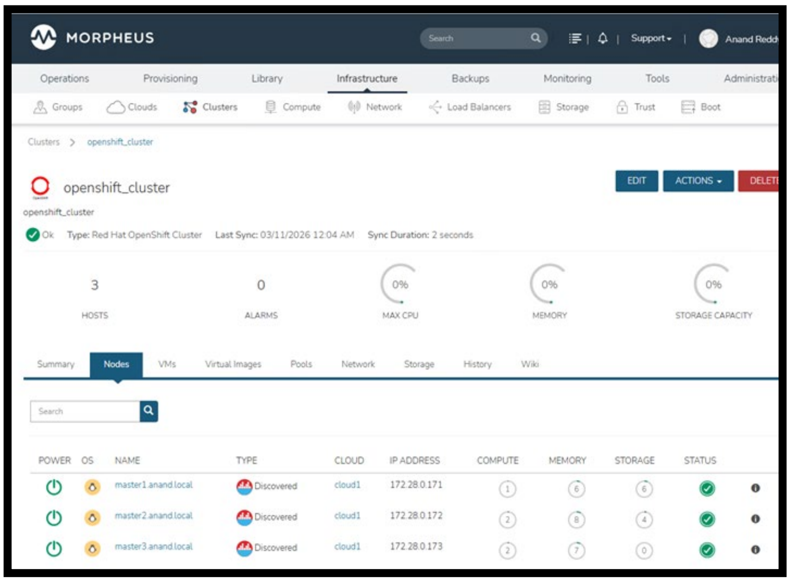

Figure C2. HPE Morpheus Enterprise managing Red Hat OpenShift Cluster

After onboarding the cluster into HPE Morpheus Enterprise, a new Cluster object is added to the **Clusters** list page. Clicking into the Cluster reveals the Cluster detail page. For more information on this feature within HPE Morpheus Enterprise, refer to the deployment guide for this reference architecture or visit the HPE Morpheus Enterprise documentation at [OpenShift Clusters | HPE Morpheus Enterprise Software Documentation v8.1.0](#).

Appendix D: Bill of materials (BOMs)

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 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 D1. Bill of materials

Component	Qty	Description
Rack and Networking Components		
Product #	Qty	Product Description
P9K40A	1	HPE 42U 600mmx1200mm G2 Enterprise Shock Rack
P9K40A 001	1	HPE Factory Express Base Racking Service
HA454A1-000	1	HPE FE Solution Package 4 SVC
R9F63A	1	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle
R9F63A 0D1	1	Factory Integrated
R9F63A B2B	1	HPE Aruba Networking CX 6300M 48G Power-to-Port Airflow 2 Fans 1 Power Supply Unit Bundle PDU
HA454A1-021	1	HPE FE Strg and Ntwking Pkg 4 SVC
R9G06A	1	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable
R9G06A B01	1	HPE Aruba Networking 50G SFP56 to SFP56 0.65m Direct Attach Copper Cable
R9F61A	1	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit
R9F61A B2B	1	HPE Aruba Networking CX 6300M 12VDC 250W 100-240VAC Power-to-Port Airflow Power Supply Unit PDU
R9F61A 0D1	1	Factory Integrated
R9F57A	1	HPE Aruba Networking 1U Universal 4-post Rack Mount Kit
R9F57A 0D1	1	Factory Integrated
R9F59A	2	HPE Aruba Networking 4-post Rack Kit
R9F59A 0D1	2	Factory Integrated
R9F67A	2	HPE Aruba Networking CX 8325-32C Power-to-Port Airflow 6 Fans 2 Power Supply Units Bundle
R9F67A 0D1	2	Factory Integrated

Component	Qty	Description
R9F67A B2B	2	HPE Aruba Networking CX 8325-32C Power-to-Port Airflow 6 Fans 2 Power Supply Units Bundle PDU
HA454A1-021	2	HPE FE Strg and Ntwking Pkg 4 SVC
R9F77A	2	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable
R9F77A B01	2	HPE Aruba Networking 100G QSFP28 to QSFP28 1m Direct Attach Copper Cable
Solution Head Nodes		
P71437-B21	3	HPE ProLiant Compute DL320 Gen12 SFF Configure-to-order Server
P73834-B21	3	Intel Xeon 6767P 2.4GHz 64-core 350W Processor for HPE
P73834-B21 0D1	3	Factory Integrated
P69727-B21	48	HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
P69727-B21 0D1	48	Factory Integrated
P75084-B21	3	HPE ProLiant Compute DL3XX Gen12 1U 8SFF x1 Tri-Mode U.3 L-Shaped Backplane Kit
P75084-B21 0D1	3	Factory Integrated
P40500-B21	24	HPE 3.84TB SATA 6G Read Intensive SFF BC Multi-Vendor SSD
P40500-B21 0D1	24	Factory Integrated
P71430-B21	3	HPE ProLiant Compute DL3X0 Gen12 x16 PCIe Primary Riser Kit
P71430-B21 0D1	3	Factory Integrated
P51178-B21	3	Broadcom BCM5719 Ethernet 1Gb 4-port BASE-T Adapter for HPE
P51178-B21 0D1	3	Factory Integrated
P01366-B21	3	HPE 96W Smart Storage Lithium-ion Battery with 145mm Cable Kit
P01366-B21 0D1	3	Factory Integrated
P58335-B21	3	HPE MR408i-o Gen11 x8 Lanes 4GB Cache OCP SPDM Storage Controller
P58335-B21 0D1	3	Factory Integrated
P76603-B21	21	HPE ProLiant Compute DL320 Gen12 1U Closed-loop Liquid Cooling FIO Fan Kit
P67240-B21	6	HPE 1000W M-CRPS Titanium Hot Plug Power Supply Kit
P67240-B21 0D1	6	Factory Integrated
P78131-B21	6	HPE C13 - GB1002 250V 10Amp 1.83m FIO Power Cord
BD505A	3	HPE iLO Advanced 1-server License with 3yr Support on iLO Licensed Features
BD505A 0D1	3	Factory Integrated
P75090-B21	3	HPE ProLiant Compute DL320 Gen12 8SFF x1 U.3 OROC Cable Kit
P75090-B21 0D1	3	Factory Integrated
P76981-B21	3	HPE ProLiant Compute iLO7 DC-SCM FIO Management Module Kit
P52349-B21	3	HPE Easy Install Rail 1 Kit
P52349-B21 0D1	3	Factory Integrated
P71433-B21	3	HPE ProLiant Compute DL3XX Gen12 NS204i-u Rear Cable Kit
P71433-B21 0D1	3	Factory Integrated
P73325-B21	3	HPE ProLiant Compute Localization FIO Kit
P76605-B21	3	HPE ProLiant Compute DL320 Gen12 1U Closed-loop Liquid Cooling FIO Heat Sink Kit
P78279-B21	3	HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device
P78279-B21 0D1	3	Factory Integrated
P79558-B21	3	HPE ProLiant Compute 25C System Inlet Ambient Operating Temperature Configuration Tracking
P79633-B21	3	HPE ProLiant Compute DAC ACC Networking Cable Operating Configuration Tracking
HA113A1	1	HPE Installation Service
HA113A1 5A6	3	HPE ProLiant DL/ML Install SVC
HU4A6A3	1	HPE 3Y Tech Care Essential Service
HU4A6A300JB	3	HPE DL320 Gen12 Support

Component	Qty	Description
Storage Options Red Hat OpenShift Data Foundation nodes		
P73282-B21	3	HPE ProLiant Compute DL380 Gen12 SFF NC Configure-to-order Server
P73282-B21 0D1	3	Factory Integrated
P73282-B21 ABA	3	HPE ProLiant Compute DL380 Gen12 SFF NC CTO Svr
HA454A1-001	3	HPE FE ProLiant Svr Pkg 4 SVC
P74571-B21	6	Intel Xeon 6530P 2.3GHz 32-core 225W Processor for HPE
P74571-B21 0D1	6	Factory Integrated
P69727-B21	96	HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
P69727-B21 0D1	96	Factory Integrated
P75740-B21	3	HPE ProLiant Compute DL3XX Gen12 8SFF x1 U.3 Tri-Mode Drive Cage Kit
P75740-B21 0D1	3	Factory Integrated
P63871-B21	6	HPE 1.6TB SAS Mixed Use SFF BC Self-encrypting FIPS 140-2 PM7 SSD
P63871-B21 0D1	6	Factory Integrated
P75741-B21	6	HPE ProLiant Compute DL3XX Gen12 8SFF x4 U.3 Tri-Mode Drive Cage Kit
P75741-B21 0D1	6	Factory Integrated
P50230-B21	24	HPE 3.2TB NVMe Gen4 High Performance Mixed Use SFF BC U.3 PM1735a SSD
P50230-B21 0D1	24	Factory Integrated
P48802-B21	3	HPE ProLiant DL380 Gen11 2U x8/x16/x8 Secondary Riser Kit
P48802-B21 0D1	3	Factory Integrated
P48803-B21	3	HPE ProLiant DL380 Gen11 2U x16/x16/x16 Primary Riser Kit
P48803-B21 0D1	3	Factory Integrated
P47785-B21	3	HPE MR216i-p Gen11 x16 Lanes without Cache PCI SPDM Plug-in Storage Controller
P47785-B21 0D1	3	Factory Integrated
P42041-B21	6	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P42041-B21 0D1	6	Factory Integrated
P44712-B21	6	HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kit
P44712-B21 0D1	6	Factory Integrated
P78145-B21	6	HPE C13 - C14 250V 10Amp 2m FIO Power Cord
E5Y43A	3	HPE OneView for ProLiant Compute DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU
P72205-B21	3	HPE ProLiant Compute DL3XX/ML350 Gen12 CPU2 to Rear OCP Slot B x8 Cable Kit
P72205-B21 0D1	3	Factory Integrated
P76453-B21	3	HPE ProLiant Compute DL380 Gen12 8SFF SFF x4 UMB PCIe Box 1/2 Cable Kit
P76453-B21 0D1	3	Factory Integrated

Component	Qty	Description
P76461-B21	3	HPE ProLiant Compute DL380 Gen12 8SFF x4 Direct Attach Box 1 for 2 Processors Cable Kit
P76461-B21 0D1	3	Factory Integrated
P77489-B21	3	HPE ProLiant Compute DL380 Gen12 8SFF x4 Direct Attach Box 2/3 for 1P Cable Kit
P77489-B21 0D1	3	Factory Integrated
P48820-B21	3	HPE ProLiant DL380/DL560 Gen11 2U High Performance Fan Kit
P48820-B21 0D1	3	Factory Integrated
P48922-B21	3	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21 0D1	3	Factory Integrated
P08040-B21	3	HPE iLO Common Password FIO Setting
P42104-B21	3	HPE ProLiant Platform Certificate and IDevID iLO FIO Setting
P52341-B21	3	HPE ProLiant DL3XX Gen11 Easy Install Rail 3 Kit
P52341-B21 0D1	3	Factory Integrated
P73325-B21	3	HPE ProLiant Compute Localization FIO Kit
P74792-B21	6	HPE ProLiant Compute DL380 Gen12 Performance Heat Sink Kit
P74792-B21 0D1	6	Factory Integrated
P77931-B21	3	HPE ProLiant Compute DL380 Gen12 16SFF x4 Direct Attach Balanced FIO Bundle Kit
P79552-B21	3	HPE ProLiant Compute 30C System Inlet Ambient Operating Temperature Configuration Tracking
Storage Options HPE Alletra Storage MP B10000(Switchless)		
S0B84A	1	HPE Alletra Storage MP B10100 Base Configuration
S0B84A 0D1	1	Factory Integrated
S3Q02A	1	HPE Alletra Storage ArcusOS per TB 5-year LTU
S3Q02A 0D1	1	Factory Integrated
581817-B21	1	HPE Configurator Defined Build Instruction Option
R7C75A	1	HPE Alletra Storage MP 10000 2U Chassis
R7C75A 0D1	1	Factory Integrated
S0S39A	2	HPE Alletra Storage MP B10140 100GbE Controller Node
S0S39A 0D1	2	Factory Integrated
R7C82A	4	HPE Alletra Storage MP 10/25GbE 4-port Host Bus Adapter
R7C82A 0D1	4	Factory Integrated
Q2P64B	16	HPE 25Gb SFP28 Short Wave Extended Temperature 1-pack Pull Tab Optical Transceiver
Q2P64B 0D1	16	Factory Integrated
S2A68A	2	HPE Alletra Storage MP 100GbE 2-port OCP Host Bus Adapter
S2A68A 0D1	2	Factory Integrated

Component	Qty	Description
R7C76A	2	HPE Alletra Storage MP C14 1600W AC Power Supply
R7C76A 0D1	2	Factory Integrated
R9H67A	18	HPE Alletra Storage MP 3.84TB NVMe SFF Self-encrypting SSD
R9H67A 0D1	18	Factory Integrated
R9R52A	1	HPE C13 - C14 250V 10Amp Black 1.4m WW Power Cord
R9R52A 0D1	1	Factory Integrated
R9S00A	1	HPE C13 - C14 250V 10Amp Gray 1.4m WW Power Cord
R9S00A 0D1	1	Factory Integrated
S0A98A	1	HPE Storage Data Encryption LTU
S0A98A 0D1	1	Factory Integrated
R8M59A	2	HPE 100Gb QSFP28 to QSFP28 0.5m Direct Attach Copper Cable
R8M59A 0D1	2	Factory Integrated
Workload Option General Purpose		
P72176-B21	3	HPE ProLiant Compute DL360 Gen12 10SFF/20EDSFF Hybrid NC Configure-to-order Server
P72176-B21 0D1	3	Factory Integrated
P72176-B21 ABA	3	HPE DL360 G12 10SFF/20EDSFF Hyb CTO Svr
HA454A1-001	3	HPE FE ProLiant Svr Pkg 4 SVC
P74571-B21	6	Intel Xeon 6530P 2.3GHz 32-core 225W Processor for HPE
P74571-B21 0D1	6	Factory Integrated
P69727-B21	48	HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
P69727-B21 0D1	48	Factory Integrated
P72223-B21	3	HPE ProLiant Compute DL3XX Gen12 1U 2SFF x4 Tri-Mode U.3 Stacking Backplane Kit
P72223-B21 0D1	3	Factory Integrated
P42044-B21	3	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE
P25960-B21 0D1	3	Factory Integrated
P42041-B21	3	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P48908-B21	3	HPE ProLiant DL3X0 Gen11 1U High Performance Fan Kit
P48908-B21 0D1	3	Factory Integrated
P03178-B21	3	HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
P03178-B21 0D1	3	Factory Integrated
P78145-B21	3	HPE C13 - C14 250V 10Amp 2m FIO Power Cord
E5Y43A	3	HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU

Component	Qty	Description
P72211-B21	3	HPE ProLiant Compute DL360 Gen12 2SFF Stacking x4 CPU2 Box1/CPU1 Box1/2 Signal DA Cable Kit
P72211-B21 0D1	3	Factory Integrated
P48922-B21	21	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21 0D1	21	Factory Integrated
P07818-B21	3	HPE DDR4 DIMM Blank Kit
P07818-B21 0D1	6	Factory Integrated
P08040-B21	3	HPE iLO Common Password FIO Setting
P52343-B21	3	HPE Easy Install Rail 5 Kit
P52343-B21 0D1	3	Factory Integrated
P72209-B21	3	HPE ProLiant Compute DL360 Gen12 10SFF/20EDSFF Hybrid Backplane Power Cable Kit
P72209-B21 0D1	3	Factory Integrated
P73325-B21	3	HPE ProLiant Compute Localization FIO Kit
P74787-B21	6	HPE ProLiant Compute DL3XX Gen12 High Performance Heat Sink Kit
P74787-B21 0D1	3	Factory Integrated
P77198-B21	3	HPE ProLiant Compute DL3XX Gen12 1U NS204i-u Front Enablement Kit
P77198-B21 0D1	3	Factory Integrated
P79558-B21	3	HPE ProLiant Compute 25C System Inlet Ambient Operating Temperature Configuration Tracking
P79633-B21	3	HPE ProLiant Compute DAC ACC Networking Cable Operating Configuration Tracking
P81162-B21	3	HPE NS204i-u v2 960GB NVMe SED Hot Plug Boot Optimized Storage Device
P81162-B21 0D1	3	Factory Integrated
Workload Option Memory Optimized		
P72176-B21	3	HPE ProLiant Compute DL360 Gen12 10SFF/20EDSFF Hybrid NC Configure-to-order Server
P72176-B21 0D1	3	Factory Integrated
P72176-B21 ABA	3	HPE DL360 G12 10SFF/20EDSFF Hyb CTO Svr
HA454A1-001	3	HPE FE ProLiant Svr Pkg 4 SVC
P74576-B21	6	Intel Xeon 6737P 2.9GHz 32-core 270W Processor for HPE
P74576-B21 0D1	6	Factory Integrated
P69727-B21	48	HPE 32GB (1x32GB) Dual Rank x8 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
P69727-B21 0D1	48	Factory Integrated
P72223-B21	3	HPE ProLiant Compute DL3XX Gen12 1U 2SFF x4 Tri-Mode U.3 Stacking Backplane Kit
P72223-B21 0D1	3	Factory Integrated

Component	Qty	Description
P42044-B21	3	Mellanox MCX631102AS-ADAT Ethernet 10/25Gb 2-port SFP28 Adapter for HPE
P42041-B21	3	Mellanox MCX631432AS-ADAI Ethernet 10/25Gb 2-port SFP28 OCP3 Adapter for HPE
P25960-B21 0D1	3	Factory Integrated
P48908-B21	3	HPE ProLiant DL3X0 Gen11 1U High Performance Fan Kit
P48908-B21 0D1	3	Factory Integrated
P03178-B21	6	HPE 1000W Flex Slot Titanium Hot Plug Power Supply Kit
P03178-B21 0D1	6	Factory Integrated
P78145-B21	6	HPE C13 - C14 250V 10Amp 2m FIO Power Cord
E5Y43A	3	HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU
P72211-B21	3	HPE ProLiant Compute DL360 Gen12 2SFF Stacking x4 CPU2 Box1/CPU1 Box1/2 Signal DA Cable Kit
P72211-B21 0D1	3	Factory Integrated
P48922-B21	3	HPE ProLiant DL3XX Gen11 Intrusion Cable Kit
P48922-B21 0D1	3	Factory Integrated
P08040-B21	3	HPE iLO Common Password FIO Setting
P52343-B21	3	HPE Easy Install Rail 5 Kit
P52343-B21 0D1	3	Factory Integrated
P72209-B21	3	HPE ProLiant Compute DL360 Gen12 10SFF/20EDSFF Hybrid Backplane Power Cable Kit
P72209-B21 0D1	3	Factory Integrated
P73325-B21	3	HPE ProLiant Compute Localization FIO Kit
P74787-B21	6	HPE ProLiant DL3XX Gen12 High Performance Heat Sink Kit
P74787-B21 0D1	6	Factory Integrated
P77198-B21	3	HPE ProLiant Compute DL3XX Gen12 1U NS204i-u Front Enablement Kit
P77198-B21 0D1	3	Factory Integrated
P79558-B21	3	HPE ProLiant Compute 25C System Inlet Ambient Operating Temperature Configuration Tracking
P79633-B21	3	HPE ProLiant Compute DAC ACC Networking Cable Operating Configuration Tracking
P81162-B21	3	HPE NS204i-u v2 960GB NVMe SED Hot Plug Boot Optimized Storage Device
P81162-B21 0D1	3	Factory Integrated
Workload Option GPU Optimized		
P76706-B21 ABA	3	HPE DL380a Gen12 8DW/16SW CTO Svr
HA454A1-024	3	HPE Factory Express Level 4 SVC
P73837-B21	6	Intel Xeon 6787P 2.0GHz 86-core 350W Processor for HPE
P73837-B21 0D1	3	Factory Integrated

Component	Qty	Description
P69729-B21	96	HPE 96GB (1x96GB) Dual Rank x4 DDR5-6400 CAS-52-52-52 EC8 Registered Smart Memory Kit
P69729-B21 0D1	6	Factory Integrated
P74712-B21	3	HPE ProLiant Compute DL380a Gen12 4EDSFF FIO Drive Cage Kit
P61191-B21	8	HPE 3.2TB NVMe Gen5 High Performance Mixed Use E3S EC1 EDSFF SPDM CM7 SSD
P61191-B21 0D1	12	Factory Integrated
P74714-B21	6	HPE ProLiant Compute DL380a Gen12 Switchboard FIO Enablement Kit
S3U30C	24	NVIDIA H200 NVL 141GB PCIe Accelerator for HPE
S3U30C 0D1	24	Factory Integrated
P76927-B21	3	HPE ProLiant Compute DL380a Front NIC FIO Enablement Kit
P45641-B23	12	HPE InfiniBand NDR/Ethernet 400Gb 1-port OSFP PCIe5 x16 MCX75310AAS-NEAT Adapter
P45641-B23 0D1	12	Factory Integrated
P65333-B21	1	HPE InfiniBand NDR200/Ethernet 200GbE 2-port QSFP112 PCIe5 x16 MCX755106AC-HEAT Adapter
P65333-B21 0D1	1	Factory Integrated
P67252-B21	24	HPE 2400W M-CRPS Titanium Hot Plug Power Supply Kit
P67252-B21 0D1	24	Factory Integrated
P78384-B21	24	HPE C19 - C20 250V 16Amp 2.5m FIO Power Cord
E5Y43A	3	HPE OneView for ProLiant DL Server including 3yr 24x7 Support FIO Bundle Physical 1-server LTU
P74716-B21	3	HPE ProLiant Compute DL380a Gen12 4EDSFF Direct Attach Cable for NVIDIA
P74716-B21 0D1	3	Factory Integrated
P83526-B21	12	HPE ProLiant Compute DL380a Gen12 GPU 16-pin Cable Kit
P83526-B21 0D1	12	Factory Integrated
P08040-B21	3	HPE iLO Common Password FIO Setting
P69770-B21	3	HPE ProLiant Compute DL380a Gen12 Ball Bearing Rail Kit
P69770-B21 0D1	3	Factory Integrated
P73325-B21	3	HPE ProLiant Compute Localization FIO Kit
P75011-B21	3	HPE ProLiant Compute DL380a Gen12 8 Double Wide NIC FIO Configuration
P75284-B21	3	HPE DL380a Gen12 NS204i-u Front Cage Kit
P75284-B21 0D1	3	Factory Integrated
P78279-B21	3	HPE NS204i-u v2 480GB NVMe Hot Plug Boot Optimized Storage Device
P78279-B21 0D1	3	Factory Integrated
S4A90C	6	NVIDIA 4-way NVLink Bridge for H200 NVL
S4A90C 0D1	6	Factory Integrated

Component	Qty	Description
HPE Support Services		
HA113A1	1	HPE Installation Service
HA113A1 5MW	3	HPE Aruba 6xxxN8xxx Install Swt SVC
H6J85A	1	HPE Rack Hardware Kit
H6J85A 0D1	1	Factory Integrated
P9L16A	1	HPE G2 Rack 42U 1200mm Side Panel Kit
P9L16A 0D1	1	Factory Integrated
Q1H95A	1	HPE Storage 1U Rack Accessories Kit
Q1H95A 0D1	1	Factory Integrated
R7A11AAE	6	HPE Compute Ops Management Standard 3-year Upfront ProLiant SaaS
HU4A6A5	1	HPE 5Y Tech Care Essential Service
HU4A6A5 SVN	6	HPE One View w/Ilo Support
HU4A6A50C4U	3	HPE DL360 Gen12 Support
HU4A6A50C4V	3	HPE DL380 Gen12 Support
R9G32AAE	1	HPE Aruba Networking Fabric Composer Device Management Service Tier 3 Switch 3y Subscription E-STU
H7J34A5	1	HPE 5yr Foundational Care 24x7 Service
H7J34A5 ZSG	1	HPE Aruba Networking 6300M 48 SW Support
H7J34A5 ZND	2	HPE Aruba Networking 8325-32 SW Support
R9G27AAE	2	HPE Aruba Networking Fabric Composer Device Management Service Tier 4 Switch 3y Subscription E-STU
S3Q02AAE	69	HPE Alletra Storage MP B10000 per TB 5-year Software and Support SaaS
HU4A5A5	1	HPE 5Y Tech Care Critical with Comprehensive Defective Material Retention Service
HU4A5A5 R2M	6	HPE iLO Advanced Non Blade Support
HU4A5A5 SVN	3	HPE One View w/Ilo Support
HU4A5A500JB	6	HPE DL320 Gen12 Support
HU4A5A5008L	1	HPE Alletra Storage MP 2U Chassis Supp
HU4A5A5008Q	4	HPE Alletra Stg 10/25GbE 4-port HBA Supp
HU4A5A5008S	18	HPE Alletra Stg MP 3.84TB NVMe SSD Supp
HU4A5A5008W	1	HPE Alletra Stg MP Base Config Supp
HU4A5A5008Y	2	HPE Alletra STG MP 32c Blck CntrlNd Supp
HU4A5A5009A	2	HPE Alletra STG MP 100GbE 2p OCPHBA Supp
HU4A5A50C4W	3	HPE DL380a Gen12 Support
H38NHAS	1	HPE Alletra Storage MP B10000 SVC
H33XYA1	30	HPE Learn Credits for Storage SVC
H33XSA1	20	HPE Learn Credits for Compute IT SVC

Component	Qty	Description
HPE Morpheus Enterprise		
HA124A1#V38	1	Morpheus SW Production QuickStart
H46SBA1	1	Morpheus SW Production QuickStart
S6E66AAE	6	HPE Morpheus Software Enterprise per Socket (15 WLE) for PCE per 3-year 24x7 Support E-LTU

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 Compute DL320 Gen12
hpe.com/servers

Red Hat OpenShift Container Platform
access.redhat.com/documentation/en-us/openshift_container_platform/4.20 /

Red Hat OpenShift Container Storage
access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.20

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, use the Feedback button and icons (located at the bottom of an opened document) on the Hewlett Packard Enterprise Support Center portal (hpe.com/support/hpesc) to send any errors, suggestions, or comments. All documents information is captured by the process.

Visit HPE.com

a50015363enw

HEWLETT PACKARD ENTERPRISE

HPE.com