



# **HPE REFERENCE CONFIGURATION FOR RED HAT OPENSIFT CONTAINER PLATFORM ON HPE SYNERGY AND HPE STORAGE SYSTEMS**



# CONTENTS

Executive summary.....	3
Introduction.....	4
Solution overview.....	4
Physical bare metal configuration.....	4
Virtualized configuration.....	5
HPE Synergy Composable Infrastructure security controls.....	6
Solution layout.....	10
Solution components.....	12
Hardware.....	12
Software.....	14
Best practices and configuration guidance for the solution.....	17
Solution cabling.....	17
Networking.....	18
Storage.....	19
Server profile.....	22
Capacity and sizing.....	24
Deployment overview.....	24
Setup iPXE, TFTP, and DHCP for RHCOS.....	24
Configure a load balancer for Red Hat OpenShift 4 nodes.....	25
Configure DNS.....	25
Configure firewall ports.....	25
Start Red Hat OpenShift Container Platform 4 user provisioned infrastructure setup.....	25
Download Red Hat OpenShift Container Platform 4 software version.....	25
Configure the HPE Synergy Compute for iPXE boot.....	26
Worker node.....	26
HPE Container Storage Interface.....	27
Red Hat OpenShift Container Storage.....	27
Provision the local volume.....	28
Test and validate the setup.....	28
Physical worker node labeling in OpenShift.....	28
Securing and monitoring OpenShift with Sysdig SaaS.....	29
Post-validation deployment.....	29
Summary.....	30
Appendix A: Bill of materials.....	31
Resources and additional links.....	34



## EXECUTIVE SUMMARY

In today's digital world, organizations are under increasing pressure to deliver applications faster while reducing costs. As these applications grow more complex, this puts stress on IT infrastructure, IT teams, and processes. To remain competitive, organizations must adapt quickly, and developers need to be more effective, efficient, and agile. Container technology provides the right application platform to help organizations become more responsive and iterate across multiple IT environments and develop, deploy, and manage applications faster. Implementing a containerized environment across existing infrastructure is a complex undertaking that can require weeks or months to mobilize particularly for enterprises. To accelerate container application delivery, Hewlett Packard Enterprise and Red Hat® are collaborating to optimize Red Hat OpenShift® Container Platform 4 on the HPE platform, including HPE Synergy, the industry's first composable infrastructure.

Red Hat OpenShift Container Platform 4 on HPE Synergy provides an end-to-end fully integrated container solution that can be configured within hours after being assembled. This eliminates the complexities associated with implementing a container platform across an enterprise data center. This also provides the automation of hardware and software configuration to quickly provision and deploy a containerized environment at scale. Red Hat OpenShift Container Platform 4 provides organizations a reliable platform for deploying and scaling container-based applications. HPE Synergy provides the flexible infrastructure required to run the container platform to dynamically provision and scale applications, whether they run on virtual machines or containers, or hosted on-premises, in the cloud, or as a part of a hybrid cloud solution.

This Reference Configuration provides guidance for deploying, scaling, and managing the Red Hat OpenShift environment with the persistent volume on HPE Synergy Composable Infrastructure and storage.

The Reference Configuration describes how to:

- Leverage strength in composability of HPE Synergy along with the capabilities of Red Hat OpenShift Container Platform 4 to provide a simplified deployment experience
- Optimized layout of a Red Hat OpenShift Container Platform 4 configuration using bare metal hosts and virtualized hosts running RHEL Core OS
  - Automated process in the infrastructure deployment and post-deployment validation helped to save the time by 60-70%
  - Automated deployment process resulted in a reduction in the manual tasks by 85%
- Configure persistent storage for containers using the choice of HPE Storage systems (Nimble, 3PAR, and D3940)

The solution demonstrates the following benefits of utilizing HPE Synergy for Red Hat OpenShift Container Platform 4:

- Use an enterprise-grade composable storage systems such as HPE Nimble Storage, HPE 3PAR StoreServ, and HPE Synergy D3940 Storage Module for persistent storage of containers to enable speed, portability, and agility for traditional enterprise applications and data.
- HPE Composable Infrastructure solution provides a layered view of security controls. The objective of this layered security view is to ensure that customers become aware of the depth of security risk that infrastructure can face and to make them aware of the depth of defense that is built in the HPE Synergy Composable Infrastructure design.
- A business-driven container application provides basic data protection offered by HPE Storage systems such as HPE Nimble Storage, HPE 3PAR StoreServ, and HPE Synergy D3940 Storage Module.

**Target audience:** This document is intended for Chief Information Officers (CIOs), Chief Technology Officers (CTOs), data center managers, enterprise architects, and implementation personnel who wish to learn more about Red Hat OpenShift Container Platform 4 on HPE Synergy Composable Infrastructure. This document assumes that the target audience is familiar with HPE Synergy, Red Hat OpenShift Container Platform 4, solutions, and core networking.

**Document purpose:** The purpose of this document describes the benefits and technical details of deploying Red Hat OpenShift Container Platform 4 on HPE Synergy Composable Infrastructure, the implementation details, and processes. For more information on implementation details and processes, see Deployment guide at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy>.



## INTRODUCTION

This Reference Configuration describes the deployment of a highly available and secure Red Hat OpenShift Container Platform 4 on HPE Synergy Composable Infrastructure. This also includes the details of the configuration of the environment. This document demonstrates how Red Hat OpenShift Container Platform 4 can be deployed and configured to take advantage of the HPE Synergy Composable Infrastructure and leverage HPE Storage systems such as HPE Nimble, HPE 3PAR, and HPE Synergy D3940 Storage. The Red Hat OpenShift Container Platform 4 deployment on HPE Synergy Composable Infrastructure configuration consists of the following:

- Three (3) OpenShift Container Platform master nodes
- Three (3) OpenShift Container Platform worker nodes
- Six (6) HPE Synergy 480 Gen10 Compute Modules run Red Hat Enterprise Linux® CoreOS (RHCOS) to support the deployment
- Persistent volume leverages HPE Storage systems such as HPE Nimble, HPE 3PAR, and HPE Synergy D3940 Storage to provide persistent storage for containers and registry as well as data management

---

### NOTE

Due to the ephemeral nature of containers, protecting persistent data associated with the containers becomes a crucial task. In this Reference Configuration, the pods which require persistent data use persistent volume with disks from the HPE Storage systems.

---

## SOLUTION OVERVIEW

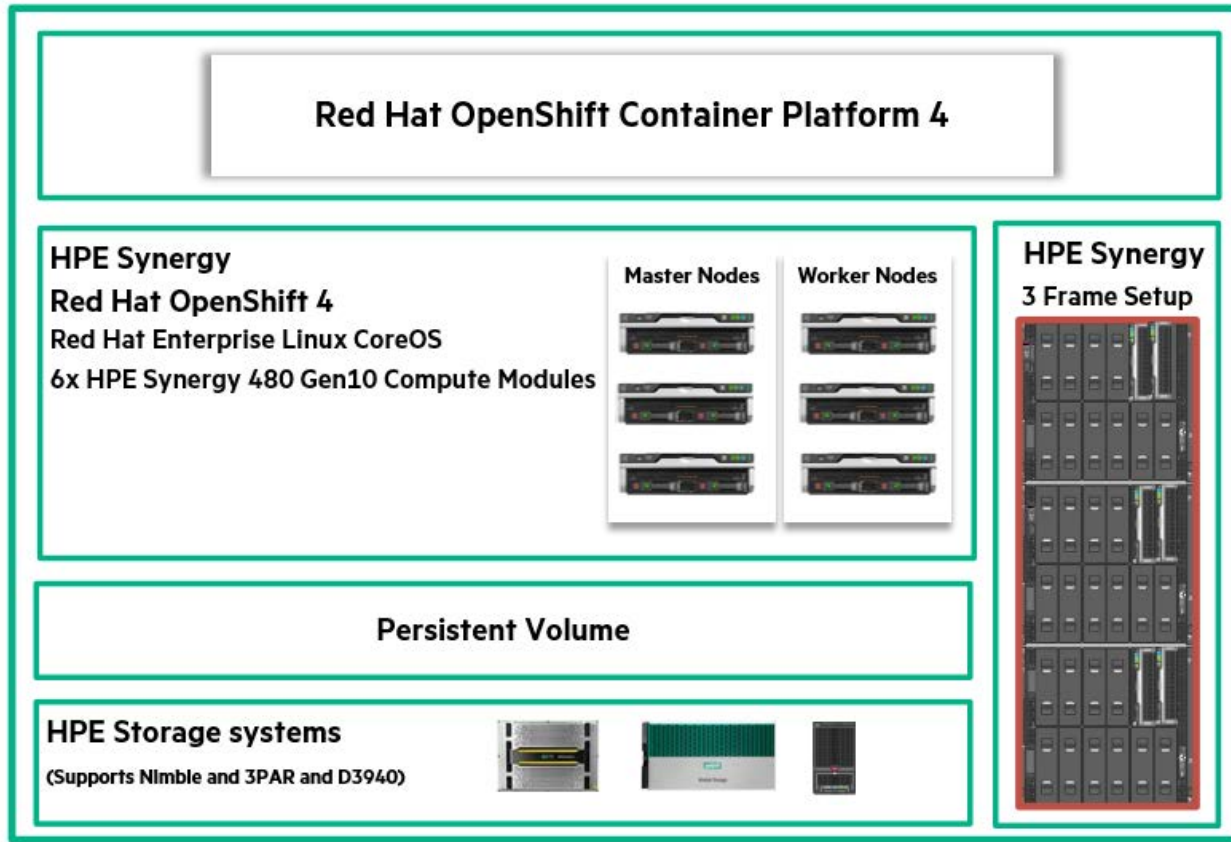
This provides an overview of Red Hat OpenShift Container Platform 4 on HPE Synergy using HPE Storage systems. This Reference Configuration describes the deployment of Red Hat OpenShift Container Platform 4 on both physical and virtual compute resources.

### Physical bare metal configuration

In a bare metal configuration, the master nodes are deployed in a highly available configuration running on three (3) HPE Synergy 480 Gen10 Compute Modules. Load balancing can be deployed as a virtual machine or as a physical appliance. Red Hat OpenShift worker nodes are deployed on the bare metal with three (3) HPE Synergy 480 Gen10 Compute Modules running RHCOS version 4. The OpenShift install tool is run to generate ignition files that contain information about the hosts that will be provisioned. The Core OS for the worker nodes is then booted with the help of iPXE and the ignition files are passed with the OS image during installation. HPE Storage systems such as HPE Nimble, HPE 3PAR, and HPE Synergy D3940 Storage Modules provide support for ephemeral and persistent container volume via persistent volume.



Figure 1 provides an overview of the Red Hat OpenShift Container Platform 4 solution layout in a bare metal setup configuration.



**FIGURE 1.** Solution layout in a bare metal setup

**NOTE**

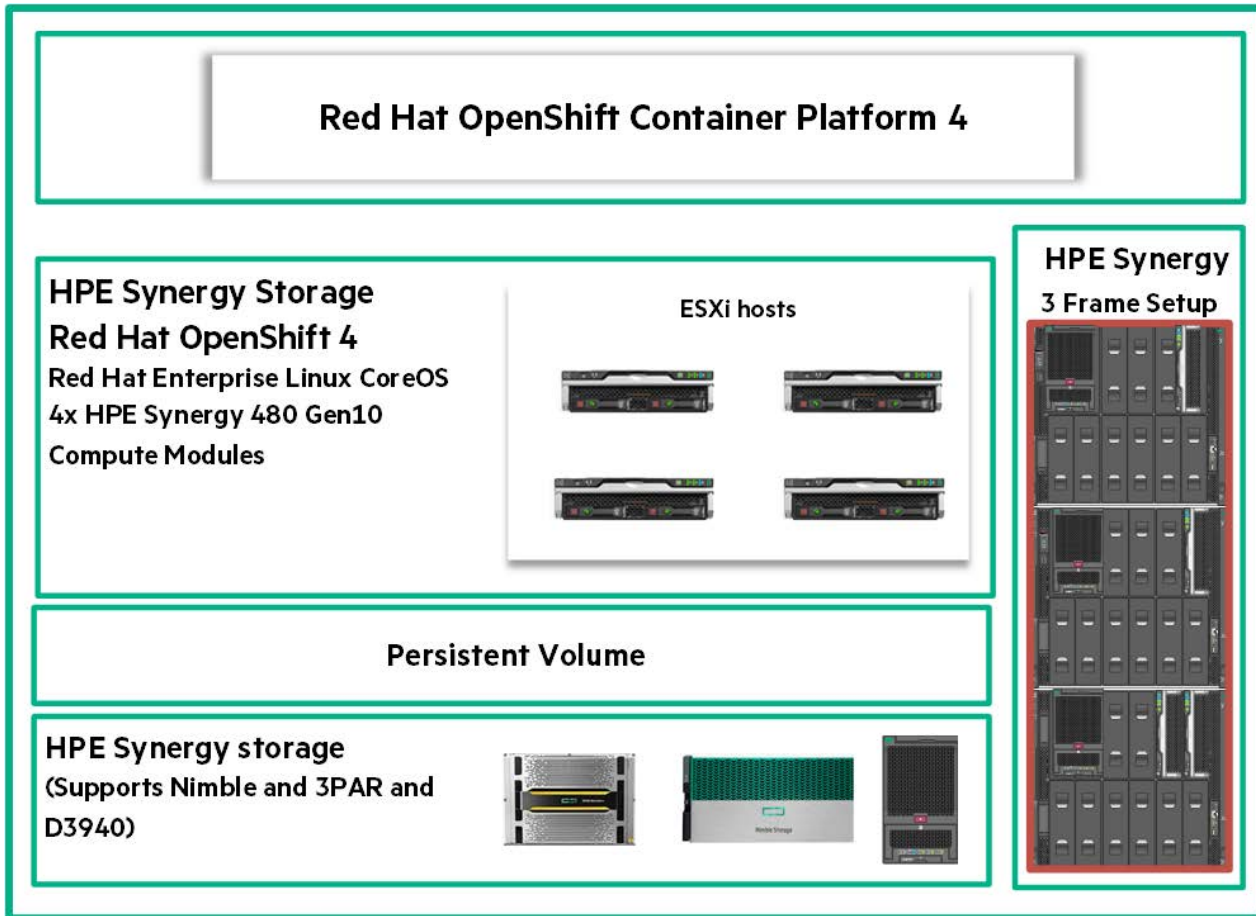
The number of physical nodes represented in the solution layout is only a started block to grow, based on customer needs. This is subject to change based on customer requirements.

**Virtualized configuration**

In a virtualized configuration, the OpenShift master nodes and worker nodes are deployed as VMware vSphere® virtual machines running on three (3) or more HPE Synergy 480 Gen10 Compute Modules. All virtual machines run RHCOS version 4. The OpenShift install tool is run to generate ignition files that contain information about the hosts that will be provisioned. The CoreOS operating system for the worker nodes is then PXE booted and the ignition files are passed with the OS image during installation. HPE Synergy D3940 Storage Module provides support for both ephemeral and persistent volume.



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



**FIGURE 2.** Solution layout in a virtualized setup

**HPE Synergy Composable Infrastructure security controls**

Hewlett Packard Enterprise has security features and functionalities built into servers from the hardware level to the firmware. Customers require a holistic view of the security controls available in the composable infrastructure to make the most of their investment. HPE Synergy Composable Infrastructure enables IT organizations to accelerate application and service delivery with fluid resource pools, made up of compute, storage, and fabric with software-defined intelligence. Each resource within the composable infrastructure is made up of multiple products such as compute modules, which in turn has multiple components such as Integrated Lights-Out (iLO), Unified Extensible Firmware Interface (UEFI), and so on. Another example is the management device such as HPE Synergy Composer, which shows its functions using HPE OneView and the HPE Synergy frame link modules. With many products available within the composable infrastructure, it is important to understand the security controls available within each of them and how they can be used to avoid potential security breaches.

This solution provides a layered view of security controls that are available to Hewlett Packard Enterprise customers.



Figure 3 shows the layered security view across various composable infrastructure components.

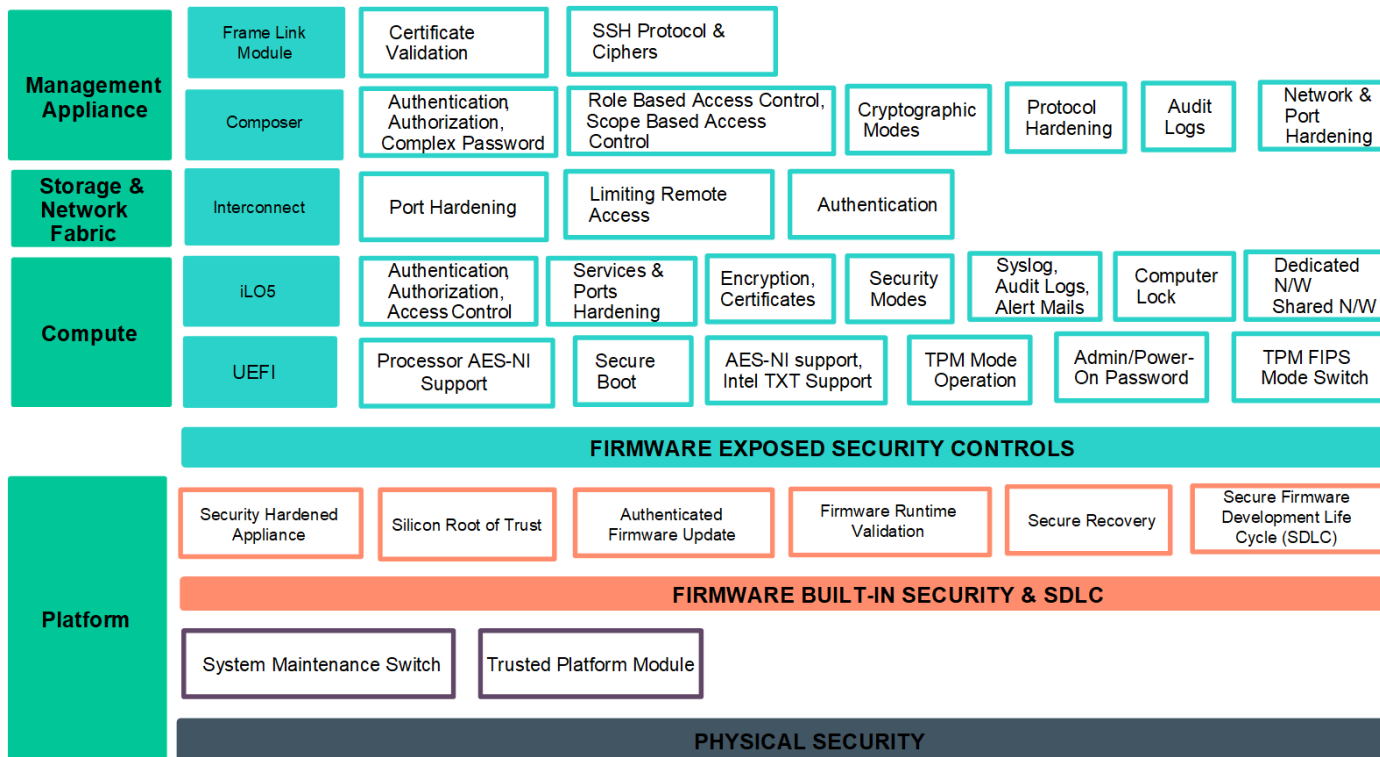


FIGURE 3. Layered security view for composable infrastructure

Table 1 describes the security control layers of physical security, firmware or hardware built-in security, and firmware exposed security control layer. The objective of this layered security view is to ensure that the customers will be aware of the depth of security risk that infrastructure can face. This also provides the knowledge of the in-depth defense that is built-in to the design of HPE Composable Infrastructure. Each security control at each layer is designed to comply with the requirements of some security tenets. The security tenets are a set of security principles that ensure the security within the information systems (for example, Authentication, Authorization, Access Control, Password Policies, Cryptographic Ciphers, Secure Protocols, Forensic Analytics – Logs, Alerts, Threat Modelling, Security Certifications, and Standards).

TABLE 1. Physical and Firmware security controls within HPE Composable Infrastructure

Security Controls Category	Description
<b>Physical security control</b>	Physical security describes measures designed to ensure the physical protection and detection of threat event in the infrastructure
<b>Firmware/Hardware built-in security control</b>	The security technologies built-in firmware makes it more secure for any communication with the underlying hardware and safe for user data at rest or transit. The threat modelling followed by Hewlett Packard Enterprise provides enhanced security to the infra components.
<b>Firmware exposed security control</b>	This is the exhaustive list of security controls that lets the customers: <ul style="list-style-type: none"> <li>• Define the boundaries for accessing various infra components</li> <li>• Set quantum safe ciphers for encryption</li> <li>• Generate alert and log changes to infra</li> </ul>



## Firmware built-in security controls

Hewlett Packard Enterprise has used a variety of technologies to ensure that the built-in firmware security controls provide the highest level of infrastructure security. This section provides a brief overview of the security controls that Hewlett Packard Enterprise has built-in to the firmware used by HPE Synergy and how these security controls offer an added advantage for HPE Synergy customers.

### Silicon Root of Trust

The iLO5 chipset contains a first of its kind Silicon Root of Trust for the HPE Synergy Gen10 compute platform which is included with the iLO standard license. Silicon Root of Trust provides an inextricably tied link between the silicon and firmware making it impossible to insert any malware, virus, or compromised code that would corrupt the boot process. The Silicon Root of Trust enables the boot process to provide a secure start. When the system boots, the iLO5 chip validates and boots its own firmware first, then validates the system BIOS. Since the Silicon Root of Trust is inextricably tied into the iLO5 hardware, every validated signature throughout the boot process can be trusted. However, in the unlikely event that iLO5 finds tampering or corruption at any point in the process, trusted firmware is immediately available for secure recovery. On startup, if iLO5 finds its own firmware has been compromised, it will load its own authenticated firmware from an integrated backup. Regardless of the license, the iLO5 firmware recovery is always available and always automatic. Remember that the Silicon Root of Trust in hardware is how the iLO5 firmware is verified. So, it can always be trusted. If the iLO5 finds that the system BIOS has been compromised, customers can connect to iLO5 and manually recover the authenticated firmware. Since the Silicon Root of Trust is embedded in the hardware itself, iLO5 can detect any compromised firmware as far back as the supply chain process. Hewlett Packard Enterprise can address platform security all the way back to the supply chain because it designs the hardware and firmware of iLO5 entirely and controls the iLO5 production process. Unlike other companies, Hewlett Packard Enterprise does not outsource the server management controller. Hewlett Packard Enterprise also has strict internal processes that dictate the firmware approval process. This gives customers an unprecedented level of assurance that no hackers have compromised the firmware before the server is received.

### Secure recovery

Secure recovery is included in the iLO advanced premium security edition license. This works alongside Silicon Root of Trust to automatically recover firmware back to a known good state in the unlikely event that it is compromised. The Silicon Root of Trust enables the secure start process. As the system boots and iLO5 verifies the series of digital signatures, iLO5 can access trusted firmware immediately and recover it to a good state, if it finds tampering or corruption in its own firmware or the system BIOS. First, if iLO5 finds its own firmware has been compromised, it will load its own authenticated firmware from an integrated backup. Regardless of the license, the secure recovery of iLO5 firmware is always available and always automatic. Second, if iLO5 finds the system BIOS has been compromised, it will try to recover from a backup copy. If the backup copy is also compromised and the customer has upgraded to the iLO advanced premium security edition license, iLO5 can automatically recover the authentic firmware. The standard license provides the opportunity for manual recovery. The Silicon Root of Trust is the foundation for the entire secure state and secure recovery process. This enables HPE Gen10 servers to be the world's most secure industry standard server and provides the extraordinary ability to not only verify the digital signature throughout the boot process but also to recover securely if any firmware is compromised.

### Firmware runtime validation

With the iLO advanced premium security edition license, the iLO5 chipset enables runtime validation of firmware. With firmware runtime verification, the iLO5 chipset performs the same checking process that happens during the boot process continually while the server is running. With the frequency of once a day, iLO5 along with its Silicon Root of Trust runs a background verification check on the iLO5 firmware, UEFI, and other firmware loaded after including the system programmable logic devices (SPLD), Innovation Engine (IE), and Management Engine (ME).

### Authenticated firmware updates

The iLO5 chipset expands the number of firmware items that customers can update directly and securely in the HPE Gen10 servers. This is a standard feature of the iLO5. Firmware items that can be securely validated and updated from the iLO now include SPLDs, HPE ProLiant power interface control utility (PowerPIC) firmware, the Intel® IE and ME, and other low-level system components. The iLO5 contains a firmware repository stored on non-volatile flash memory (NAND), which allows components such as the service pack for HPE ProLiant (SPP) and other firmware updates to be applied and installed offline through iLO5.

### Best practices followed by Hewlett Packard Enterprise to deliver security hardened HPE Synergy Composer appliance

Hewlett Packard Enterprise follows secure development lifecycle and used a security assessment tool called Comprehensive Applications Threat Analysis (CATA) to identify and remediate security defects in the appliance operating system.





---

**NOTE**

The design of the appliance is based on CATA fundamentals and underwent a CATA review.

---

The factors that contribute to appliance security hardening are as follows:

- Appliance is hardened to enforce mandatory access control. This means users of HPE Synergy are provided the role-based access control (RBAC) that enables an administrator to establish access control and authorization for users based on their responsibilities.
- Important services of the appliance run with required privileges. This implies HPE Synergy Composer is governed by scope-based access control that enables an administrator to establish access control for users by allowing a role to be restricted to a subset of resources managed by the appliance.
- The appliance is configured and maintained by a firewall, which blocks unused ports. Restricting the usage of all non-essential ports reduces the attack surface for HPE Synergy Composer.
- The appliance operating system bootloader is password protected. This means HPE Synergy Composer cannot be compromised by someone attempting to boot in single-user mode.
- The appliance is designed to operate in an isolated management LAN. Hewlett Packard Enterprise recommends the creation of a private management LAN, known as air-gapped, and keep that separate from production LANs, using VLAN or firewall technology or both.
- Hewlett Packard Enterprise supports the digital signature of all software or firmware updates to ensure their integrity and authenticity. This implies that when the customer is re-imaging the composer to quickly bring it to a specific firmware revision level, the reimaging process is used to verify the digital signature.
- Operating system level users are not allowed to access the appliance, with the following exceptions:
  - A special preset command can be used only if the Infrastructure administrator password is lost or forgotten. This command requires that you contact your authorized support representative to obtain a one-time password.
  - A setting that enables an authorized support representative to obtain a one-time password, so that they can log in to the appliance console (and only the console) to perform advanced diagnostics. Customers can either enable or disable access with this setting.
- Hewlett Packard Enterprise closely monitors the security bulletins for threats to appliance software components and, if necessary, issues software updates.

**Data protection**

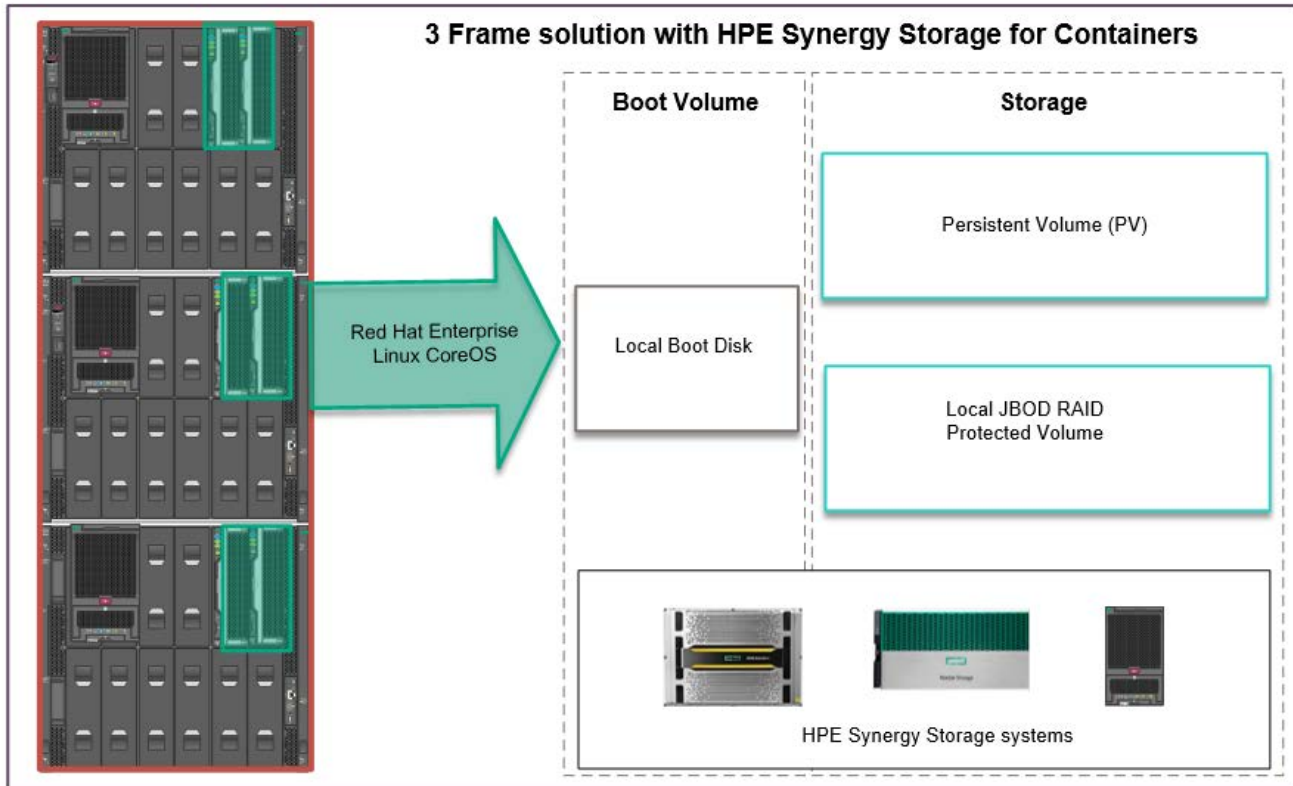
The business requirements should define the container application data protection architecture. These requirements include factors such as the speed of recovery, the maximum permissible data loss, and data retention needs. The data protection plan must also take into consideration the various regulatory requirements for data retention and restoration. Finally, different data recovery scenarios must be considered, ranging from the typical foreseeable recovery resulting from user or application errors to disaster recovery scenarios that include the complete loss of a site. For basic data protection at a hardware layer, HPE Synergy D3940 Storage offers RAID protected volume. For more information, see [HPE Synergy 12Gb SAS Storage User Guide](#). Simplified data protection management by HPE Nimble Storage provides high-performance and array-based data protection capabilities that have been integrated with leading third-party offerings to reduce operational complexity and simplify data protection management.

For advanced data protection features such as scheduling snapshots, application-aware and application-consistent snapshots, replication of snapshots, local and cloud snapshots are recommended for production use-cases. It is recommended to go with the third-party data protection providers that provide Kubernetes application backup.



## SOLUTION LAYOUT

Figure 4 highlights the solution at a high level.




**FIGURE 4.** High-level overview of the solution layout

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

- LDAP or Active Directory
- DHCP
- DNS
- NTP
- TFTP
- iPXE



Figure 5 shows the layout of the hardware within the racks along with storage devices.



**FIGURE 5.** Hardware layout within the rack along with HPE Storage systems

**NOTE**

The figure depicts the hardware layout in the test environment. However, this is subject to change based on the client’s requirement.



## SOLUTION COMPONENTS

### Hardware

Table 2 shows the list of hardware components used in this solution.

**TABLE 2.** List of hardware components

Component	Qty	Description
<b>HPE Synergy components</b>		
HPE Synergy 12000 Frame	3x	Three (3) HPE Synergy 12000 Frames house the infrastructure used for the solution.
HPE Synergy Composer	2x	Two (2) HPE Synergy Composers for core configuration and lifecycle management for the Synergy components.
HPE Virtual Connect SE 40Gb F8 Module	2x	Two (2) HPE Virtual Connect SE 40Gb F8 modules provide network connectivity into and out of the frames.
HPE Synergy 12G SAS Connection Module	6x	Six (6) HPE 12G SAS Connection Modules (two (2) per frame).
HPE Synergy 480 Gen10 Compute Module	6x	Three (3) bare metal master nodes and three (3) bare metal worker nodes as described in this document.
<b>HPE Storage systems</b>		
HPE Synergy D3940 Storage Module	3x	Three(3) HPE Synergy D3940 12Gb SAS CTO Drive Enclosure with 40 SFF (2.5in) Drive Bays
HPE Nimble Storage	1x	One (1) Array for persistent volume
HPE 3PAR StoreServ	1x	One (1) HPE 3PAR array
<b>Switches</b>		
HPE FlexFabric 5945 2-Slot Switch	2x	Each switch contains one each of the HPE 5945 modules listed as follows:
HPE 5945 24p SFP+ and 2p QSFP+ module	2x	One (1) module per HPE FlexFabric 2-Slot Switch
HPE 5945 8p QSFP+ module	2x	One (1) module per HPE FlexFabric 2-Slot Switch

### NOTE

The HPE Storage systems mentioned in Table 2 is for reference purposes. Use the correct storage information based on the deployment configuration.

### HPE Synergy

HPE Synergy, the first platform built for composable infrastructure empowers IT to create and deliver new value instantly and continuously. This single infrastructure reduces operational complexity for traditional workload and increases operational velocity for the new breed of applications and services. Through a single interface, HPE Synergy composes compute, storage, and fabric pools into the configuration for any application. It also enables a broad range of applications from bare metal to virtual machines, to containers, and operational models such as hybrid cloud and DevOps. HPE Synergy enables IT to rapidly react to new business demands.

HPE Synergy Frames contain a management appliance called the HPE Synergy Composer which hosts HPE OneView. HPE Synergy Composer manages the composable infrastructure and delivers:

- Fluid pools of resources, where a single infrastructure of compute, storage, and fabric boots up ready for workload and demonstrates the self-assimilating capacity
- Software-defined intelligence, with a single interface that precisely composes logical infrastructure at near-instant speed and demonstrates template-driven frictionless operations
- Unified API access, which enables simple line-of-code programming of every infrastructure element, easily automates IT operational processes and effortlessly automates applications through infrastructure deployment

HPE Synergy Composer provides the enterprise-level management to compose and deploy system resources to meet your application needs. This management appliance uses software-defined intelligence to aggregate compute, storage, and fabric resources in a manner that scales to your application needs, instead of being restricted to the fixed ratio of traditional resource offerings. HPE Synergy template-based provisioning enables fast time to service with a single point for defining compute module state, pooled storage, network connectivity, and boot image.



HPE OneView is a comprehensive unifying platform designed for converged infrastructure management. A unifying platform increases the productivity of every member of the internal IT team across servers, storage, and networking. HPE OneView provides organizations an efficient way to work by streamlining processes, incorporating best practices, and creating a new holistic way to work. It is designed for open integration with existing tools and processes to extend these efficiencies.

HPE OneView is instrumental for the deployment and management of HPE servers and enclosure networking. It collapses infrastructure management tools into a single resource-oriented architecture that provides direct access to all logical and physical resources of the solution.

Logical resources include server profiles and server profile templates, enclosures and enclosure groups, and logical interconnects and logical interconnect groups. Physical resources include compute modules, interconnects, and storage modules.

HPE OneView offers a uniform way for administrators to interact with resources by providing a RESTful API foundation. The RESTful APIs enable administrators to utilize a growing ecosystem of integrations to further expand the advantages of the integrated resource model. This model removes the need for the administrator to enter and maintain the same configuration data more than once and keep all versions up to date. It encapsulates and abstracts many underlying tools behind the integrated resource model. This enables the administrator to operate with new levels of simplicity, speed, and agility to provision, monitor, and maintain the solution.

Within the context of the solution, HPE OneView for Synergy is utilized to:

- Configure the profiles of the HPE Synergy compute modules
- Apply and maintain compliance for firmware across the HPE Synergy infrastructure
- Configure networking from the HPE Synergy compute modules to inbound and outbound destinations

### **HPE Synergy 480 Gen10 Compute Module**

The HPE Synergy 480 Gen10 Compute Module delivers an efficient and flexible two-socket server to support the most demanding workload. The compute module is powered by Intel Xeon® Scalable family of processors, up to 3TB DDR4, and large storage capacity within a composable architecture.

HPE Synergy 480 Gen10 Compute Module:

- Provides the most secure server with exclusive HPE Silicon Root of Trust. This secure server protects your applications and assets against downtime associated with hacks and viruses
- Offers customer choice for greater performance and flexibility with Intel Xeon Scalable family of processors on the HPE Synergy 480 Gen10 architecture
- Offers intelligent system tuning with processor smoothing and workload matching to improve processor throughput and an increase in overall performance by 8% over an earlier generation
- Features a maximum memory footprint of 3TB for large in-memory database and analytic applications
- Features a hybrid HPE Smart Array for both RAID and HBA zoning in a single controller

The HPE Synergy 480 Gen10 Compute Module provides the required compute to power this solution running both Red Hat Virtualization for the core management pieces of Red Hat OpenShift and RHCOS to host the worker nodes.

The bill of materials found in [Appendix A](#) of this document outlines the configuration of the HPE Synergy 480 Gen10 Compute Modules used in this solution.

### **HPE Synergy D3940 Storage Module**

The HPE Synergy D3940 Storage Module is a direct-attached storage module with 40 Small Form Factor (SFF) drive bays designed to use in HPE Synergy 12000 Frames. Through the HPE Synergy 12Gb SAS connection module, it provides composable storage for up to 10 compute modules in a single frame. HPE Synergy storage is optimized to use either as a direct-attached storage array or with software-defined storage solutions, such as Red Hat OpenShift Container Storage.



HPE Synergy storage enables a variety of workloads by permitting multiple drive types to be configured in the same storage module. HPE Synergy D3940 Storage Modules support a family of 12G SAS or 6G SATA HDD and SSD Smart Drives. Storage modules connect to compute modules within a frame through the HPE Synergy 12 Gb SAS Connection Module. Any number of drive bays in a storage module can be composed of any compute module containing a Smart Array controller connected to the SAS fabric. This allows efficient utilization of available drives. HPE Synergy storage can scale up to 200 SFF drives with five storage modules in a single HPE Synergy 12000 Frame. Adding a second I/O adapter and second SAS connection module provides a redundant path to SAS drives inside the storage module, ensuring high availability. The modular design of the HPE Synergy D3940 Storage Module allows to slide out from the frame to service drives or I/O adapters without interrupting the operation of other drives within the module. Additionally, the HPE Synergy D3940 Storage Module is optimized for solid-state drives using a high-performance SAS connection with sixteen 12 Gb/s SAS lanes. This allows a single HPE Synergy storage module to deliver as much as 8 times the bandwidth of other JBOD options reaching up to 2M IOPs.

### HPE Nimble Storage

HPE Nimble Storage AF40 is used to provide persistent, block storage in this solution. The HPE Nimble Storage array for Docker data provides the storage volume, to host the repository, to store container images, and to provide persistent volume for applications.

### HPE Nimble Storage persistent volume

Using the HPE-CSI-driver for containers, HPE Nimble Storage can provide persistent volume for the Kubernetes cluster in this solution. With HPE Nimble Storage, DevOps, IT teams, and others can be assured that their environment exceeds enterprise-grade requirements for supporting backup or protection and disaster recovery of data.

### HPE 3PAR

HPE 3PAR storage provides the world's most intelligent storage for mission-critical apps that provides on-demand experience for instant access to data and consumable as a service. The HPE 3PAR StoreServ flash-optimized data storage system modernizes the data center to handle unpredictable workloads effortlessly. Some of the other supported features include rapid and automated provisioning, multitenant design, hardware-accelerated deduplication, and compression.

HPE 3PAR StoreServ 8440 Storage provides shared and dedicated storage for a variety of purposes within this solution including housing virtual machines and persistent volume.

## Software

This section describes the software versions utilized in the solution along with any special installation or configuration requirements. Table 3 lists the software version used in this solution.

**TABLE 3.** Software version

Component	Version
Red Hat Enterprise Linux CoreOS (RHCOS)	4.4
Red Hat OpenShift Container Platform	4
HPE Nimble OS	5.0.8
HPE 3PAR OS	3.3.1

### iPXE

iPXE is an open-source implementation of the Preboot eXecution Environment (PXE) client firmware and bootloader. It is used to enable servers without built-in PXE support to boot from the network. In this solution, the master and worker nodes configured for OpenShift Container Platform 4 deployment are booted with CoreOS leveraging the iPXE configured in the OCP environment. The details of iPXE configuration and procedure is described in the Deployment guide at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable>.



**Red Hat Enterprise Linux CoreOS (RHCOS)**

RHCOS represents the next generation of single-purpose container operating system technology. RHCOS combines the quality standards of Red Hat Enterprise Linux (RHEL) with the automated and remote upgrade features from Linux containers. RHCOS is supported only as a component of Red Hat OpenShift Container Platform 4 for all Red Hat OpenShift Container Platform machines. RHCOS is the only supported operating system for Red Hat OpenShift Container Platform control plane or master machines. While RHCOS is the default operating system for all cluster machines, you can create compute, or worker machines that use Red Hat Enterprise Linux 7 as their operating system.

This solution is built on RHCOS. Each Red Hat OpenShift Container Platform 4 control plane and worker nodes are running on RHCOS as the dedicated physical server node functions.

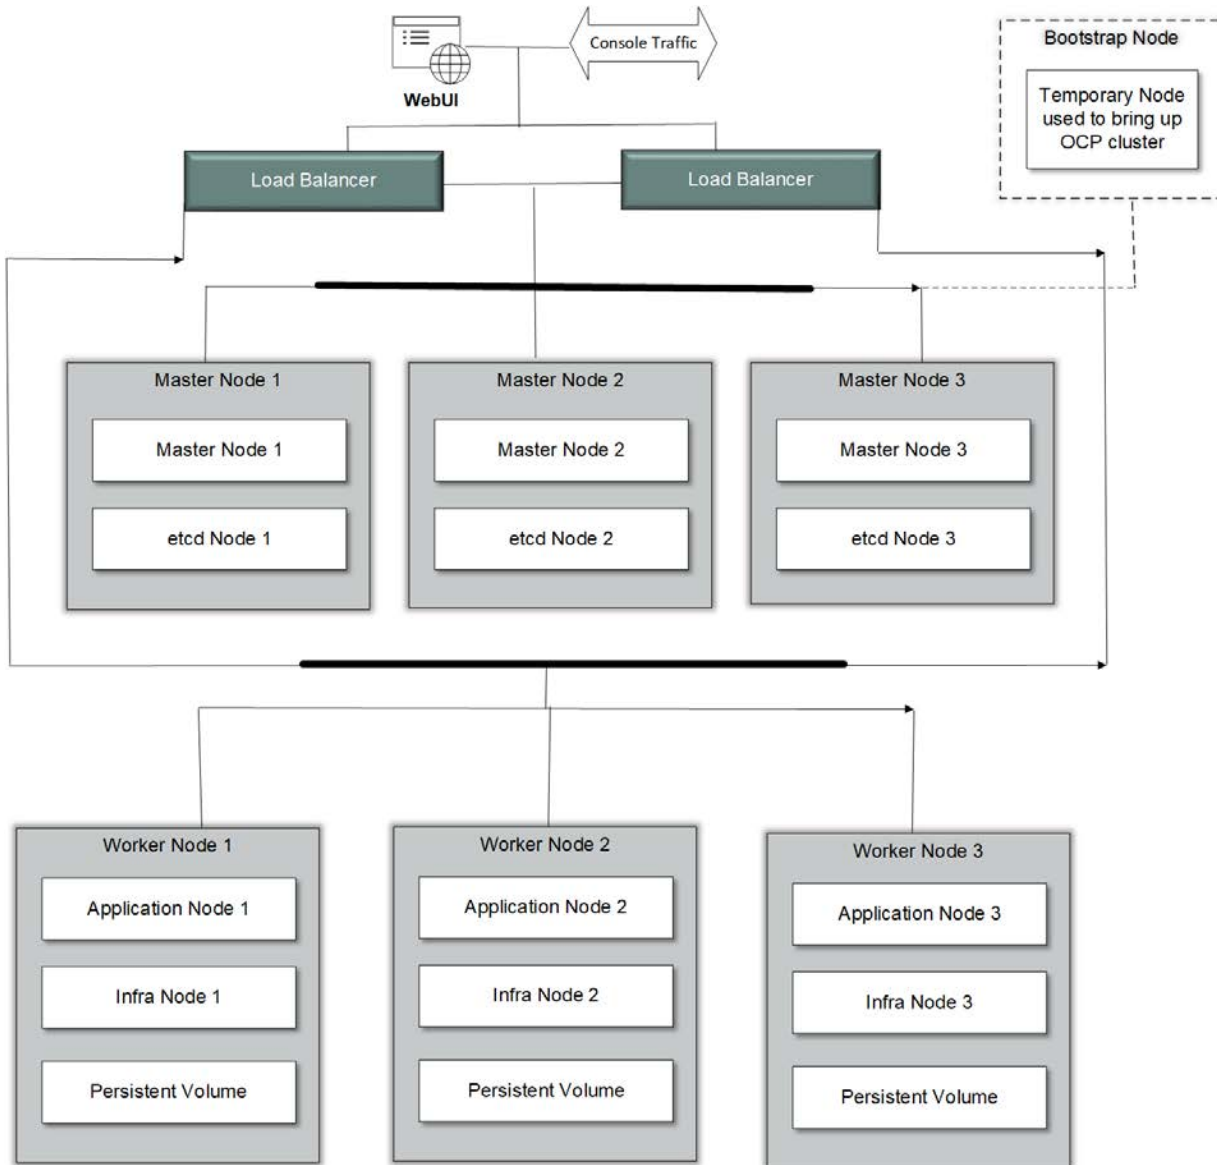
**Red Hat OpenShift Container Platform 4**

Red Hat OpenShift Container Platform 4 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 Container Platform helps businesses achieve greater value by delivering modern and traditional applications with shorter development cycles and lower operating costs. Red Hat OpenShift Container Platform 4 is built on open source innovation and industry standards, including Kubernetes and RHCOS, the enterprise Linux container distribution.

Red Hat OpenShift Container Platform 4 can be provisioned with persistent storage by using local volumes. Local persistent volume allows you to access local storage devices, such as a disk or partition, by using the standard PVC interface. Local volume can be used without manually scheduling pods to nodes because the system is aware of the volume node constraints. However, the local volume is still subject to the availability of the underlying node and is not suitable for all applications.



Figure 6 describes how the individual Red Hat OpenShift Container Platform 4 pieces are laid out.



**FIGURE 6.** Red Hat OpenShift Container Platform 4 layout






## BEST PRACTICES AND CONFIGURATION GUIDANCE FOR THE SOLUTION

This section discusses the high-level cabling and configuration of the solution hardware and software.

### Solution cabling

Figure 7 describes the cabling configuration of the three (3) HPE Synergy 12000 Frames as well as the HPE FlexFabric 5945 Switches and Intelligent Resilient Fabric (IRF) within the context of this solution. These cables carry frame management, inter-frame and interconnect traffic between frames.



**FIGURE 7.** Frame and switch cabling within the solution


### NOTE

For details on the cabling for the HPE Storage systems used in this solution, see the Deployment Guide located at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy>.



### Networking

Figure 8 documents the cabling of the solution from the HPE Virtual Connect SE 40 GB modules to the switches. All egressing Ethernet networks are carried on a single bridge-aggregation group (BAGG). Top of rack switching was used in the creation of this solution, but end of row switching is equally effective in HPE Synergy environments and can reduce overall solution costs by reducing the number of physical switches.



**FIGURE 8.** Network cabling from the HPE Synergy 12000 Frames to the switches

---

#### NOTE

It is recommended to configure redundant network(s) for production environment.

---



Table 4 describes the configuration of the networks as defined within HPE OneView for HPE Synergy and the bandwidth associated with each network. Network is carried outbound on a single BAGG except the Synergy Management network.

**TABLE 4.** Networks defined within HPE OneView for HPE Synergy

Network Name	Type	VLAN Number	Purpose	Requested Bandwidth (Gb)	Maximum Bandwidth (Gb)
Management	Ethernet	1193	Solution management	5	20
Data_Center	Ethernet	2193	Application access, authentication and other user networks	10	20
Synergy Management	Ethernet	193	OneView for Synergy	2	20
iPXE_Boot	Ethernet	2193	Boot for compute	3	20
iSCSI_VLAN_A	Ethernet	3193	iSCSI_VLAN_A	10	40
iSCSI_VLAN_B	Ethernet	3194	iSCSI_VLAN_B	10	40

Table 5 explains the cabling of the HPE Virtual Connect Interconnect Modules to the HPE FlexFabric 5945 Switch and highlights what networks are carried on the connections.

**TABLE 5.** Uplink set mapping

Uplink Set	Synergy Source	Switch Destination
Network	Enclosure 1 Port Q3	FortyGigE1/1/1
	Enclosure 1 Port Q4	FortyGigE2/1/1
	Enclosure 2 Port Q3	FortyGigE1/1/2
	Enclosure 2 Port Q4	FortyGigE2/1/2
iSCSI_SAN_A	Enclosure 1 Port Q5	FortyGigE1/1/5
	Enclosure 1 Port Q6	FortyGigE1/1/6
iSCSI_SAN_B	Enclosure 2 Port Q5	FortyGigE2/1/5
	Enclosure 2 Port Q6	FortyGigE2/1/6

Utilizing HPE Synergy, the networks within the solution can traverse the HPE Synergy infrastructure in an east-west fashion across high speed, low latency links both within and between HPE Virtual Connect Modules. The communication between the Red Hat OpenShift Container Platform 4 management pieces remains within the HPE Synergy Frames.

## Storage

The HPE Synergy D3940 Storage Module provides SSDs and optional Hard Disk Drives (HDDs) consumed by the Local Storage Operator, if not utilizing local disks within the HPE Synergy 480 Gen 10 Compute Modules. It can also provide boot volumes.

HPE Nimble Storage and HPE 3PAR dynamically provides Persistent Volumes (PVs) for containers using Dynamic Volume Provisioner which is integrated with the HPE Container Storage Interface Driver.



Figure 9 describes the logical storage layout used in the solution.

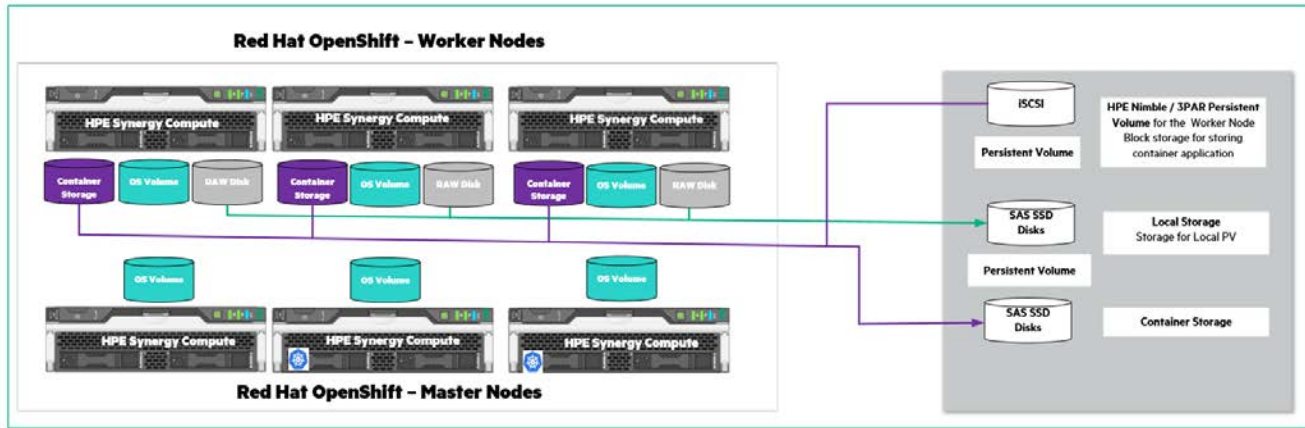


FIGURE 9. Logical storage layout within the solution

Table 6 lists all volumes used within the solution for the storage systems and highlights what storage provides the capacity and performance for each function.

TABLE 6. Details of the volume

Source	Volume/Disk Function	Hosts	Shared/Dedicated
HPE Synergy D3940 Storage	Remote Volume	OpenShift worker nodes	Shared
	Operating System (Optional)	All Nodes	Dedicated
HPE Nimble	Remote Volume	OpenShift worker nodes	Shared
	Operating System (Optional)	All Nodes	Dedicated
HPE 3PAR	Remote Volume	OpenShift worker nodes	Shared
	Operating System (Optional)	All Nodes	Dedicated

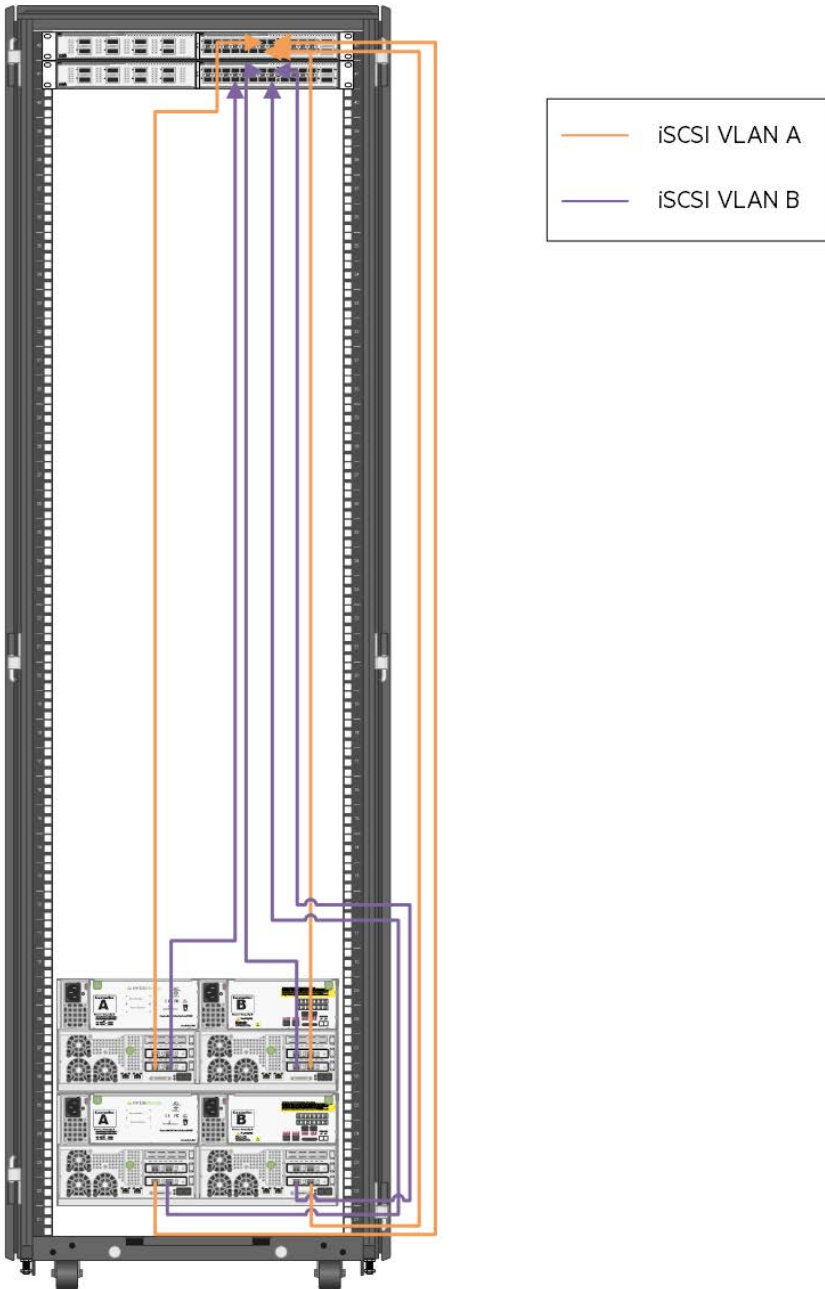
**NOTE**

The OpenShift Container registry data is hosted on the persistent application data store.



**Cabling of the HPE Nimble Storage**

An HPE Nimble Storage AF40 array provides shared and dedicated storage for a variety of purposes within this solution. Figure 10 shows the cabling of the HPE Nimble Storage AF40 to the HPE switching utilized in this solution. Note that this diagram shows the storage and switching in the same rack to provide clarity. As implemented for this solution, the switching resided in the HPE Synergy rack. The orange and purple wires in the figure represent the separate iSCSI VLANs.




**FIGURE 10.** Cabling of the HPE Nimble Storage arrays to the HPE FF 5940 switches



### Cabling of the HPE 3PAR StoreServ

HPE 3PAR StoreServ 8440 Storage provides shared and dedicated storage within this solution. Figure 11 shows the cabling of the HPE Synergy 12000 Frames to the HPE Fibre Channel switching and cabling of the HPE 3PAR StoreServ 8440 Storage to the HPE Fibre Channel switching utilized in this solution. Cabling for each array follows the same guidelines for maximum availability.



**FIGURE 11.** Cabling of the HPE Synergy interconnects and HPE SN6600B switching

### Server profile

HPE Synergy Composable Infrastructure using HPE Virtual Connect provides the construct of a server profile. A server profile allows a suite of configuration parameters including network and storage connectivity, BIOS tuning, boot order configuration, local storage configuration applied programmatically to compute resources. These features along with firmware upgrade, iLO updates, and BIOS settings updates are automated and are controlled by input files by the user. These templates are key to deliver the “infrastructure-as-code” capabilities of the HPE Synergy platform. For the purpose of this solution, a single template was created that was applied to all OpenShift worker compute modules.



The critical items configured as a part of the template supporting Red Hat OpenShift Container Platform 4 are the network connections and storage. Figure 12 describes the configuration of the network interfaces as part of the profile template for the worker nodes. There are two (2) redundant Ethernet networks that are used for network communication.

### Connections







ID	Name	Network	Port	Boot
● 1	Management	<u>TwentyNet</u> VLAN2193	Mezzanine 3:1-a	Not bootable  
	Type	Ethernet		
	MAC address	52:4D:1F:20:03:02 (v)		
	Requested virtual functions	Auto		
	Requested bandwidth	10 Gb/s		
	Link aggregation group	None		
● 4	iSCSI_A	<u>iSCSI_SAN_A</u> VLAN3193	Mezzanine 3:2-a	Not bootable  
	Type	Ethernet		
	MAC address	52:4D:1F:20:03:06 (v)		
	Requested virtual functions	None		
	Requested bandwidth	2 Gb/s		
	Link aggregation group	None		
● 5	iSCSI_B	<u>iSCSI_SAN_B</u> VLAN3194	Mezzanine 3:1-d	Not bootable  
	Type	Ethernet		
	MAC address	52:4D:1F:20:03:07 (v)		
	Requested virtual functions	None		
	Requested bandwidth	2 Gb/s		
	Link aggregation group	None		

FIGURE 12. Server connections as part of the profile for OpenShift nodes

HPE Storage modules are used to create the HPE Data Fabric as shown in Figure 13. The solid-state disks and/or hard disk drives provided by HPE Nimble, HPE 3PAR, and HPE Synergy D3940 is consumed by the HPE Data Fabric, and optionally, by compute nodes as boot devices.

### Local Storage

Integrated storage controller 

Managed by OneView

Write cache Managed manually


Name	Type	Logical Drive ID	RAID Level	Number of Drives	Size GB	Drive Technology	Boot	Accelerator
disk01	Logical drive	1	RAID 1	2	n/a	not specified	<input checked="" type="checkbox"/>	Managed manually 

FIGURE 13. Server connections as part of the profile for OpenShift nodes



## Capacity and sizing

Sizing for a Red Hat OpenShift Container Platform 4 environment varies depending on the requirements of the organization and type of deployment. The installer should read and understand Red Hat recommendations around scalability and performance before installation. This ensures the need for their environment is addressed. Red Hat publishes documentation around scalability and performance for each OpenShift Container Platform release. For more information on the documentation for OpenShift Container Platform 4, see [https://docs.openshift.com/container-platform/4.4/scalability\\_and\\_performance/recommended-host-practices.html](https://docs.openshift.com/container-platform/4.4/scalability_and_performance/recommended-host-practices.html). This URL points to 4.4 version of the Red Hat OpenShift being installed.

## Non-root privileges and sudo command

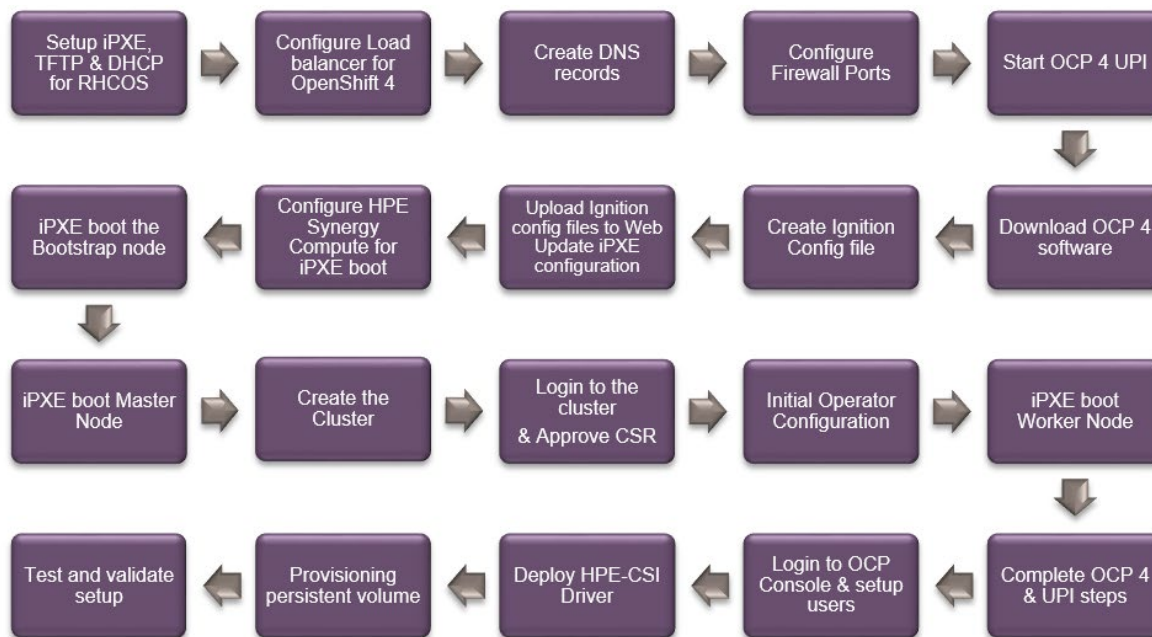
The industry-wide security best practice is not to use the root user account for the administration of the Linux distro-based servers. However, certain operations require root user privileges. Sudo command is used in such scenarios. The sudo command allows you to run programs with the security privileges of another user (by default, as the root). It prompts the user for their password and confirms the request to execute a command by checking a file called sudoers, which the system administrator configures. Using the sudoers file, system administrators can give certain users or groups access to some or all commands without those users having to know the root password. It also logs all commands and arguments for traceability.

## NOTE

The root user is not active by default in RHCOS, so the root login is not available. Instead, user can log in as core user.

## DEPLOYMENT OVERVIEW

Figure 14 shows the overall deployment flow for the solution.



**FIGURE 14.** Red Hat OpenShift Container Platform 4 deployment overview

## Setup iPXE, TFTP, and DHCP for RHCOS

In this setup, the machine is booted by leveraging the iPXE server. This prepares the iPXE and TFTP server to able to boot RHCOS. This is the initial stage and DHCP is an integral part of the PXE boot process. So, configuring the DHCP is also important. This configuration can be done using sudo access. The details to configure the iPXE set up is listed in the Deployment guide at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable>.





## Configure a load balancer for Red Hat OpenShift 4 nodes

In multi-node cluster deployment of OpenShift, the load balancer is mandatory. Hewlett Packard Enterprise has leveraged HAProxy load balancing required traffic. This configuration can be done using sudo access. For commercial load balancer such as F5 Big-IP or any other OpenShift Container Platform 4 supported load balancer, you need to visit the manufacture website. For more details on configuring sudo to allow non-root users to execute root level commands and for information on HAProxy configuration, see the HPE solutions for Red Hat OpenShift Container Platform GitHub at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable>.

## Configure DNS

In user-provisioned infrastructure (UPI), DNS records are required for each machine. These records must be able to resolve the hostnames of all other machines in a Red Hat OpenShift Container Platform cluster. This component also can be configured using sudo access for Linux-based DNS solution or Windows-based DNS solution. For more information on Role-based access control, see [Windows Role based Access](#). For third-party DNS solutions, you need to visit the manufacture website. This provides more details to configure the sudo to allow non-root users to execute root level commands. For more information, see section [User-provisioned DNS requirements](#).

## Configure firewall ports

In user-provisioned infrastructure (UPI), network connectivity between machines allows cluster components to communicate within the Red Hat OpenShift Container Platform cluster. Hence, the required ports must be open between Red Hat OpenShift cluster nodes. This component also can be configured using sudo access for Linux-based firewall. For third-party firewall solutions, you need to visit the manufacture website. This provides details on configuring sudo to allow non-root users to execute root level commands. For more information, see section [Networking requirements for user-provisioned infrastructure](#).

## Start Red Hat OpenShift Container Platform 4 user provisioned infrastructure setup

In user provisioned infrastructure (UPI), it begins with installing a bastion host. This setup uses RHEL 7.6 virtual machine as a bastion host. This bastion host is used for deployment and management of the Red Hat OpenShift Container Platform 4 version clusters. The setup and configuration can be completed using sudo user access. For more information, see section [Generating an SSH private key and adding it to the agent](#).

## Download Red Hat OpenShift Container Platform 4 software version

Download OpenShift Container Platform 4. Check the access token for your cluster and install it on the bastion host. The bastion host is used for deploying and managing the OpenShift Container Platform 4 version clusters. The setup and configuration can be completed using sudo user access. For more information, see section [Obtaining the installation program](#).

## Create ignition config files

This step begins with the creation of the `install-config.yaml` in a new folder. Use the OpenShift install tool to convert the yaml to the ignition config files required to install the Red Hat OpenShift Container Platform 4. There is no system modification done on the bastion host or the provisioning server. This setup can be completed using sudo access. For more information, see [https://docs.openshift.com/container-platform/4.4/installing/installing\\_bare\\_metal/installing-bare-metal.html#installation-initializing-manual\\_installing-bare-metal](https://docs.openshift.com/container-platform/4.4/installing/installing_bare_metal/installing-bare-metal.html#installation-initializing-manual_installing-bare-metal).

In a virtualized setup to deploy OpenShift Container Platform 4, the ignition config files for master, worker, and bootstrap nodes are converted to Base64 encoding. For more information on creating an installation configuration file, see [https://docs.openshift.com/container-platform/4.4/installing/installing\\_vsphere/installing-vsphere.html](https://docs.openshift.com/container-platform/4.4/installing/installing_vsphere/installing-vsphere.html).

## Upload ignition config files to the web

This step involves uploading the ignition config files to an internal website that allows anonymous access to the PXE boot process. Update the PXE default file to point to the website location of the ignition file. The action required in this step can be done using sudo user. For more information, see section [Creating Red Hat Enterprise Linux CoreOS \(RHCOS\) machines by PXE or iPXE booting](#).

In a virtualized setup to deploy OpenShift Container Platform 4, a template for OVA image is created. This template is used for creating nodes on the cluster. The ignition config files are provided on each node while provisioning the VMs. For more information on Create template for OVA Image, see [https://docs.openshift.com/container-platform/4.4/installing/installing\\_vsphere/installing-vsphere.html](https://docs.openshift.com/container-platform/4.4/installing/installing_vsphere/installing-vsphere.html).



## Configure the HPE Synergy Compute for iPXE boot

The configuration involves setting up the server profile in HPE Synergy Composer for iPXE boot and for the required storage. Hewlett Packard Enterprise uses the HPE Synergy Composer to create the server profiles and templates. The access to the composer UI is that of a non-root user. Hence, from a security standpoint, no root access is being used for HPE Synergy Composer access. For more information, see HPE OneView 4.2 User Guide [Server profiles and server profile templates section](#).

### Bootstrap Node

The bootstrap node is a temporary node that is used to bring up the OpenShift cluster. After the cluster is up, this machine can be decommissioned, and the hardware will be reused. The iPXE boot process must use bootstrapping information as a part of the iPXE boot parameter to install the RHCOS on this node.

### Master Node

The master node uses the iPXE image for RHCOS after the bootstrap node. The iPXE boot process must use the `master.ign` information as part of the iPXE boot parameter to install the RHCOS on this node. The root user is not active by default in RHCOS. So, root login is not available. Instead, log in as the core user.

### Create the cluster

The four nodes, one bootstrap and three master nodes boot up and are available at the login prompt for RHCOS. Use the OpenShift install tool to complete the bootstrap process. For more information, see [Creating the cluster](#). This action is taken using the sudo user logged in on the bastion host or provision server.

### Login to the cluster

After the bootstrap process has completed successfully, login to the cluster. The `kubeconfig` file is present in the `auth` directory where the ignition files are created on the bastion host. Export the `cluster_kubeconfig` file and login to your cluster as a default system user. The `kubeconfig` file contains information about the cluster that is used by the CLI to connect a client to the correct cluster and API server. The file is specific to a cluster and is created during the Red Hat OpenShift Container Platform installation. After logging in, approve the pending OpenShift CSR for the nodes. For more information, see section [Logging in to the cluster](#).

### Initial operator configuration

After the control plane initializes, you must immediately configure operators that are not available. This ensures their availability (for example image-registry). For more information, see section [Initial Operator Configuration](#). This action is taken using the sudo user logged in on the bastion host or provision server.

### Worker node

This step involves decommissioning the bootstrap node and deleting the associated HPE Synergy server profile. Boot the compute nodes associated with worker node profile that has the second volume for local storage setup using iPXE. The root user is not active by default in RHCOS. So, root login is not available, instead, log in as the core user.

### Complete the installation of Red Hat OpenShift Container Platform 4 and higher versions for user provisioned infrastructure

After the worker node boots up successfully, use the `oc get nodes` from the bastion host. The admin can see the worker nodes as part of the OpenShift cluster. Run the OpenShift install tool to complete the installation. For more information, see section [completing installation on user-provisioned infrastructure](#). After this process is completed, it will provide the URL for the Red Hat OpenShift Container Platform 4 version of the console along with the temporary user `kubeadmin` and temporary password for login.

### Log in to the Red Hat OpenShift Container Platform 4 and higher versions of the console

Log in to the OpenShift Container Platform 4 version of the console using the URL, username, and password provided in the complete Red Hat OpenShift Container Platform 4 user-provisioned infrastructure. Set up a new user with the cluster admin privileges. For more information, see section [Understanding authentication](#).

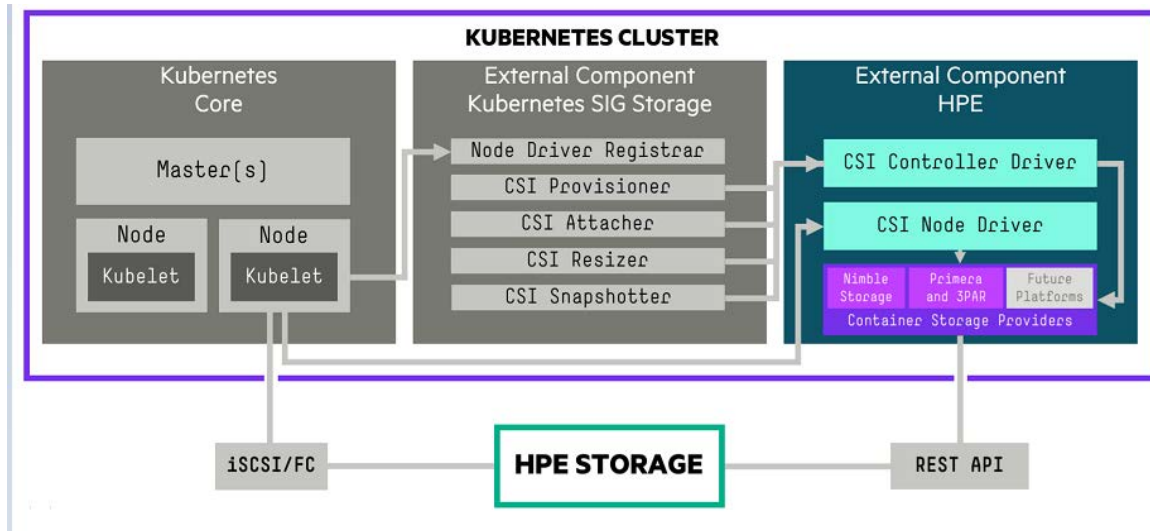


## HPE Container Storage Interface

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

The CSI specification includes constructs to manage snapshots as native Kubernetes objects and create new Persistent Volume Claim (PVC) by referencing those objects. Other capabilities include PVC expansion, inline ephemeral volumes, and the ability to present raw block storage to pods.

Figure 15 shows the HPE CSI Driver for Kubernetes architecture.



**FIGURE 15.** HPE CSI driver for Kubernetes architecture

For more information, see [https://scod.hpdev.io/csi\\_driver/index.html](https://scod.hpdev.io/csi_driver/index.html).

## Red Hat OpenShift Container Storage

Red Hat OpenShift Container Storage (OCS) is software-defined storage that is optimized for container environments. It runs as an operator on Red Hat OpenShift Container Platform 4 to provide highly integrated and simplified persistent storage management for containers. Red Hat OpenShift Container Storage 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.

Red Hat OpenShift Container Storage is deployed as a minimal cluster of three (3) worker nodes. Spread the nodes across three different availability zones to ensure availability. Red Hat OpenShift Container Storage makes use of the Red Hat OpenShift Container Platform default storage class. The OpenShift Container Storage in our test environment was configured on the virtualized set up with HPE Synergy D3940 as storage. The details of OCS configuration and procedure on storage sizing is described in the Deployment guide at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable>.

### NOTE

The storage class on Red Hat OpenShift Container Platform nodes are configured automatically. In case a different storage class is used then it must be changed back to the appropriate storage class as required.



## Provision the local volume

Local volume cannot be created by dynamic provisioning. Instead, the persistent volume must be created by the local storage operator. This provisioner will look for any devices, both filesystem and block volumes, at the path specified in the defined resource. Local volume must be statically created as a PVC that can be accessed by the pod. After a local volume has been mapped to a PVC, it can be specified inside of a resource. For more information, see [Provision the local volumes](#).

---

### NOTE

Local volume can only be used as a statically created persistent volume.

---

## Test and validate the setup

It involves the creation of storage class, persistent volume claim, persistent volume, and pods as per your workload requirements on Red Hat OpenShift Container Platform 4 using local storage operator as a persistent storage provider.

## Physical worker node labeling in OpenShift

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

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

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



## Securing and monitoring OpenShift with Sysdig SaaS

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

The key benefits are:

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

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

---

### NOTE

The Sysdig functionality is not turned on by default in this solution. For more information on how to enable Sysdig, see the Sysdig configuration section listed in HPE solutions for Red Hat OpenShift Container Platform GitHub at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable/>.

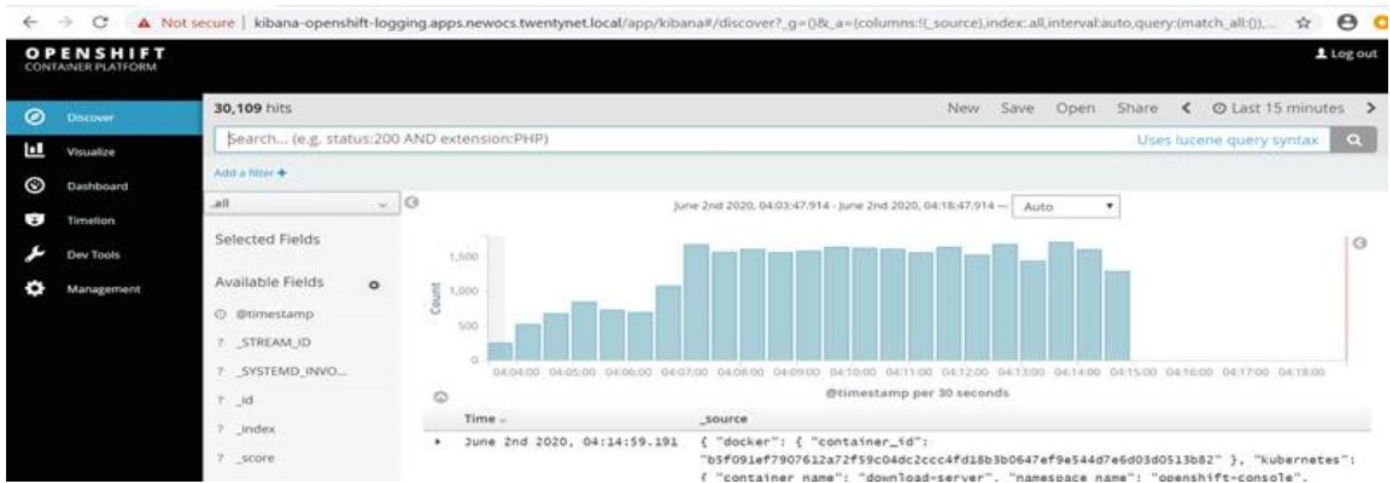
---

## Post-validation deployment

After the successful deployment of Red Hat OpenShift Container Platform 4 on the identified nodes and access to the OpenShift console is successful, the following are deployed from the operator hub for monitoring the logs of the applications deployed on the cluster.

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





**FIGURE 16.** View of Kibana Dashboard

The details of installation and configuration of EFK operators, cluster operators, Kibana etc. are described in the Deployment guide at <https://github.com/HewlettPackard/hpe-solutions-openshift/tree/master/synergy/scalable>.

## NOTE

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

## SUMMARY

Red Hat OpenShift Container Platform 4 on HPE Synergy provides a container solution that eliminates many complexities associated with implementing a container platform at scale. Red Hat OpenShift Container Platform 4 provides organizations with a reliable platform for deploying and scaling container-based applications. HPE Synergy provides the flexible infrastructure needed to run the container platform to dynamically provision and scale applications, whether they run on-premises, in the cloud, or hybrid cloud.

This Reference Configuration utilizes HPE Synergy to deploy Red Hat OpenShift Container Platform 4 and persistent volume. The following benefits can be realized by implementing this solution:

- Deploying the management, etcd, and worker nodes on bare metal eliminates the overhead associated with hypervisors and thus optimizes performance.
- Using an HPE Storage systems such as HPE Nimble, HPE 3PAR, and HPE Synergy D3940 Storage for persistent storage with containers to enable speed, portability, agility, and data protection for traditional enterprise applications and data.
- Using HPE Synergy Composable Infrastructure provides a layered view of security controls. The objective to choose this layered security view is to ensure that the customers are aware of the depth of security risk that infrastructure can face and the depth of defense that is built in the HPE Synergy Infrastructure design.
- Utilizing HPE Synergy Composable Infrastructure provides the ability to programmatically define and deploy servers, storage, and networking as part of the comprehensive container solution.



## APPENDIX A: BILL OF MATERIALS

The following bill of materials contains the core components utilized in the creation of this solution. Services, support, and software are not included in the BOM and the power distribution should be customized based on customer needs.

### NOTE

Part numbers are at time of testing and subject to change. The bill of materials does not include complete support options or other rack and power requirements. If you have questions regarding ordering, please consult your HPE Reseller or HPE Sales Representative. For more details, see. [hpe.com/us/en/services/consulting.html](http://hpe.com/us/en/services/consulting.html).

**TABLE A1.** Bill of materials

Part number	Qty	Description
<b>Rack and Network Infrastructure</b>		
P9K10A	1	HPE 42U 600mmx1200mm G2 Kitted Advanced Shock Rack with Side Panels and Baying
P9K10A 001	1	HPE Factory Express Base Racking Service
H6J85A	1	HPE Rack Hardware Kit
BW932A	1	HPE 600mm Rack Stabilizer Kit
BW932A B01	1	HPE 600mm Rack include with Complete System Stabilizer Kit
AF533A	4	HPE Intelligent Modular 3Ph 14.4kVA/CS8365C 40A/208V Outlets (6) C19/Horizontal NA/JP PDU
<b>HPE Synergy Composable Infrastructure</b>		
797740-B21	3	HPE Synergy 12000 Configure-to-order Frame with 1x Frame Link Module 10x Fans
779218-B21	4	HPE Synergy 20Gb Interconnect Link Module
794502-B23	2	HPE Virtual Connect SE 40Gb F8 Module for Synergy
755985-B21	6	HPE 12G SAS Connectivity Module for Synergy
798096-B21	3	HPE 6x 2650W Performance Hot Plug Titanium Plus FIO Power Supply Kit
804353-B21	2	HPE Synergy Composer
804938-B21	3	HPE Synergy Frame Rack Rail Kit
804942-B21	3	HPE Synergy Frame Link Module
804943-B21	1	HPE Synergy Frame 4x Lift Handles
859493-B21	1	HPE Synergy Multi Frame Master1 FIO
859494-B22	1	HPE Synergy Multi Frame Master2 FIO
804101-B21	8	HPE Synergy Interconnect Link 3m Active Optical Cable
720199-B21	2	HPE Blade System c-Class 40G QSFP+ to QSFP+ 3m Direct Attach Copper Cable
861412-B21	2	HPE Synergy Frame Link Module CAT6A 1.2m Cable
861413-B21	1	HPE Synergy Frame Link Module CAT6A 3m Cable
<b>Master Nodes</b>		
871940-B21	3	HPE Synergy 480 Gen10 Configure-to-order Compute Module
873381-L21	3	HPE Synergy 480/660 Gen10 Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) FIO Processor Kit
873381-B21	3	HPE Synergy 480/660 Gen10 Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) Processor Kit
815097-B21	54	HPE 8GB (1x8GB) Single Rank x8 DDR4-2666 CAS-19-19-19 Registered Smart Memory Kit
815098-B21	18	HPE 16GB (1x16GB) Single Rank x4 DDR4-2666 CAS-19-19-19 Registered Smart Memory Kit
875478-B21	6	HPE 1.92TB SATA 6G Mixed Use SFF (2.5in) SC 3yr WTY Digitally Signed Firmware SSD
P01367-B1	3	HPE 96W Smart Storage Battery (up to 20 Devices) with 260mm Cable Kit
804424-B21	3	HPE Smart Array P204i-c SR Gen10 (4 Internal Lanes/1GB Cache) 12G SAS Modular Controller



Part number	Qty	Description
777430-B21	3	HPE Synergy 3820C 10/20Gb Converged Network Adapter
<b>Worker nodes</b>		
871943-B21	6	HPE Synergy 480 Gen10 6130 2P 64GB-R P204i-c SAS Performance Compute Module
873381-L21	6	HPE Synergy 480/660 Gen10 Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) FIO Processor Kit
873381-B21	6	HPE Synergy 480/660 Gen10 Intel Xeon-Gold 6130 (2.1GHz/16-core/125W) Processor Kit
815097-B21	108	HPE 8GB (1x8GB) Single Rank x8 DDR4-2666 CAS-19-19-19 Registered Smart Memory Kit
815098-B21	36	HPE 16GB (1x16GB) Single Rank x4 DDR4-2666 CAS-19-19-19 Registered Smart Memory Kit
P01367-B1	6	HPE 96W Smart Storage Battery (up to 20 Devices) with 260mm Cable Kit
804424-B21	6	HPE Smart Array P204i-c SR Gen10 (4 Internal Lanes/1GB Cache) 12G SAS Modular Controller
777430-B21	6	HPE Synergy 3820C 10/20Gb Converged Network Adapter
804428-B21	3	HPE Smart Array P416ie-m SR Gen10 Controller
		HPE Synergy D3940 Storage
835386-B21	3	HPE Synergy D3940 12Gb SAS CTO Drive Enclosure with 40 SFF (2.5 in.) Drive Bays
P06584-B21	18	HPE 960GB SAS 12G Read Intensive SFF (2.5in) SC 3yr Wty Digitally Signed Firmware SSD
757323-B21	3	HPE Synergy D3940 IO Adapter
<b>HPE Nimble Storage</b>		
Q8H41A	1	HPE Nimble Storage AF40 All Flash Dual Controller 10GBASE-T 2-port Configure-to-order Base Array
Q8B88B	1	HPE Nimble Storage 2x10GbE 2-port FIO Adapter Kit
Q8G27B	1	HPE Nimble Storage NOS Default FIO Software
Q8H47A	1	HPE Nimble Storage AF40 All Flash Array R2 11.52TB (24x480GB) FIO Flash Bundle
ROP84A	2	HPE Nimble Storage NEMA IEC 60320 C14 to C19 250V 15 Amp 1.8m FIO Power Cord
Q8F56A	1	HPE Nimble Storage 10GbE 2-port Spare Adapter
P9Q66A	2	HPE G2 IEC C20 Input/(8) C13 Expansion Outlets/PDU Extension Bar Kit
		HPE 3PAR Storage
Q2S13A	1	HPE 3PAR StoreServ RPS Service Processor
H6Y98B	1	HPE 3PAR 8440 4N+SW Storage Base
H6Z00A	6	HPE 3PAR 800 4-pt 16Gb FC Adapter
ROP66A	32	HPE 3PAR 8000 920GB SFF SDD+SW
P9S15A	2	HPE G2 Mtrd/Swtd 4.9kVA/L6-30P NA/J PDU
120672-b21	1	HPE Rack Ballast Kit
<b>HPE 5945 FlexFabric Switching</b>		
JQ075A	2	HPE FF 5945 2-Slot Switch
JH180A	2	HPE 5930 24p SFP+ and 2p QSFP+ Module
JH183A	2	HPE 5930 8-port QSFP+ Module
JH389A	4	HPE X712 Back (Power Side) to Front (Port Side) Airflow High Volume Fan Tray
JC680A	4	HPE 58x0AF 650W AC Power Supply
JC680A B2B	4	INCLUDED: Jumper Cable - NA/JP/TW
JG326A	2	HPE X240 40G QSFP+ QSFP+ 1m DAC Cable
JG327A	4	HPE X240 40G QSFP+ QSFP+ 3m DAC Cable





Part number	Qty	Description
<b>Red Hat OpenShift Container Platform</b>		
R1Z92AAE	1	Red Hat OpenShift Container Platform 4 for HPE Synergy 1-32 Cores 1yr Subscription 24x7
<b>Red Hat Enterprise Linux CoreOS (RHCOS) Server</b>		
J8J36AAE	6	Red Hat Enterprise Linux CoreOS (RHCOS) Server 2 Sockets 1 Guest 1 Year Subscription 24x7 Support
<b>Red Hat Enterprise Linux CoreOS (RHCOS) for Virtual data centers</b>		
G3J22AAE	11	Red Hat Enterprise Linux CoreOS (RHCOS) for Virtual Datacenters 2 Sockets 1 Year Subscription 24x7 Support

**NOTE**

Hewlett Packard Enterprise recommends you choose the appropriate storage systems, to build the solutions based on the requirement.



## RESOURCES AND ADDITIONAL LINKS

Red Hat, [redhat.com](https://www.redhat.com)

Red Hat OpenShift Container Platform 4 documentation, <https://docs.openshift.com/container-platform/4.4/welcome/index.html>

HPE Synergy, [hpe.com/info/synergy](https://www.hpe.com/info/synergy)

HPE Synergy D3940 Storage, <https://buy.hpe.com/us/en/synergy/synergy-storage/synergy-storage-modules/synergy-storage-modules/hpe-synergy-d3940-storage-module/p/1008615217>

HPE Nimble Storage, <https://www.hpe.com/us/en/storage/nimble.html>

HPE 3PAR Storage, <https://www.hpe.com/us/en/storage/3par.html>

HPE FlexFabric 5945 switching, [hpe.com/us/en/product-catalog/networking/networking-switches/pip.hpe-flexfabric-5945-switch-series.1009148840.html](https://www.hpe.com/us/en/product-catalog/networking/networking-switches/pip.hpe-flexfabric-5945-switch-series.1009148840.html)

To help us improve our documents, please provide feedback at [hpe.com/contact/feedback](https://www.hpe.com/contact/feedback).

---

© Copyright 2020-2022 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Red Hat is a registered trademark of Red Hat, Inc. in the United States and other countries. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. Intel and Xeon are trademarks of Intel Corporation in the U.S. and other countries. Microsoft and Windows are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. VMware vSphere is a registered trademark of VMware, Inc. in the United States and/or other jurisdictions. All third-party trademarks are the property of their respective owners. Redfish is a trademark of Distributed Management Task Force, Inc.