p +61 1300 113 112
e info@hyperscalers.com

Hyperconverged Red Hat OpenShift Container Platform with Data Foundation

Reference Guide

Hyperscalers with Red Hat

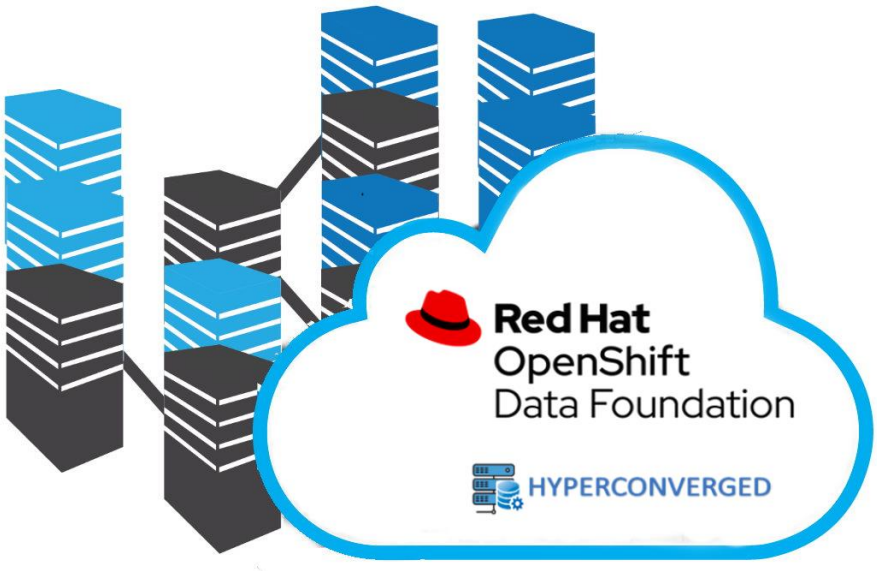

Monday, 20 March 2023

1 CONTENTS

2 INTRODUCTION

Red Hat® OpenShift® is an enterprise-ready Kubernetes container platform based on a unified architectural vision and supported by an open hybrid cloud strategy.

Using OpenShift, applications and the data centres that support them can expand safely and securely from just a few machine and application instances to thousands of instances that serve millions of clients.

Hyperscalers understands the need for Enterprise service providers and IT administrators to run and manage virtual machine (VM) and container workloads side by side on a single platform.

OpenShift can fulfill this requirement by allowing you to develop, manage, and deploy virtual machines side-by-side with containers and serverless across a common hyperconverged storage pool.

Hyperscalers has partnered with Red Hat to engineer and qualify a purpose-built, Hyperconverged Red Hat OpenShift Container Platform that incorporates all these features under a single hardware and software architecture that is jointly engineered and supported by Hyperscalers and Red Hat.

As a key part of this, the OpenShift Data Foundation supports a collection of on-demand storage and data service types that is tightly integrated within the OpenShift Platform.

The Hyperconverged OpenShift Platform enables you to leverage powerful storage technology breakthroughs that can deliver significant performance and/or capacity improvements such as you might require for your specific implementation.

These storage products have been tested within Hyperscalers rigorous laboratory environment for compatibility at all levels of the solution architecture, saving you from unexpected hardware, baseboard management, device driver and software stack issues.

Support for all components delivered within the Hyperconverged OpenShift Platform solution architecture is provided by Hyperscalers as part of initial solution delivery and during your ongoing support relationship with us.

Regardless of whether your primary focus is to leverage OpenShift continuous integration and continuous delivery (CI/CD) pipeline capabilities, and/or to support consistent management of hybrid cloud, multi-cloud, and edge deployments, the Hyperconverged Red Hat OpenShift Container Platform can provide you with the capacity, performance and availability capabilities and attributes that your organisation requires.

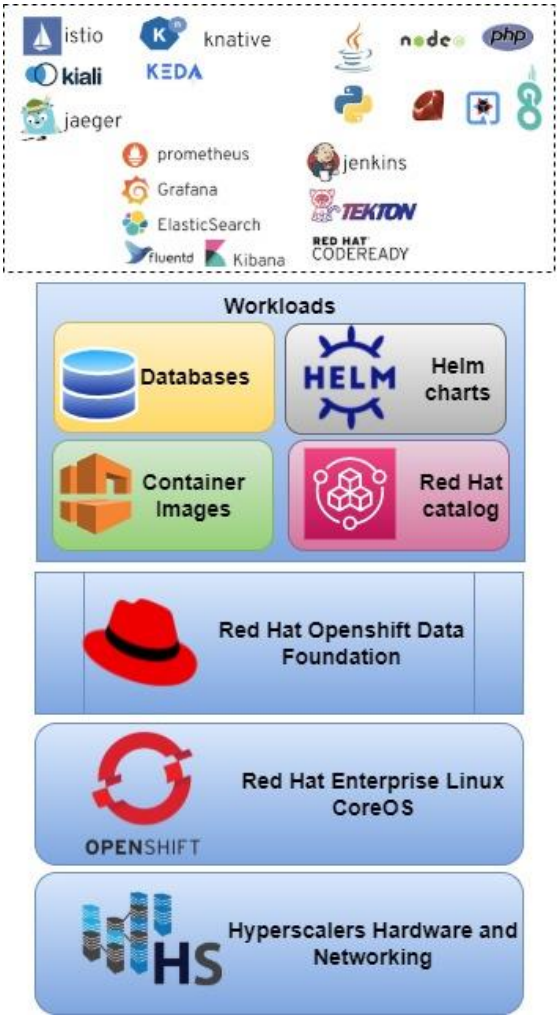

Figure 1 System Functional Black Diagram for Red Hat OpenShift with Data Foundation

Why Hyperscalers

Hyperscalers [1] is the world's first open supply chain Original Equipment Manufacturer- OEM, solving Information Technology challenges through standardization of best practices and hyperscale inspired practices and efficiencies. Hyperscalers offers choice across two open hardware architectures:

- Hyperscale - high efficiency open compute equipment as used by macro service providers
- Tier 1 Original – conventional equipment as per established Tier 1 OEM suppliers.

Each architecture is complete with network, compute, storage, and converged GP GPU infrastructure elements, and is open / free from vendor lock-in.

Hyperscalers' appliance solutions are packaged complete with hardware, software and pre-built (customisable) configurations. These were all pre-engineered using an in-house IP Appliance Design Process and validated in partnership with associated major software manufacturers. Many can be "test-driven" using Hyperscalers Lab as a Service (LaaS). Hyperscalers appliance solutions are ideally suited to IaaS PaaS and SaaS providers looking to implement their services from anywhere.

The Red Hat OpenShift Appliance by Hyperscalers is a complete package including high performance CPU, memory, and network resources coupled with highly sophisticated hyperconverged Red Hat Data foundation to provide persistent container storage to the enterprise IT workloads. The detailed deployment steps including Hyperscalers IP appliance design process is described in this document.

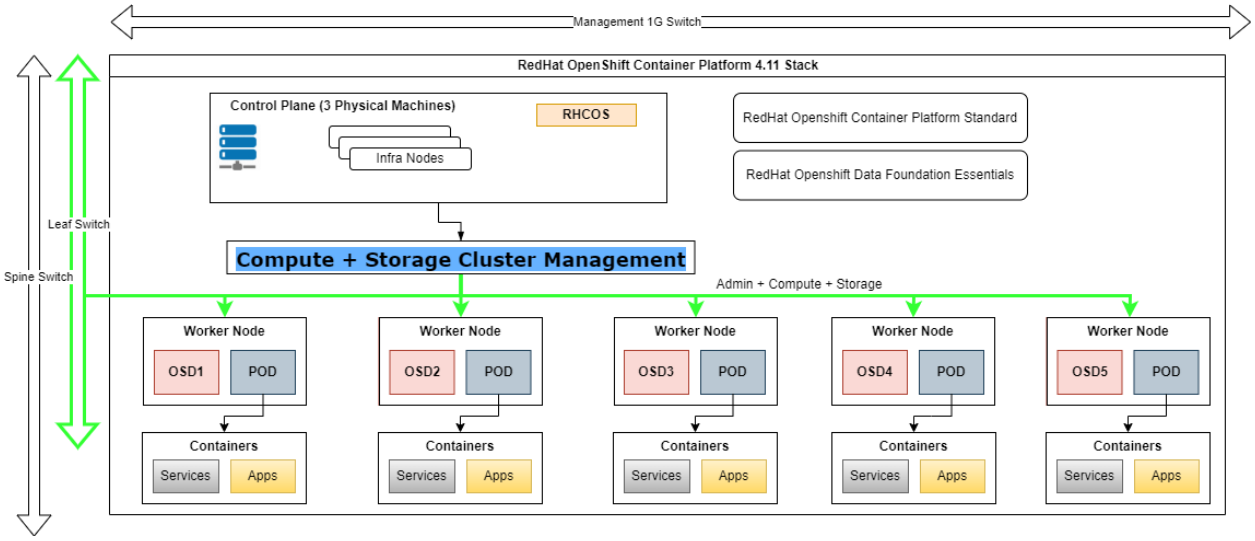

Figure 2 System Architecture for Red Hat OpenShift

Audience and Purpose

Engineers, Enthusiasts, Executives, and IT professionals with background in Computer Science/ Electronics/ Information Technology with understanding in Linux commands, Python language and basic electronics who intend to study, explore, deploy Red Hat OpenShift cluster with hyperconverged Red Hat Data Foundation platform can be benefitted from this reference guide.

Documents, Knowledge Base, and Technical Support

Hyperscalers reference architectures and appliance / solutions demonstrations are available at: https://www.hyperscalers.com/OCP-hyperscale-rack-solutions

For technical queries regarding this document and for managing virtualized, mobile, and cloud technologies, you can contact Hyperscalers technical support at support@hyperscalers.com

Red Hat provides a variety of specialized, supplementary offerings to help you achieve operational excellence with your enterprise solutions. With a Red Hat subscription, your mission-critical work is backed by experienced, knowledgeable, and committed support engineers. More documentation on the general setup and configuration of Red Hat OpenShift can be found at https://access.redhat.com/documentation/en-us/openshift_container_platform/4.11

Additional reference to the Red Hat OpenShift on open rack platforms (OCP) can be found in Hyperscalers whitepapers and reference architecture section link - https://www.hyperscalers.com/red-hat-ceph-openstack-13-world-record-saas-paas-openshift-paas-appliance-open-compute-ocp-hyperscale-how-to-build-cloud-world-record

Readers are recommended to have a prior knowledge and expertise with Kubernetes, Ceph and Linux programming to better understand the following documentation.

Contact info@hyperscalers.com for more information.

Features of Hyperconverged Red Hat OpenShift Container Platform with Data Foundation

Red Hat OpenShift is a comprehensive platform inclusive of core Kubernetes, Linux Kernel based Virtual Machine (KVM) and highly resilient Ceph storage technology. This makes Red Hat OpenShift production grade with not just a platform to run apps but a complete packaging of authentication, networking, security, monitoring, logs management, etc. Red Hat OpenShift integrates all the features of Kubernetes and provides add-on features to the platform as listed below[1]:

1. A trusted OS foundation: RHEL CoreOS or RHEL
2. Automated Operations
3. Developer Services
4. Application Services
5. Cluster Services

Red Hat OpenShift manages hybrid technologies and applications, helping you modernize existing applications and accelerate new cloud-native application development and delivery at scale across any infrastructure.

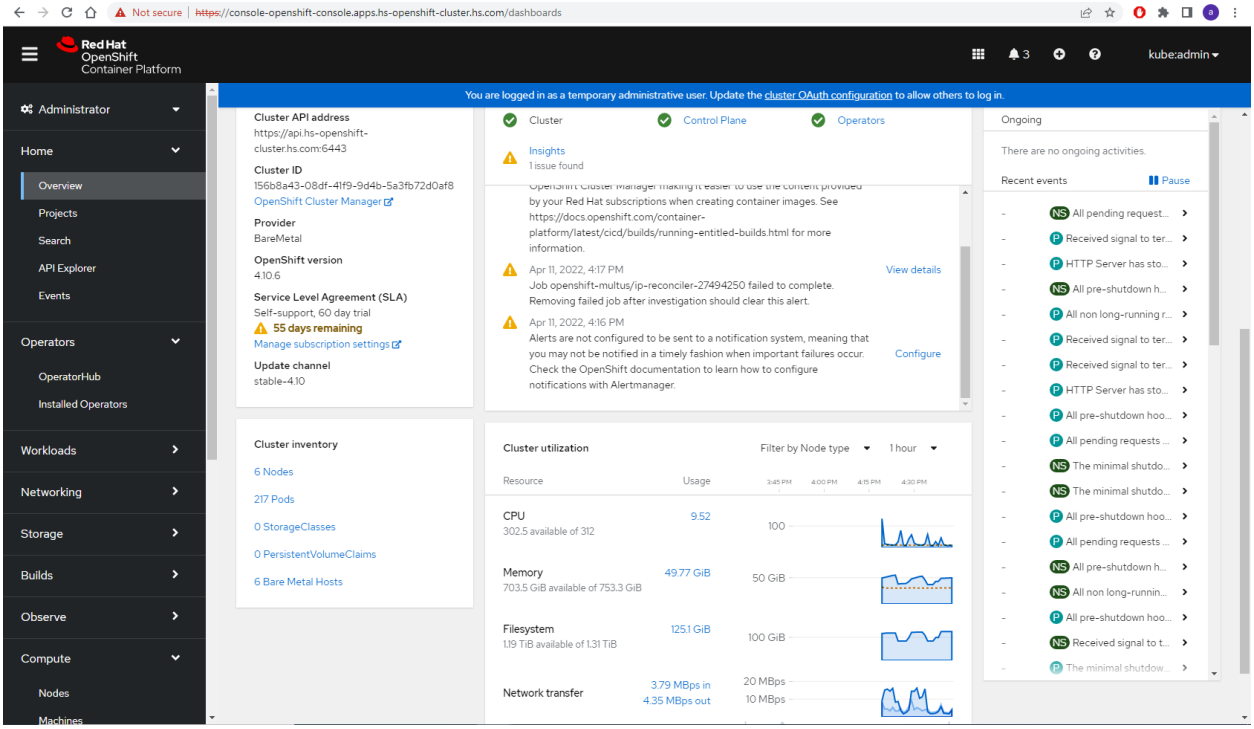

Figure 3 Red Hat OpenShift Dashboard

Benefits of Red Hat OpenShift[2] :

Scalability

Apps running on Red Hat OpenShift can scale to thousands of instances across hundreds of nodes in seconds.

Flexibility:

Red Hat OpenShift simplifies deployment and management of a hybrid infrastructure, giving you the flexibility to have a self-managed or fully managed service, running on-premises or in cloud and hybrid environments.

Open-source standards

Red Hat OpenShift incorporates Open Container Initiative (OCI) containers and Cloud Native Computing Foundation-certified Kubernetes for container orchestration, in addition to other open-source technologies.

Container portability

Container images built on the OCI industry standard ensure portability between developer workstations and Red Hat OpenShift production environments.

Enhanced developer experience

Red Hat OpenShift offers a comprehensive set of developer tools, multilanguage support, and command line and integrated development environment (IDE) integrations. Features include continuous integration/continuous delivery (CI/CD) pipelines based on Tekton and third-party CI/CD solutions, service mesh, serverless capabilities, and monitoring and logging capabilities.

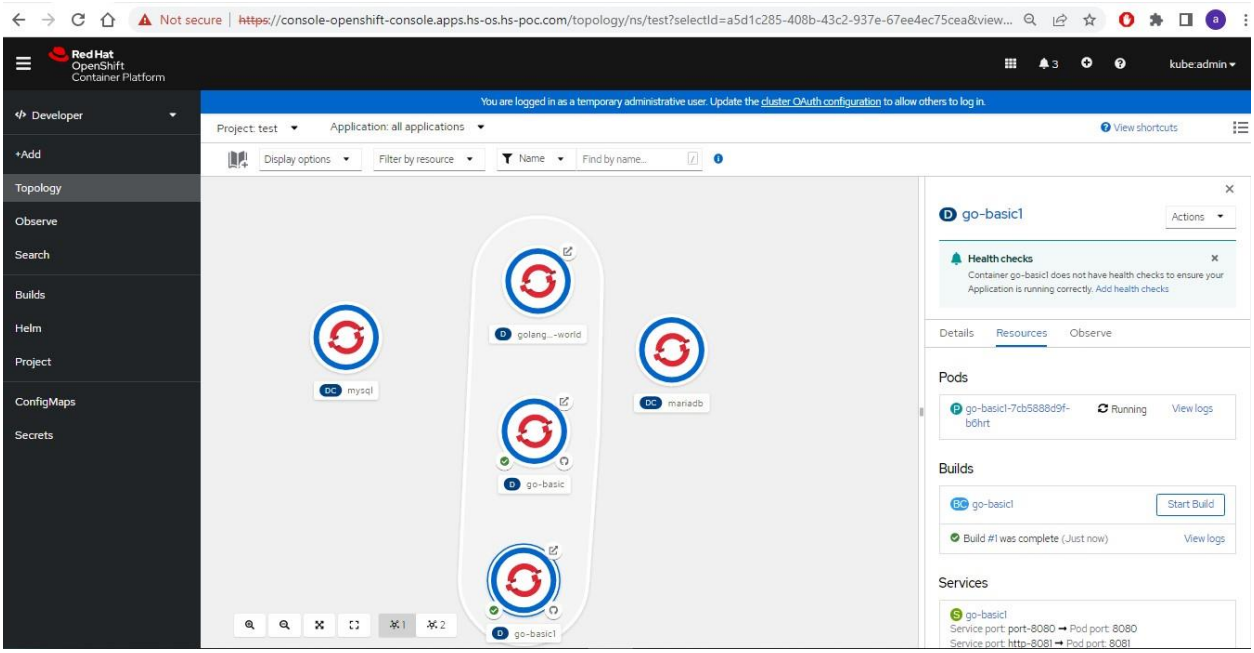

Figure 4 Apps deployed with ease

Automated installation and upgrades

Automated installation and over-the-air platform upgrades are supported in cloud with Amazon Web Services, Google Cloud Platform, IBM Cloud, and Microsoft Azure, and on-premises using vSphere, Red Hat OpenStack Platform, Red Hat Virtualization, or bare metal. Services used from the Operator Hub can be deployed fully configured and are upgradable with 1 click.

Automation

Streamlined and automated container and app builds, deployments, scaling, health management, and more are included.

Edge architecture support

Red Hat OpenShift enhances support of smaller-footprint topologies in edge scenarios that include 3-node clusters, single-node Red Hat OpenShift, and remote worker nodes, which better map to varying physical size, connectivity, and availability requirements of different edge sites. The edge use cases are further enhanced with support for Red Hat OpenShift clusters on ARM architecture, commonly used for low-power-consumption devices.

Multi-cluster management

Red Hat OpenShift with Red Hat Advanced Cluster Management for Kubernetes can easily deploy apps, manage multiple clusters, and enforce policies across clusters at scale.

Advanced security and compliance

Red Hat OpenShift offers core security capabilities like access controls, networking, and enterprise registry with built-in scanner. Red Hat Advanced Cluster Security for Kubernetes enhances this with security capabilities like runtime threat detection, full life cycle vulnerability management, and risk profiling.

Persistent storage

Red Hat OpenShift supports a broad spectrum of enterprise storage solutions, including Red Hat OpenShift Data Foundation, for running both stateful and stateless apps.

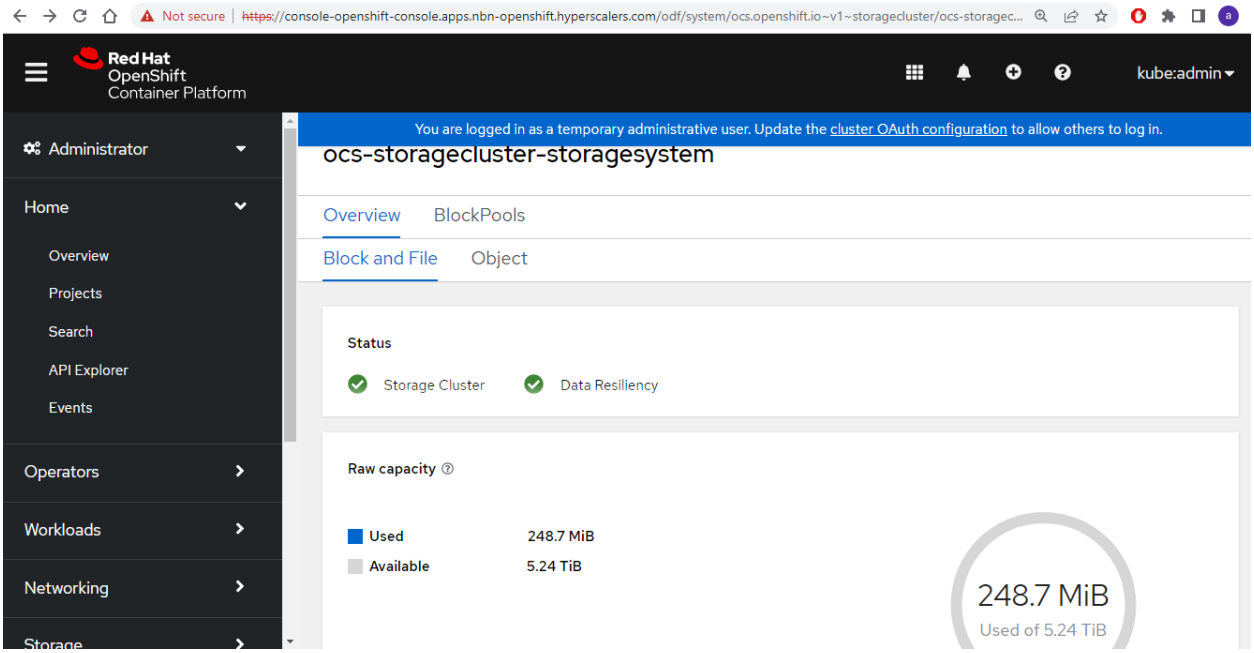

Figure 5 Block, File and Object resilient data storage

Robust ecosystem

An expanding ecosystem of partners provides a wide variety of integrations. Third parties deliver additional storage and network providers, IDE, CI, integrations, independent software vendor solutions, and more.

Power of Kubernetes

Within OpenShift Container Platform, Kubernetes manages containerized applications across a set of containers or hosts and provides mechanisms for deployment, maintenance, and application-scaling. The container runtime packages, instantiates, and runs containerized applications [3].

Install Operators:

Red Hat OpenShift platform provides several operators that are pre-engineered to perform specific applications like Elastic Search, Kafka, SSL certificate management etc. Please note that a subscription to individual services may be required to use the applications in Red Hat OpenShift environment.

OpenShift Data Foundation

Ceph based persistent storage previously called Red Hat OpenShift Container Storage—is software-defined storage for containers. It provides cluster data management capabilities that lets organizations deploy their apps and data management as needs dictate, and then adjust as they move forward.

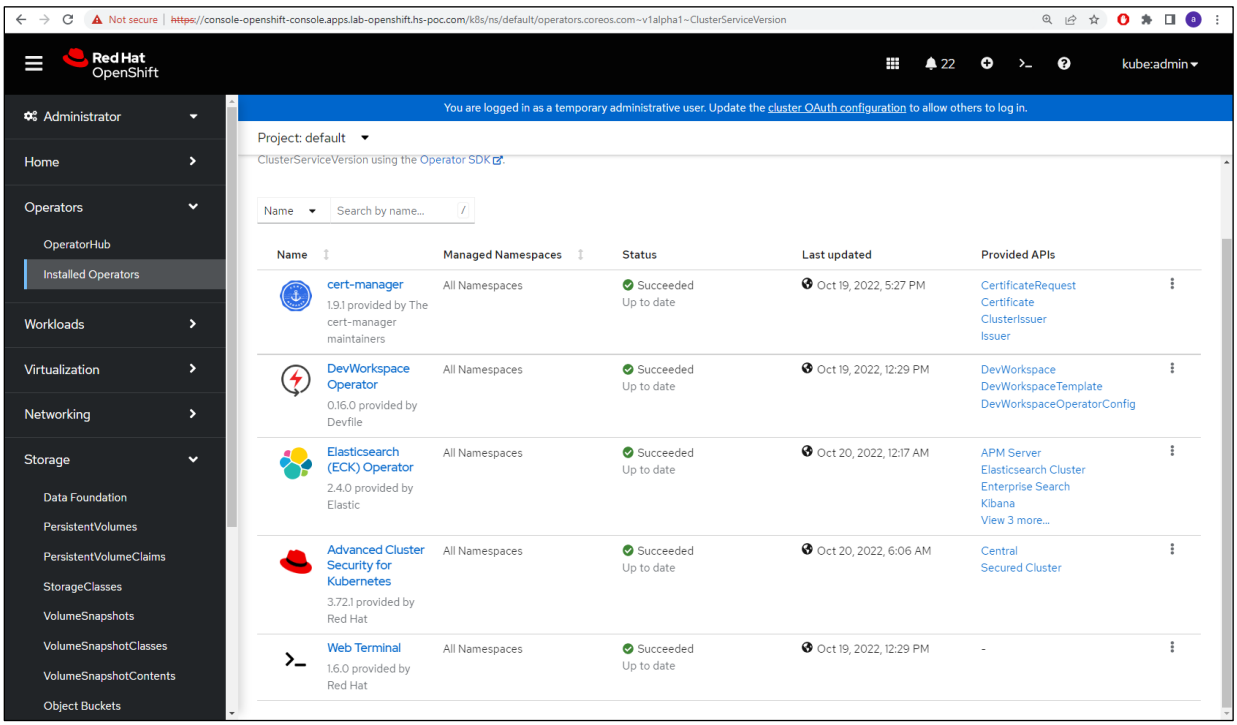

Figure 6 Red Hat OpenShift Operators

Persistent volumes, Claims and Storage classes:

In a containerized environment, the storage is classified as Storage class which is consumed by the persistent volumes using the persistent volume claims. Each of the persistent volume claim is associated with an application which decides the size of the storage allocation to them. As mentioned earlier, Ceph is the underlying storage technology for the Red Hat OpenShift that can provide a block, filesystem, and object storage classes for the application to use from. Below are some of the screenshots that show the storage platform provided by the Red Hat Data Foundation.

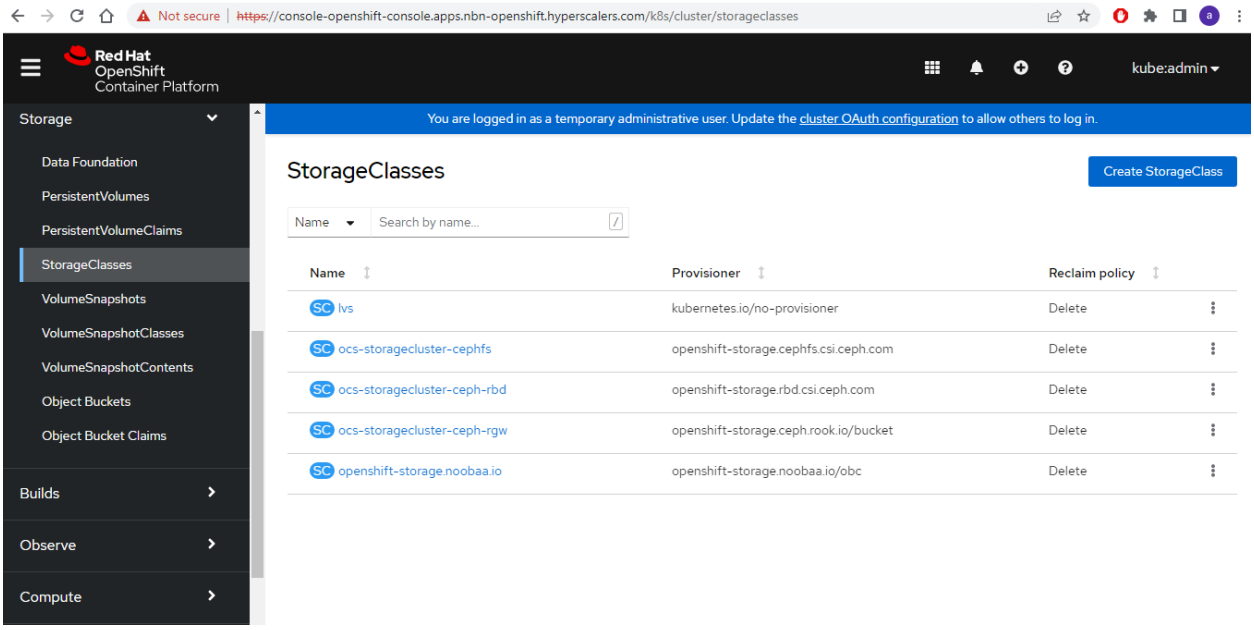

Figure 7 OpenShift Block, File and Object storage classes

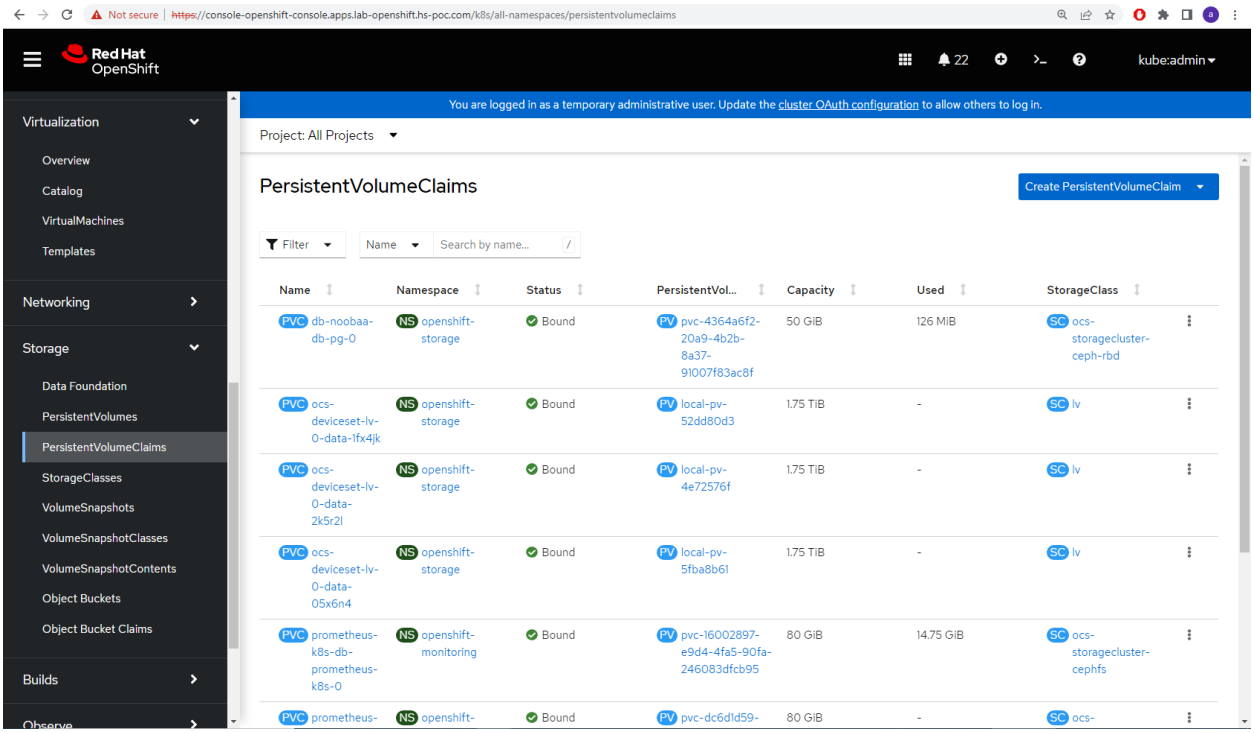

Figure 8 Red Hat Data Foundation Persistent Volume Claims

Pods:

Pods are the basic unit of the containerized resource where an application resides. Pods can be replicated for high availability and the route to the application is defined in the pods. Every pod is associated with a namespace or project, and they can be accessed via SSH.

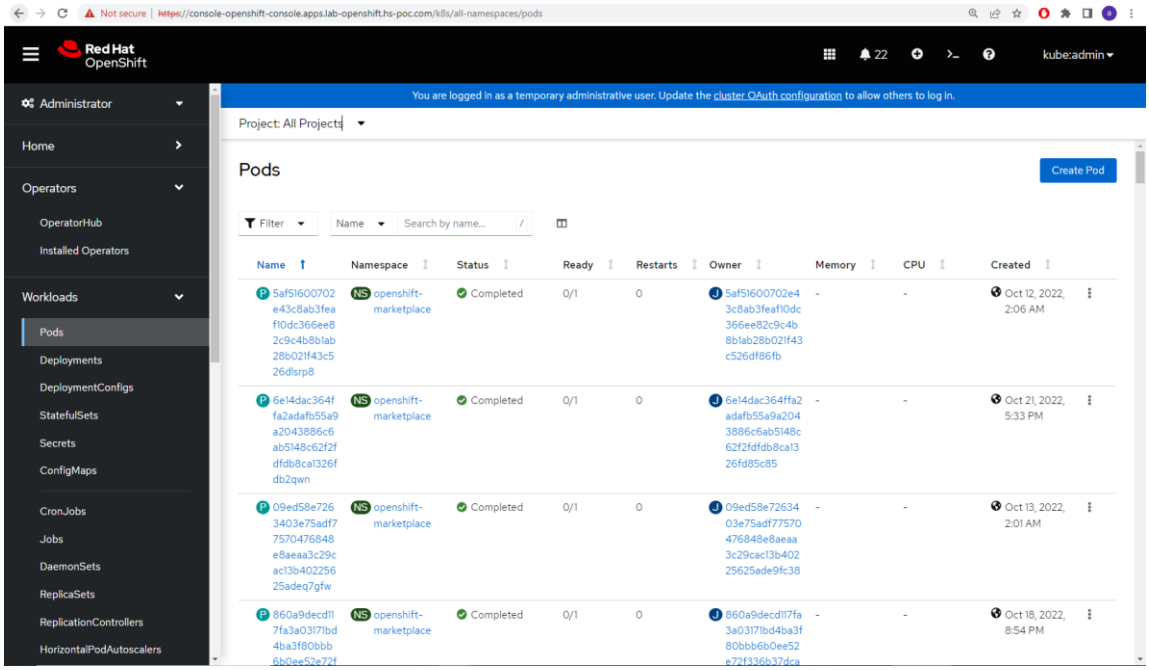

Figure 9 Pods in Red Hat OpenShift

Virtualization:

Red Hat Virtualisation is enabled on the OpenShift using the Virtualization Operator. This enables the user to have the containers alongside the virtual machines to provide adaptability for application that are built for Kubernetes containers and those apps built for virtual machines. Hence, the same physical resource can used to deploy containerised as well as virtual machine-based workloads.

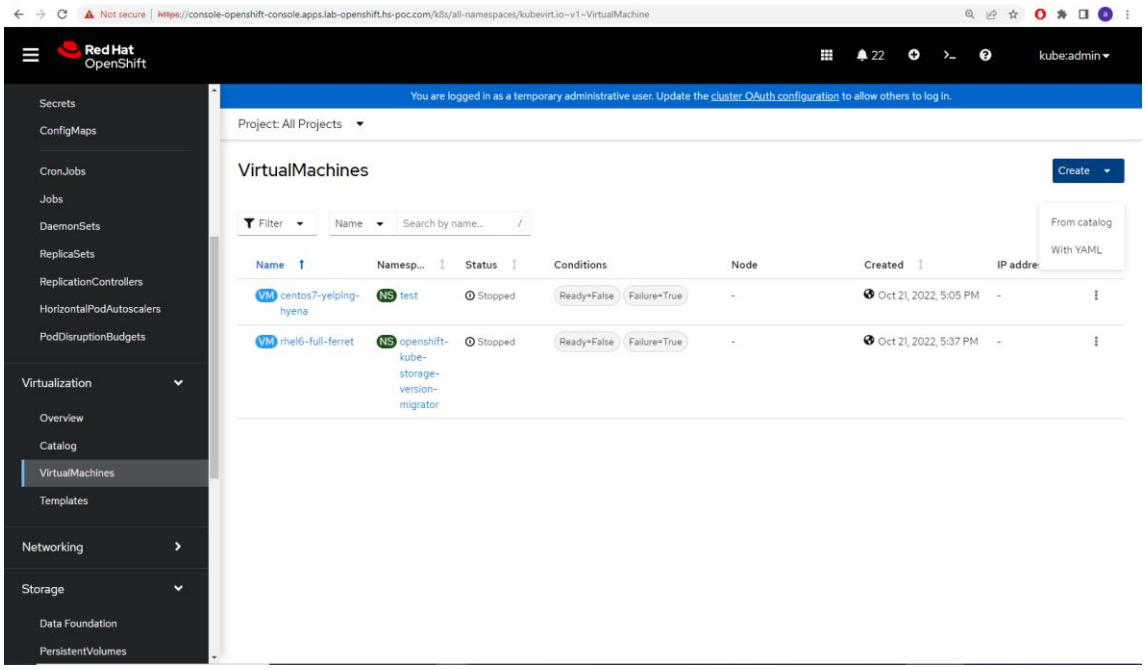

Figure 10 Virtualization in Red Hat OpenShift

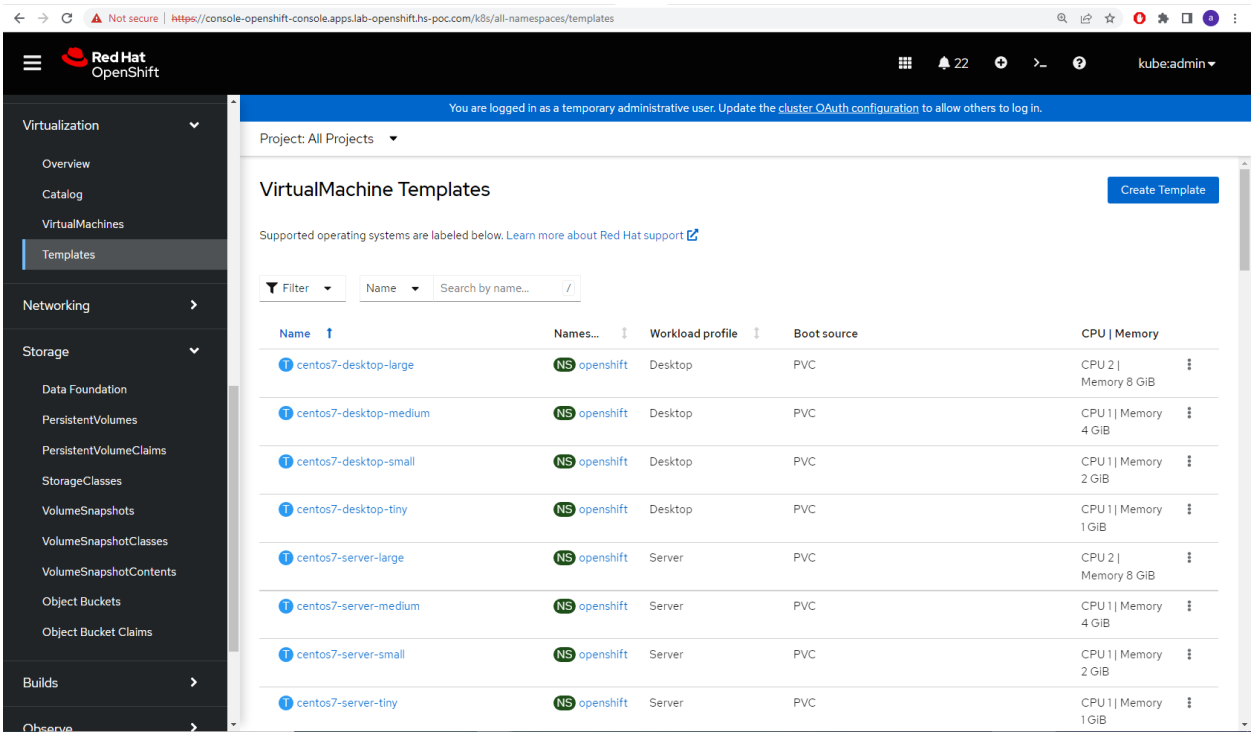

Figure 11 Virtual Machine templates in Red Hat OpenShift

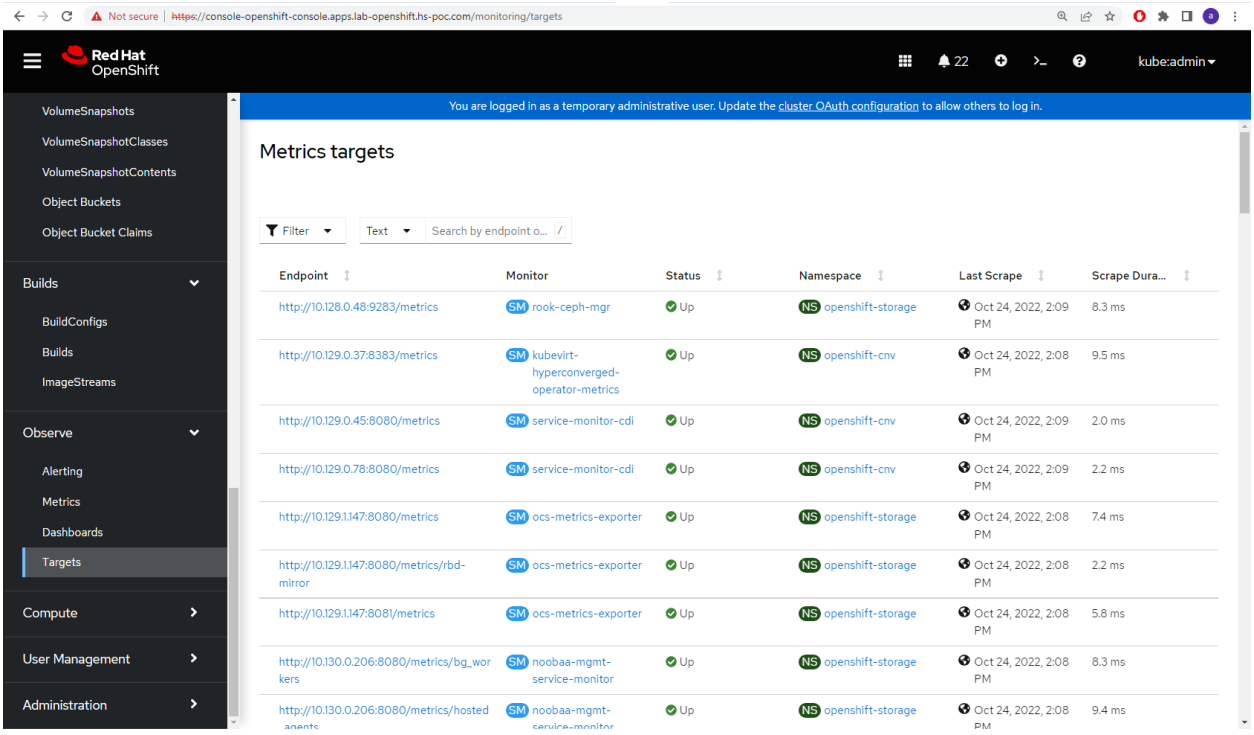

Figure 12 Access Endpoints to Monitoring and Health Metrics

Cluster Upgrades and Subscription:

Red Hat OpenShift provides a 60-day free evaluation to setup and support your cluster deployment and then we can attach a license/subscription based on purchase period for the support and services. Cluster upgrade can run on the live environment without affecting the workload, but it is highly recommended to understand the requirements of your workloads before upgrading as it could cause inter-operability issues within the apps in the container.

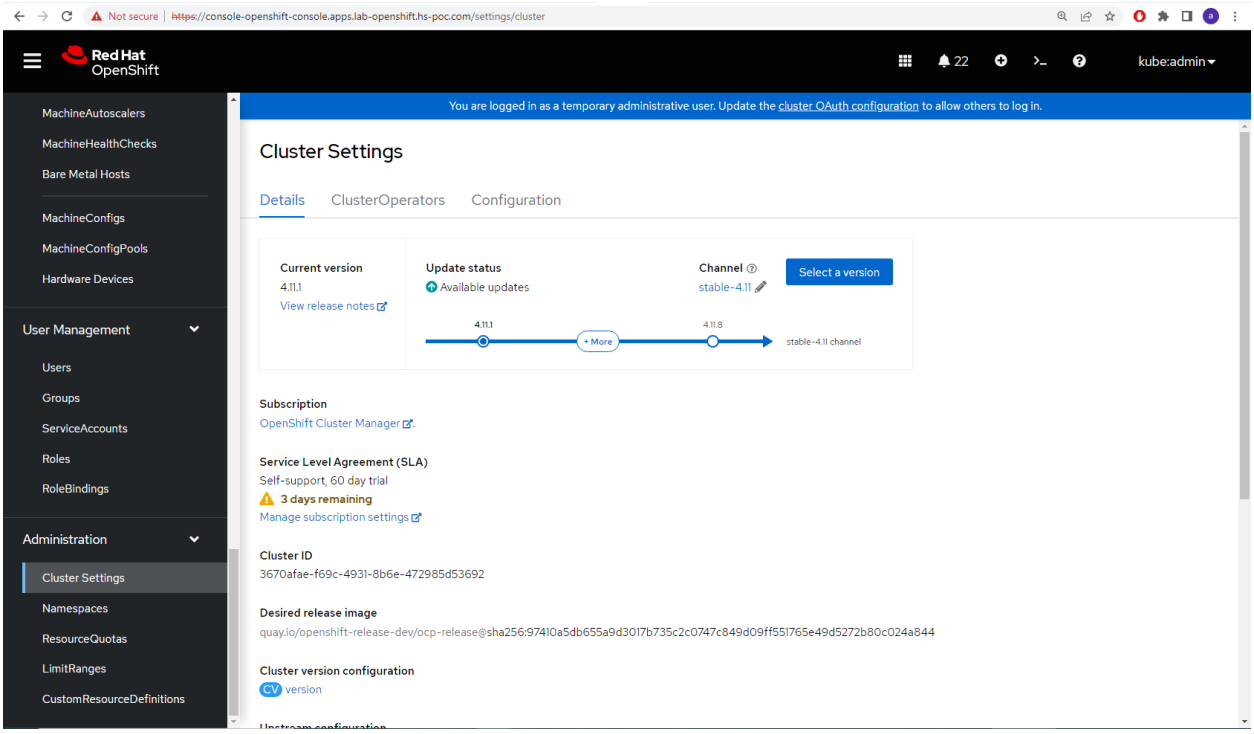

Figure 13 Red Hat OpenShift Cluster maintenance and upgrade

Important Considerations

The following documentation gives a detailed step-by-step deployment guide of Red Hat OpenShift along with Red Hat Virtualisation in a hyperconverged Red Hat Data foundation storage setup. The architecture is specific and tailored for enabling enterprises to provision a platform for their service delivery within or outside the organisation. Hyperscalers recommends the below important considerations before proceeding to the deployment phase.

1. Red Hat OpenShift requires minimum 3 master nodes to be deployed and 1 worker node for compute. These nodes can be VMs or Bare Metal server infrastructure (Whichever you choose, keep the same infrastructure for all the nodes).
2. Keep the hardware configuration consistent across all the nodes to ensure replication and high availability.
3. BIOS needs to disable secure boot to install Red Hat Core OS using Rufus flashed USB.
4. All our nodes are in high performance power profile setting from the BIOS.
5. Allocate a specific subnet and IP range on the same VLAN for all the nodes intended to be clustered.
6. Open the below ports for subnet / DHCP from your firewall

   4789 UDP, 443 TCP, 8443 TCP, 10250 TCP, 10010 TCP, 2049 TCP/UDP,

   2379 TCP, 2380 TCP, 9000 TCP, 8444 TCP, 22 TCP, 53 TCP/UDP,

   80 TCP, 8053 TCP/UDP, 1936 TCP, 2379 TCP, 2380 TCP, 4789 UDP,

   9200 TCP, 9300 TCP, 9090 TCP, 9100 TCP.

7. Ensure to have a DNS server in the same network since you must create "A" hosts for the API and APPS later during deployment and configuration.
8. Register and ensure to acquire a subscription from Red Hat portal https://cloud.redhat.com/

Digital IP Appliance Design Process

Hyperscalers has developed a Digital- IP-Appliance Design Process and associated Appliance Optimizer Utility which can enable the productization of IT-appliances for Digital-IP owners needing to hyperscale their services very quickly, reliably and at a fraction of traditional costs.

Appliance Optimizer Utility AOU

The Appliance Optimizer Utility (AOU) automates the discovery of appliance bottlenecks by pinging all layers in the proposed solution stack. A live dashboard unifies all key performance characteristics to provide a head-to-head performance assessment between all data-path layers in the appliance, as well as a comparison between holistic appliances.

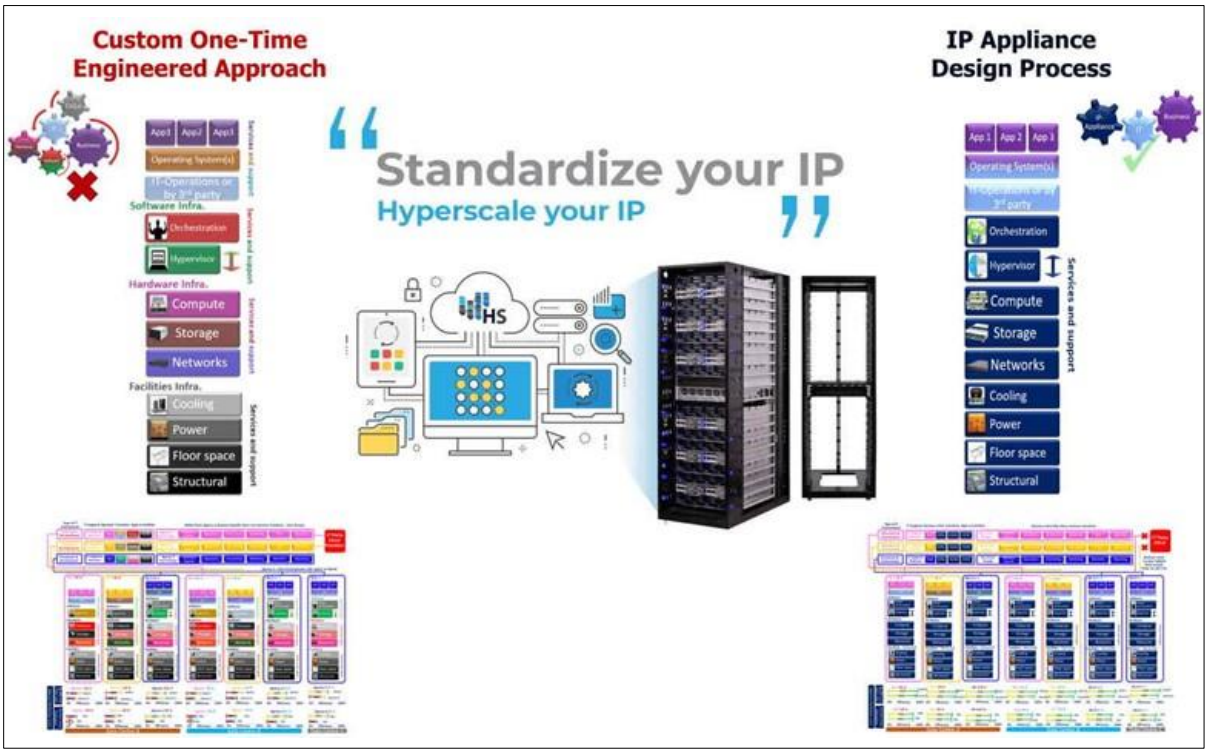

Figure 14 Digital IP-Appliance Design Process

Infrastructure Setup

To demonstrate a scalable and resilient Red Hat OpenShift cluster, we have used 3 master nodes and 6 worker nodes with 4x local storage drives per worker nodes. The hardware configuration for the build is listed below:

Master nodes:

2x Intel® Xeon® Silver 4310T Processor 10c 3.40 GHz 2.30 GHz 15 MB

8x DDR4 25600 (3200Mhz) 32GB Register Samsung M393A4K40DB3-CWE

1x NIC OCP(Pull Tab) 25Gb SFP28 2 port MELLANOX CX631432A ConnectX-6 Lx PCI-E X 8 Gen 4 MCX631432AN-ADAB

2x OS SSD - M.2 PCIe SSD 250GB Samsung MZ-V7S250 (For OS 250 GB DWPD < 1)

Worker nodes:

2x CPU - Intel® Xeon® Gold 5320 Processor 26c 3.40 GHz 2.20 GHz 39 MB 185W

8x DDR4 25600 (3200Mhz) 32GB Register Samsung M393A4K40DB3-CWE

1x NIC OCP(Pull Tab) 25Gb SFP28 2 port MELLANOX CX631432A ConnectX-6 Lx PCI-E X 8 Gen 4 MCX631432AN-ADAB

2x OS SSD - M.2 PCIe SSD 250GB Samsung MZ-V7S250 (For OS 250 GB DWPD < 1)

4x NVMe 2.5" U.2 NVMe 15mm PCI-e Gen4 X4 15.4TB Samsung PM1733 MZWLR15THALA-00007(O)

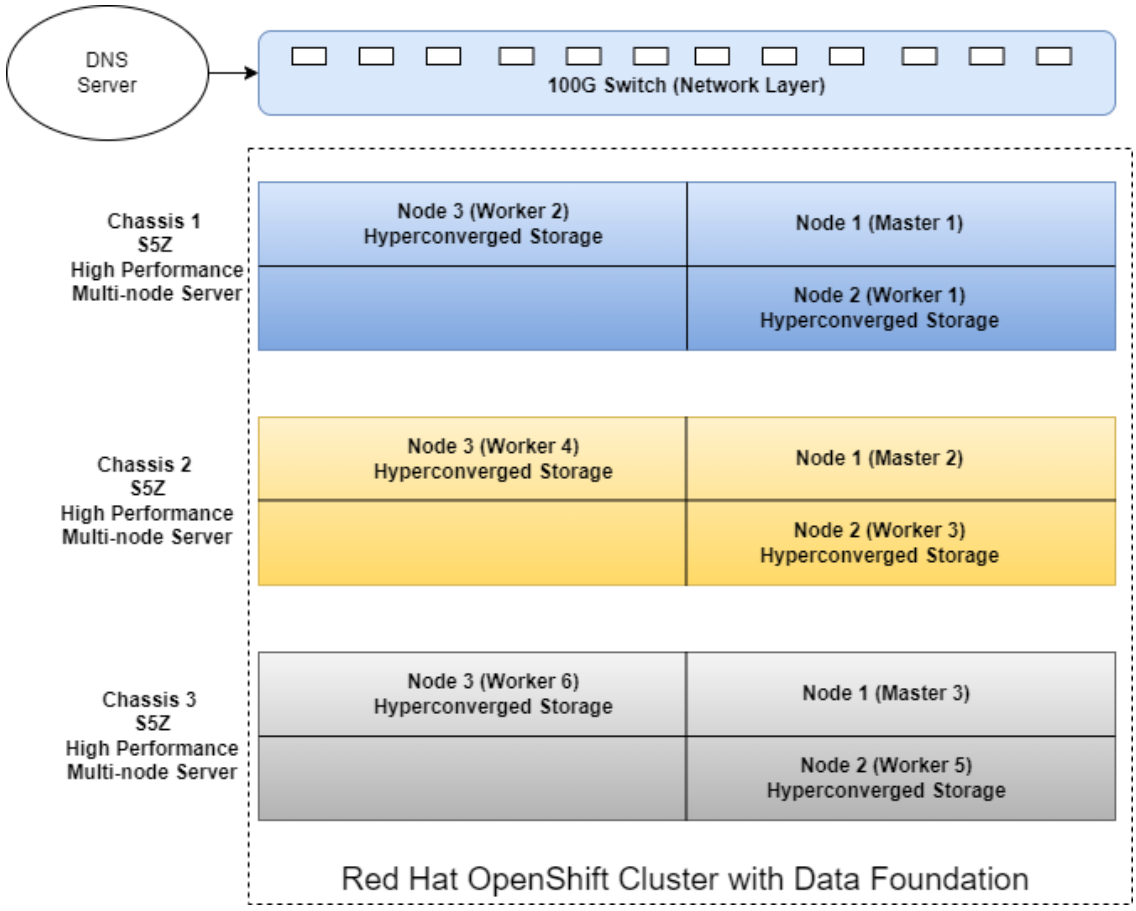

Figure 15 Physical Architecture of Red Hat OpenShift with Data Foundation

Software Stack:

Operating System: Red Hat Enterprise Linux CoreOS (RHCOS) v4.11

Cluster Deployment tool: Red Hat Hybrid Cloud Console

Hyperconverged Storage Operator: OpenShift Data Foundation v4.11.2

Virtualisation: OpenShift Virtualization v4.11.0

Building Blocks:

S5X 2.5" | D53X-1U Ultimate 1U Server for Intel Xeon 3rd Gen Processors

The S5X 2.5" (D53X-1U) based on PCIe Gen 4.0 and Intel's 3rd Generation Processor Family (Ice-lake) offers: Two (2) CPU Sockets for up to 80 cores using Intel® Xeon® Platinum 8380 Processor 40cores each. 32 Memory slots for up to 8TB DIMM or Up to 12TB DIMM+DCPM (PMEM 200 series). 12 Front Storage drive bays 2.5" hot-plug U.2 NVMe or SATA/SAS. Five (5) x PCIe 4.0 expansions slots for Network Interface Cards NIC. Two (2) M.2 onboard storage. Three (3) accelerators like NVIDIA T4 GPU.

S5K | D43K-1U Ultimate 1U Server for AMD EPYC Milan 3rd Gen Processors

Native design for AMD EPYC™ 7003 Processors, ready for PCIe 4.0 eco-system deployment. Up to 128-core within 1U form factor, optimized for HPC workloads. With 4 AMD xGMI-2 between dual EPYC™ processors up to 16GT/sec of CPU interconnect speed. Up to 5 PCIe expansion slots in a 1U chassis. Flexible I/O options with a variety of SAS mezzanine and OCP mezzanine option for diverse configurations. Flexible storage configurations, tailored for diversified software defined workloads. NUMA balanced PCIe topology for NVMe drives.

S5Z | T43Z-2U The Power of Hyper Convergence

The S5Z | T43Z-2U based on PCIe Gen 4.0 and Intel's 3rd Generation Processor Family (Ice-lake) is a high performance, multi node server offering eight (8) CPU in 2RU as part of four (4) independent nodes. Each node offers two (2) CPU Sockets for up to 80 cores using Intel® Xeon® Platinum 8380 Processor 40cores each, 16 Memory slots for up to 4TB DIMM or up to 6TB DIMM+DCPM (PMEM 200 series), four (4) 2.5" U.2 NVMe front storage drive

bays with two (2) M.2 NVMe for OS or caching, and three (3) x PCIe 4.0 expansions slots for Network Interface Cards NIC or accelerators like GPU.

S9CA | S43CA-2U AMD Density Optimized "EPYC" Multi-node Server

The QuantaPlex S43CA-2U is a multi-node server that supports the next generation of powerful AMD EPYC™ processors. Each of the 4 nodes in this compact 2U chassis is capable of supporting a single-P top bin Rome CPU, boasting a dominant 64-cores while still providing 16 DIMM slots that meets the most intense computing environment needs.

Access and Default Credentials

To access the Hyperscalers Lab as a Service (LaaS) portal, navigate to https://www.hyperscale2.com which is a repository of enterprise appliances that can be used to test drive the use cases before deploying on a mass scale.

Red Hat OpenShift with Data Foundation can be accessed from https://www.console-openshift-console.apps.openshift.hyperscale2.com (Please request for the credentials to info@hyperscalers.com and we can assist you.)

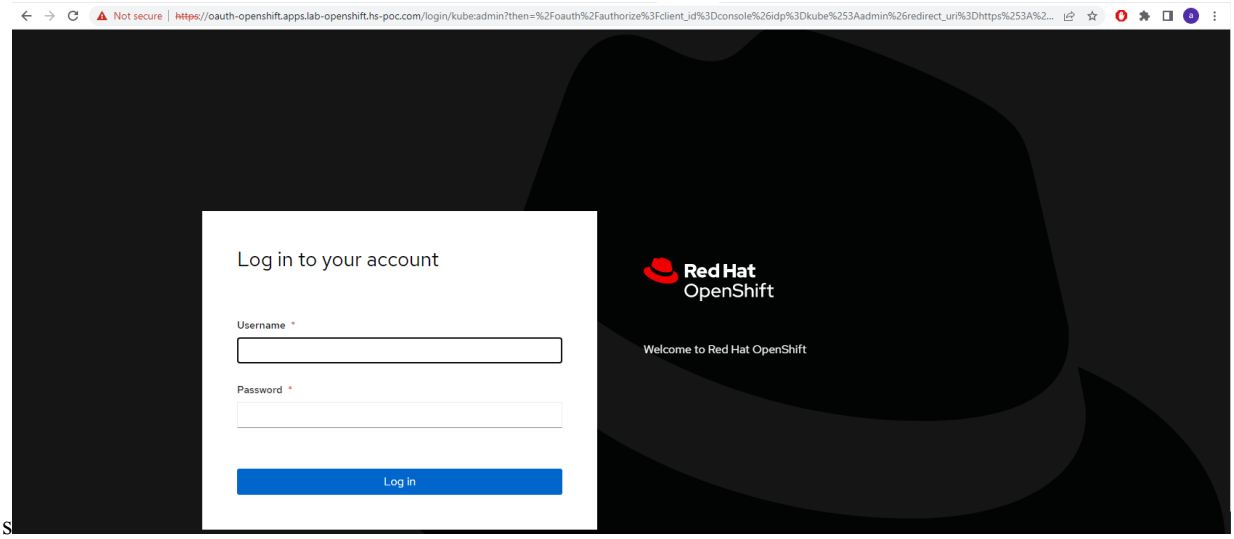

S

Figure 16 Red Hat OpenShift Login Dashboard

Terminologies

Hyperconvergence [3] - Hyperconvergence is an IT framework that combines storage, computing, and networking into a single system in an effort to reduce data centre complexity and increase scalability.

Ceph[4] - Ceph is an open-source software-defined storage platform that implements object storage on a single distributed computer cluster and provides 3-in-1 interfaces for object-, block- and file-level storage.

Containerisation [5]- Containerization is operating system-level virtualization or application-level virtualization over multiple network resources so that software applications can run in isolated user spaces called containers in any cloud or non-cloud environment, regardless of type or vendor.

Pods [6]- OpenShift Online leverages the Kubernetes concept of a pod, which is one or more containers deployed together on one host, and the smallest compute unit that can be defined, deployed, and managed. Pods are the rough equivalent of a machine instance (physical or virtual) to a container.

Master and Worker [7]- A Kubernetes cluster consists of a set of worker machines, called nodes, that run containerized applications. Every cluster has at least one worker node. The control plane called as master manages the worker nodes and the Pods in the cluster.

DNS [8]- The Domain Name System is the hierarchical and decentralized naming system used to identify computers reachable through the Internet or other Internet Protocol networks. Kubernetes clusters should have cluster Domain Naming Server (DNS). Cluster DNS is a DNS server, in addition to the other DNS server(s) in your environment, which serves DNS records for Kubernetes services. Containers started by Kubernetes automatically include this DNS server in their DNS searches.

Monitor - A Ceph Monitor (ceph-mon) maintains maps of the cluster state, including the monitor map, manager map, the OSD map, the MDS map, and the CRUSH map. These maps are critical cluster state required for Ceph daemons to coordinate with each other. Monitors are also responsible for managing authentication between daemons and clients. At least three monitors are normally required for redundancy and high availability.

Ceph OSDs [9] - An Object Storage Daemon (Ceph OSD, ceph-osd) stores data, handles data replication, recovery, rebalancing, and provides some monitoring information to Ceph Monitors and Managers by checking other Ceph OSD Daemons for a heartbeat. At least three Ceph OSDs are normally required for redundancy and high availability.

3 BASE PRODUCT DEPLOYMENT

Red Hat OpenShift is deployed using the Red Hat CoreOS ISO downloaded from OpenShift Cluster Manager-2022. OpenShift Cluster Manager provides links and steps to install Red Hat OpenShift Container Platform clusters and tools to create Red Hat OpenShift Dedicated and Red Hat OpenShift Service. OpenShift Container Platform is a self-managed hybrid cloud platform. With OpenShift Container Platform, you can create your clusters on any private or public cloud or bare metal, using your own infrastructure.

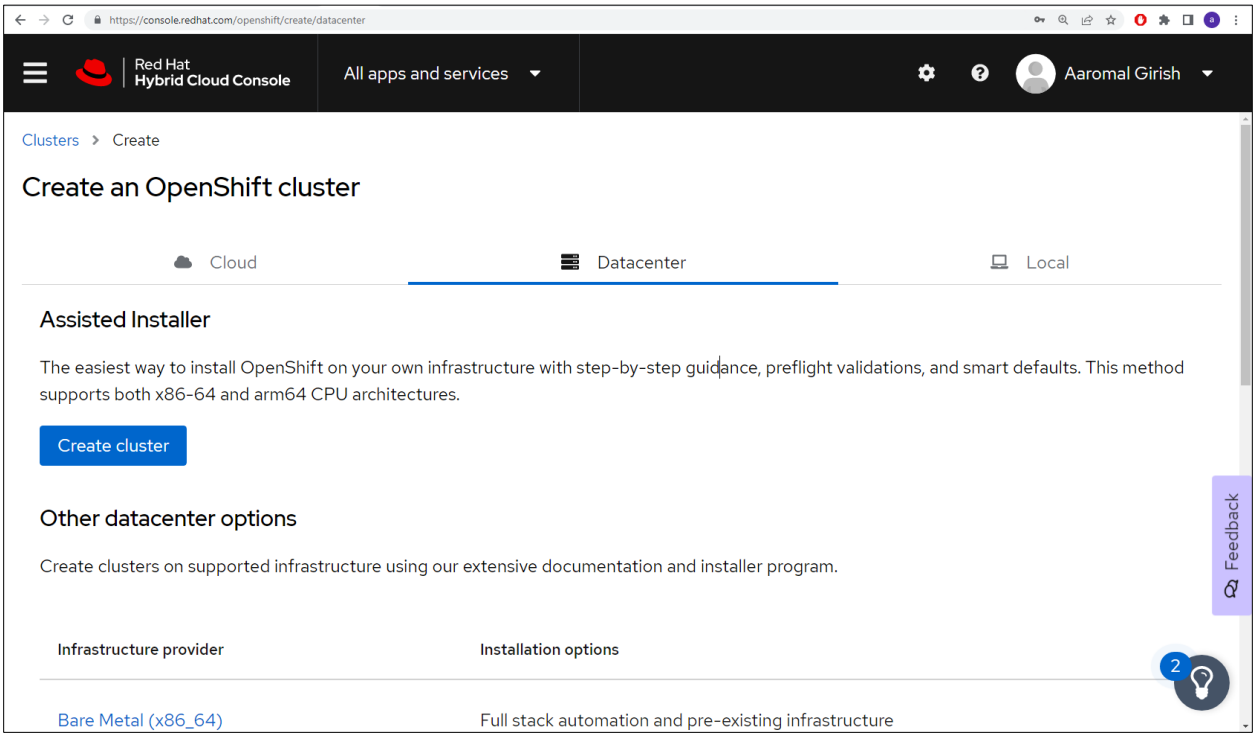

Figure 17 Red Hat Hybrid Cloud Console Assisted Installer

Preinstallation Requirements

- **Production level hardware requirements**

For production environments, the following recommendations apply:

  - **Master hosts -** In a highly available OpenShift Container Platform cluster with external etcd, a master host needs to meet the minimum requirements and have 1 CPU core and 1.5 GB of memory for each 1000 pods. Therefore, the recommended size of a master host in an OpenShift Container Platform cluster of 2000 pods is the minimum requirements of 2 CPU cores and 16 GB of RAM, plus 2 CPU cores and 3 GB of RAM, totalling 4 CPU cores and 19 GB of RAM.
  - **Worker hosts -** The size of a node host depends on the expected size of its workload. As an OpenShift Container Platform cluster administrator, you need to calculate the expected workload and add about 10 percent for overhead. For production environments, allocate enough resources so that a node host failure does not affect your maximum capacity.

- **DNS Requirements**

OpenShift Container Platform requires a fully functional DNS server in the environment. This is ideally a separate host running DNS software and can provide name resolution to hosts and containers running on the platform.

- **Aggregated resource requirements for Red Hat OpenShift Data Foundation[10]**

Red Hat OpenShift Data Foundation services consist of an initial set of base services and can be extended with additional device sets. All these Red Hat OpenShift Data Foundation services pods are scheduled by Kubernetes on OpenShift Container Platform nodes according to resource requirements. Expanding the cluster in multiples of three, one node in each failure domain, is an easy way to satisfy pod placement rules.

Deployment Mode	Base services (Ceph Monitor + OSD Nodes)	Additional device Set (Ceph OSD Nodes)
Internal	30 CPU (logical) 72 GiB memory 3 storage devices	6 CPU (logical) 15 GiB memory 3 storage devices
External	4 CPU (logical) 16 GiB memory	Not applicable

Example: For a 3-node cluster in an internal mode deployment with a single device set, a minimum of 3 x 10 = 30 units of CPU are required.

Installation Components

Below is the list of installation components to deploy a fully functional Red Hat OpenShift cluster with OpenShift Data Foundation.

1. Cluster details setup
2. OpenShift Operator Selection
3. Cluster hosts discovery
4. OS installation disk selection
5. Cluster network setup
6. Summary and review
7. Storage operator setup from console

Deployment

Step1: Login to the Hybrid Cloud Console provided by the Red Hat. This console will be used as a Cluster deployment tool for Red Hat OpenShift. Fill in the cluster name and preferred domain name to be used to access the OpenShift Cluster. Ensure that the DNS in the network has a separate domain for the deployment.

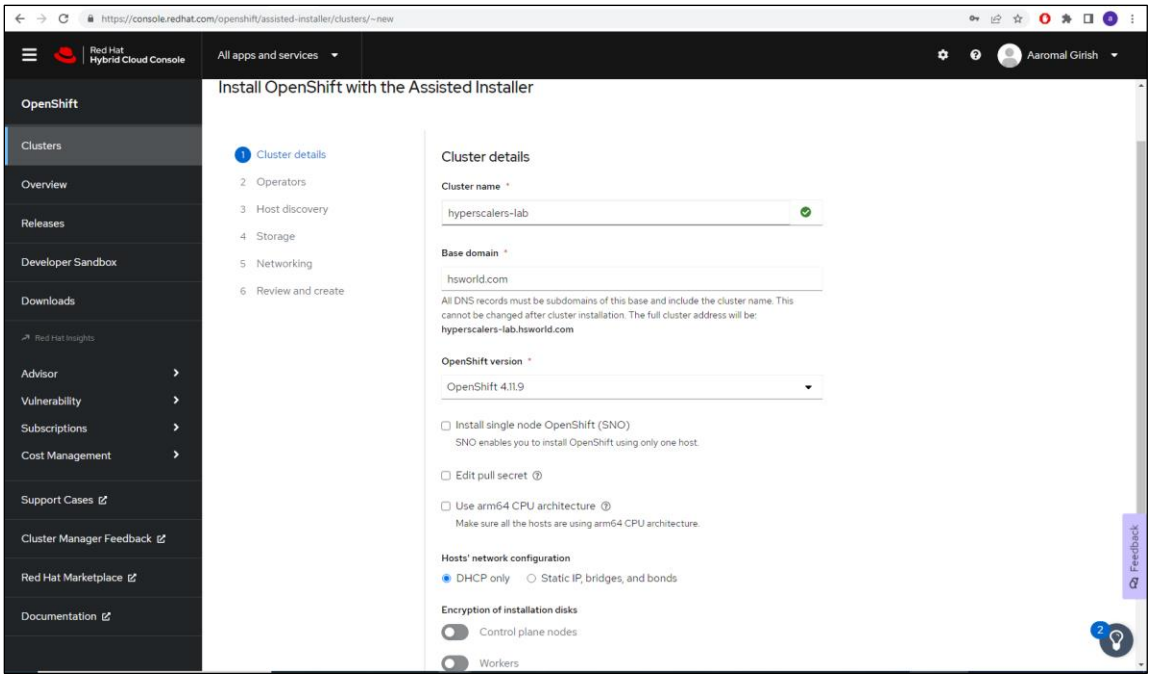

Figure 18 Cluster Details page for OpenShift Assisted Installer

Step2: Select the OpenShift Virtualisation operator from the list. Please note that Red Hat OpenShift Data Foundation can be installed after the cluster deployment as it saves time for the cluster deployment.

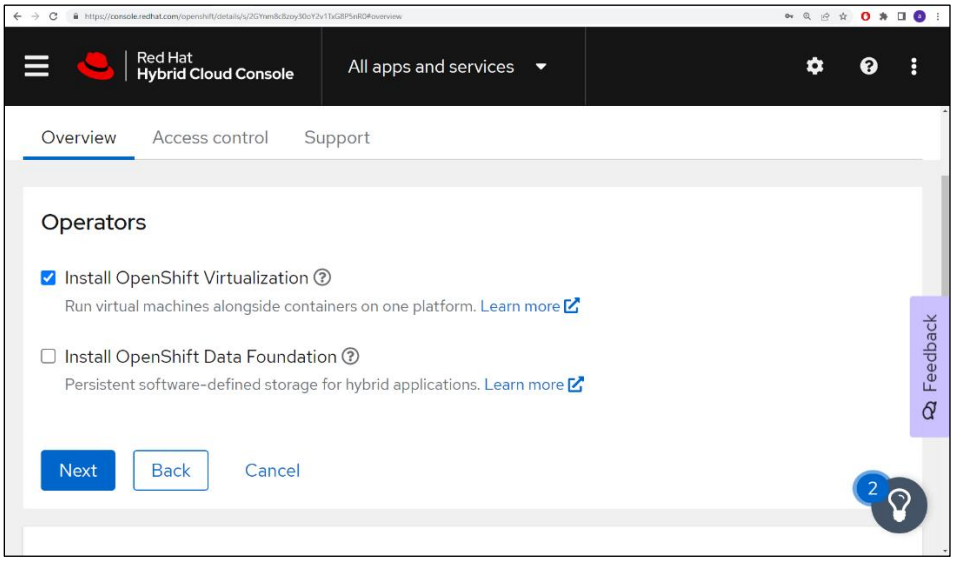

Figure 19 Operator Selection

Step3: On the next section, host discovery is started in the network to identify all the resources/servers that are booted with Red Hat Core OS discovery ISO. Click on the "Add hosts" button to download the ISO. Use Rufus to create a bootable flash USB and configure the BIOS of each node to boot from this USB for one-time boot. Ensure that the second boot option is the SSD that is used to install Red Hat Core OS.

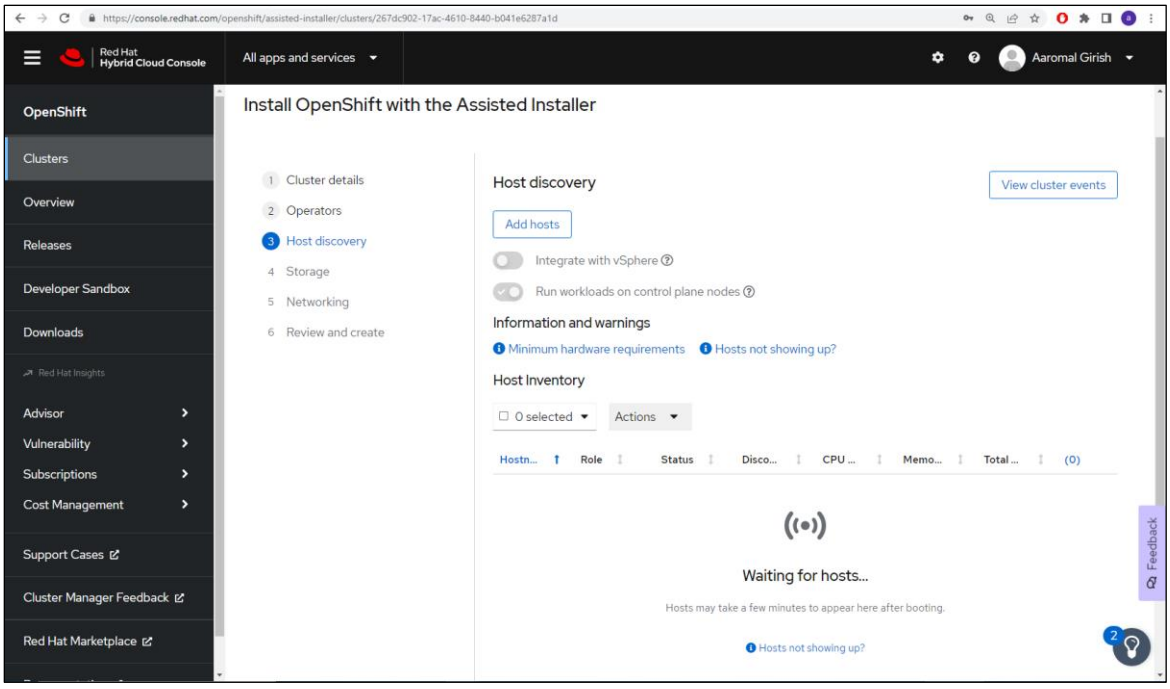

Figure 20 Host discovery in Red Hat OpenShift

Step4: In the next step, the storage for installing the OpenShift is selected. Note that the disk will be erased and formatted during the OpenShift installation. Ensure that the status of the nodes are in the ready state before proceeding to the next steps.

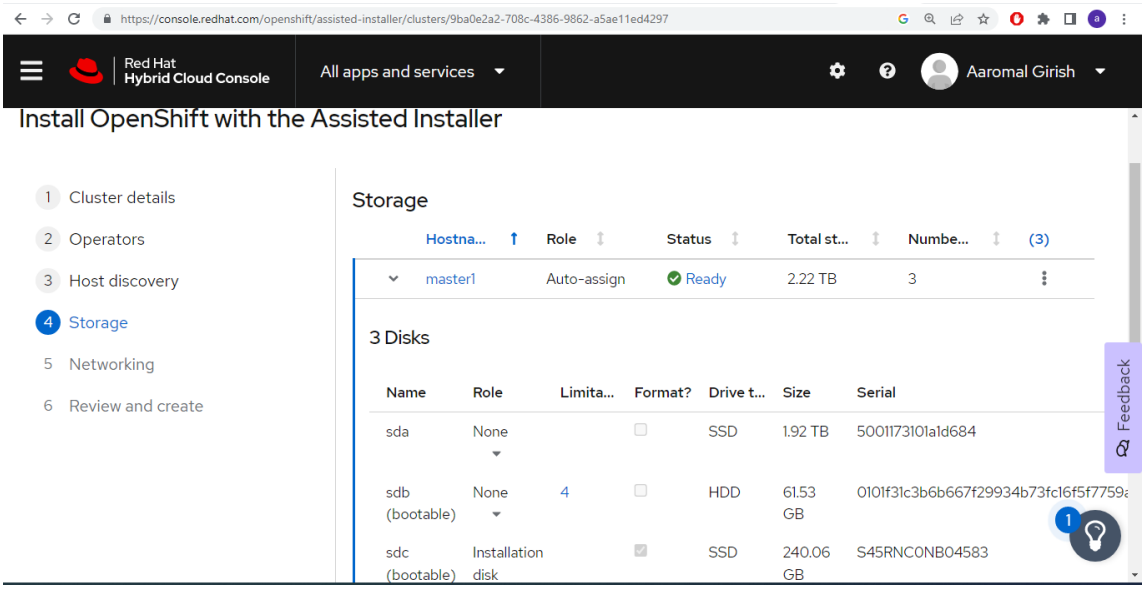

Figure 21 Installation disk selection

Step5: The next step in the deployment of Red Hat OpenShift is to configure the network to be used for the VMs, API and APPs. Choose the subnet and the IP range to use for the OpenShift processes. It is recommended to use a separate subnet with a wide range of available IPs to ensure the long-term usability.

Individual VLANs can achieve network isolation between the management networks, data networks and the storage networks. Typically, two separate switches can be implemented with one dedicated management switch and the other

switch can have two VLANs and individual subnet to cater for the data network and the storage network. This way can allow the network traffic without any bottleneck and the appliance access becomes more secure.

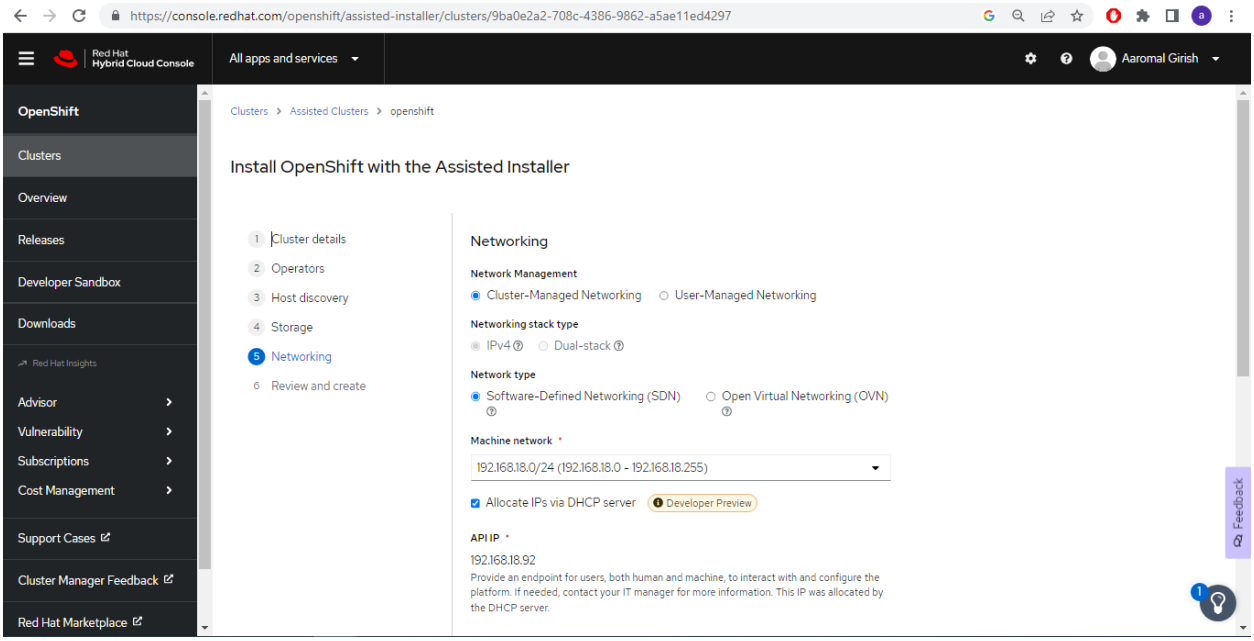

Figure 22 Cluster network selection and configuration

Step6: On the next screen, you will be asked to verify the summary of the inputs and confirm to proceed with the auto deployment of Red Hat OpenShift. The installation process takes about two hours or more based on the resource provided. During this process, the nodes reboot itself to boot from the SSD where the Red Hat Core OS (RHCOS) is installed.

Since the automated process could be hanging on for the nodes to reboot automatically, it is strongly recommended to the operators to closely watch the progress and manually reboot the nodes in case the automatic reboot fails. After this step, the nodes must boot into the physical disk rather than the live USB boot.

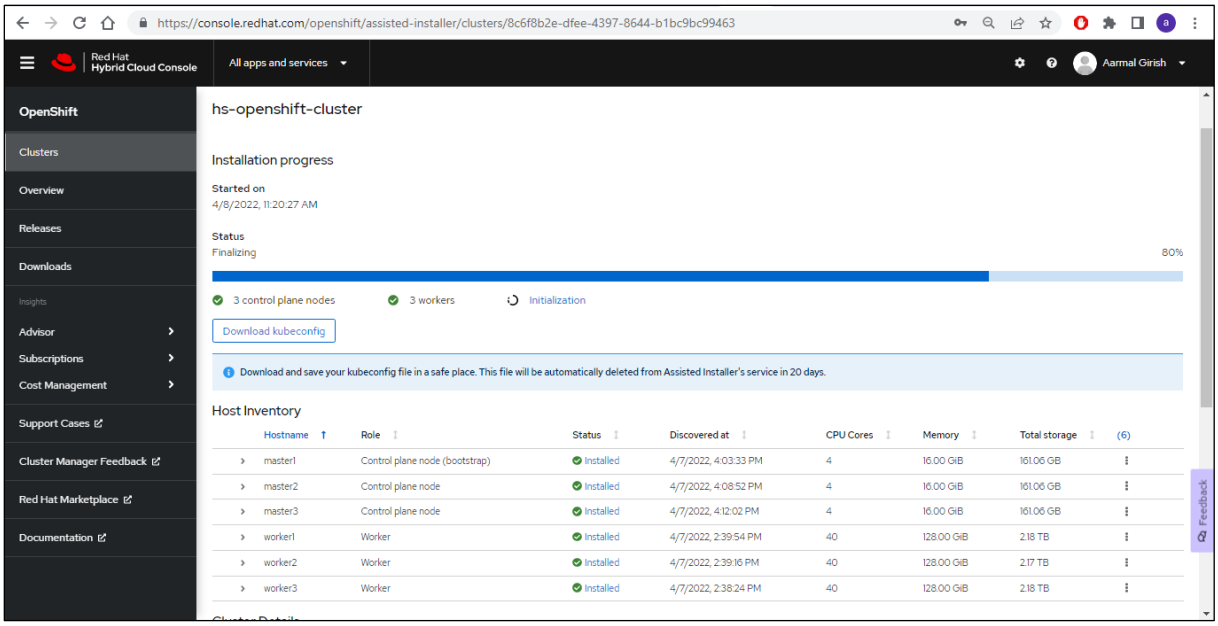

Figure 23 Hosts installed with Red Hat OpenShift

Step7: The console to the Red Hat OpenShift can be accessed from the web console URL if the DNS is configured with the redirection IP addresses for the API and *.apps. For a client machine to access the Red Hat OpenShift console without the DNS server, the content of /etc/hosts must be added with the IP redirects that will be based on your domain naming conventions in the "Not able to access the Web Console?" section. The username and the password for the console can be seen from the hybrid cloud dashboard itself.

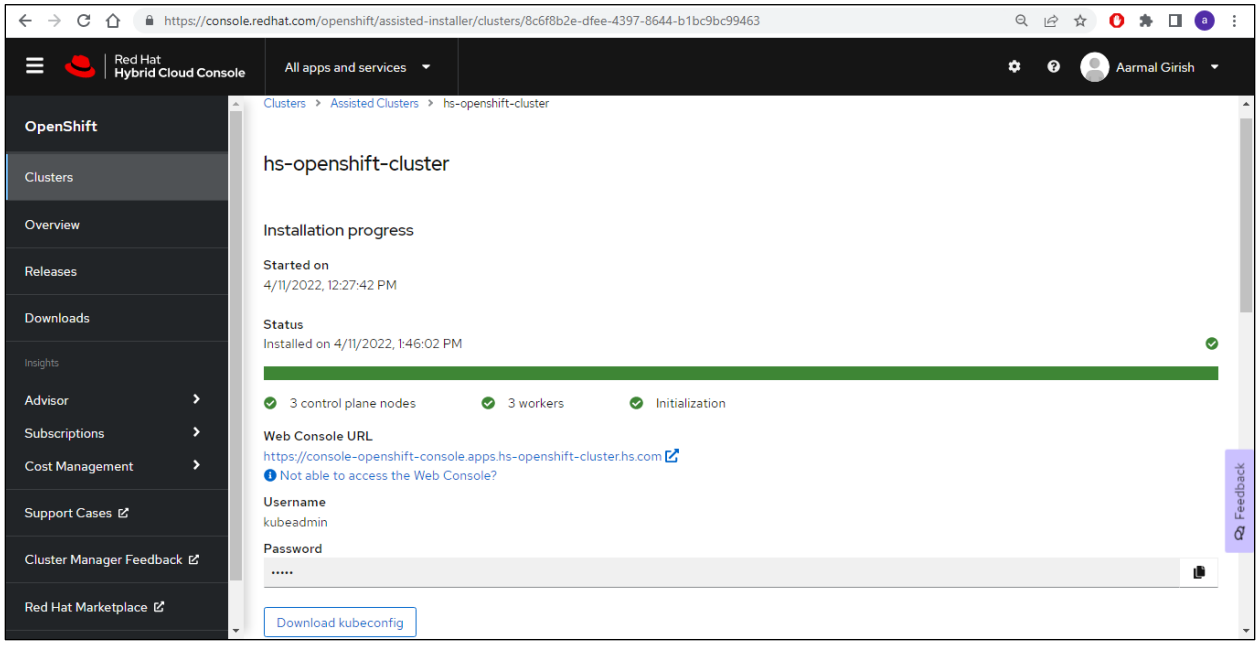

Figure 24 Red Hat OpenShift Dashboard

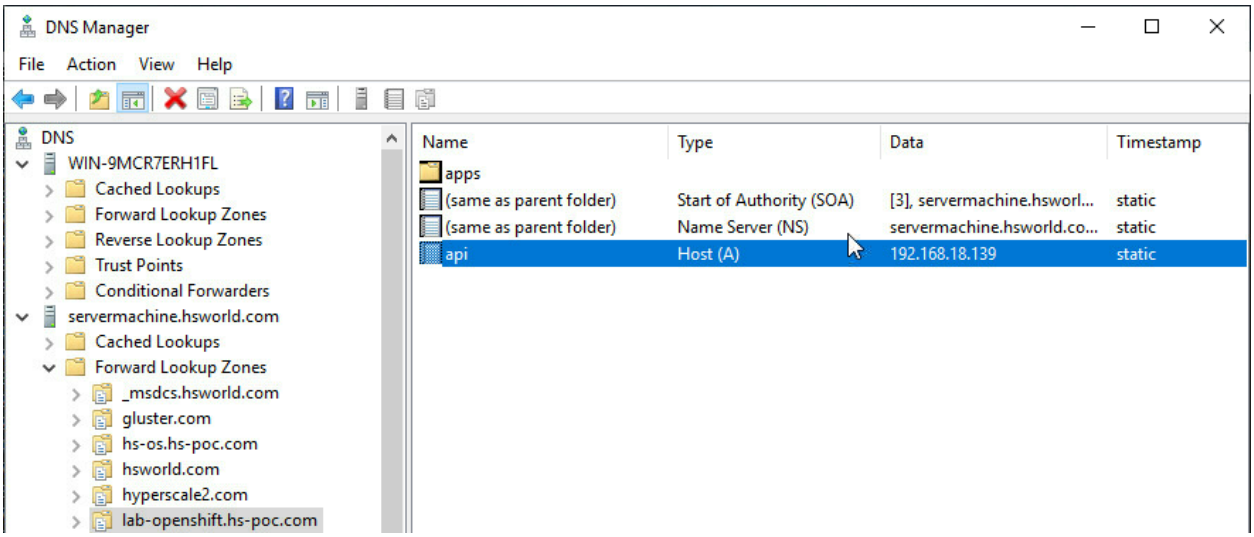

Figure 25 DNS manager with domain and "A" hosts

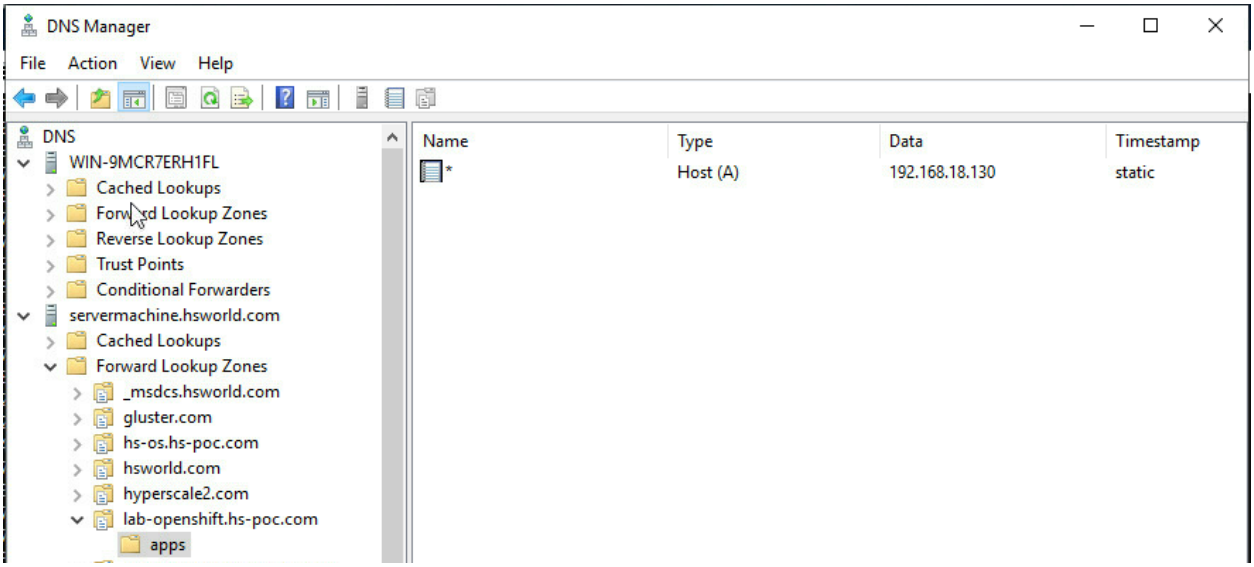

Figure 26 A host created in DNS manager for wildcard

Step8: The next step is to view the resources in the Red Hat OpenShift Cluster dashboard. This is the available virtualized compute and networking resource for the containers and applications to make use of.

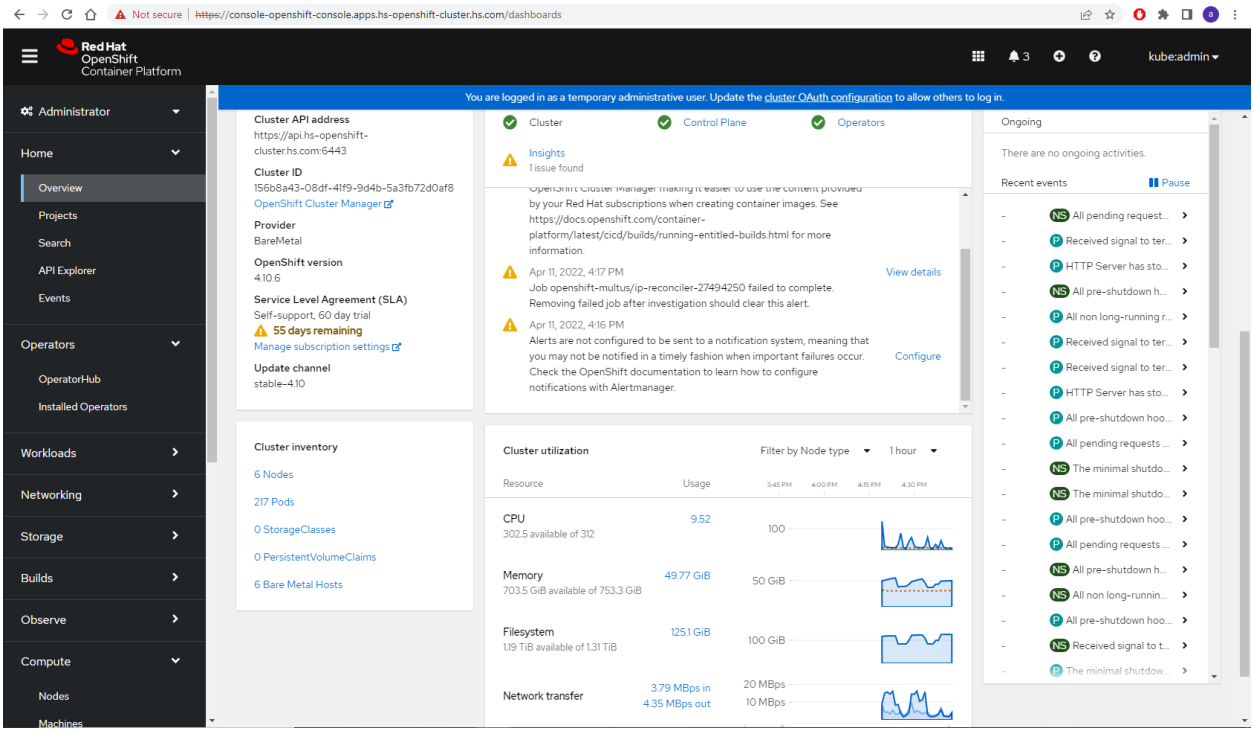

Figure 27 Red Hat OpenShift Dashboard

Step9: The host machines involved in the Red Hat OpenShift cluster can be viewed from the Dashboard -> Compute -> Nodes.

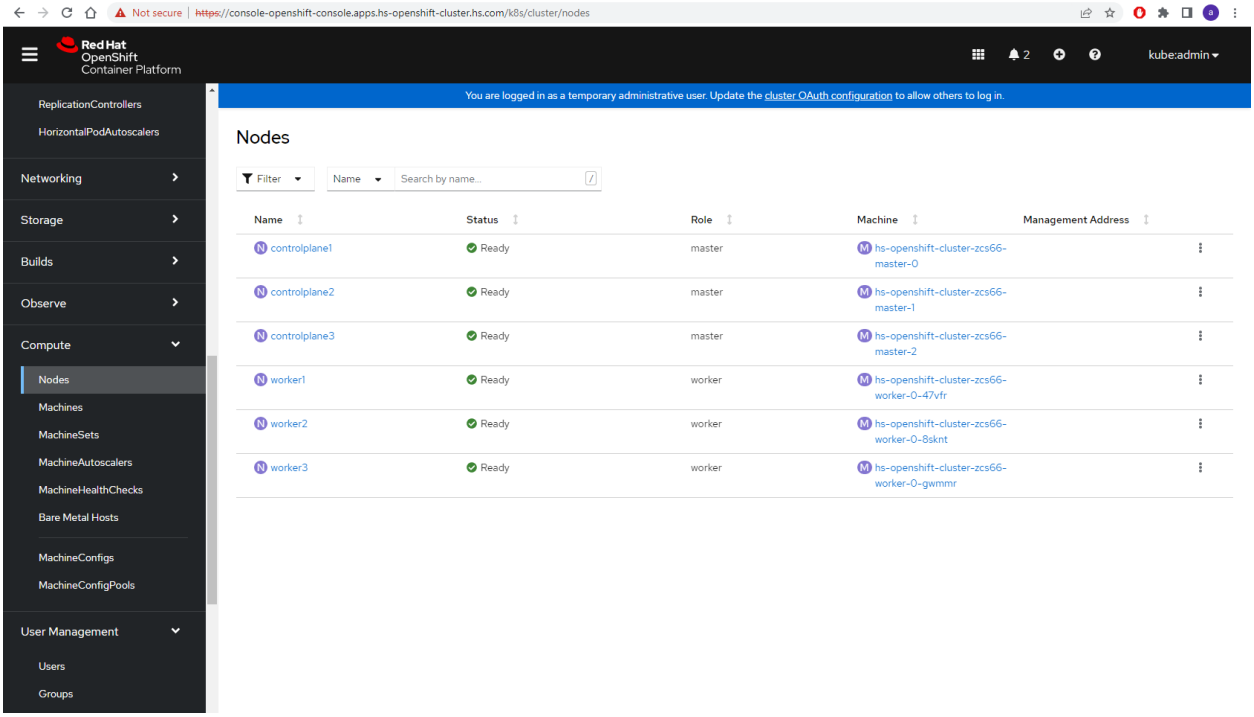

Figure 28 Red Hat OpenShift compute nodes

Step10: The Red Hat OpenShift containerized environment requires a persistent storage to assign the containers and pods to make use of. In this appliance, Red Hat's Data Foundation operator is installed to create a persistent, hyperconverged storage disks form each of the cluster hosts.

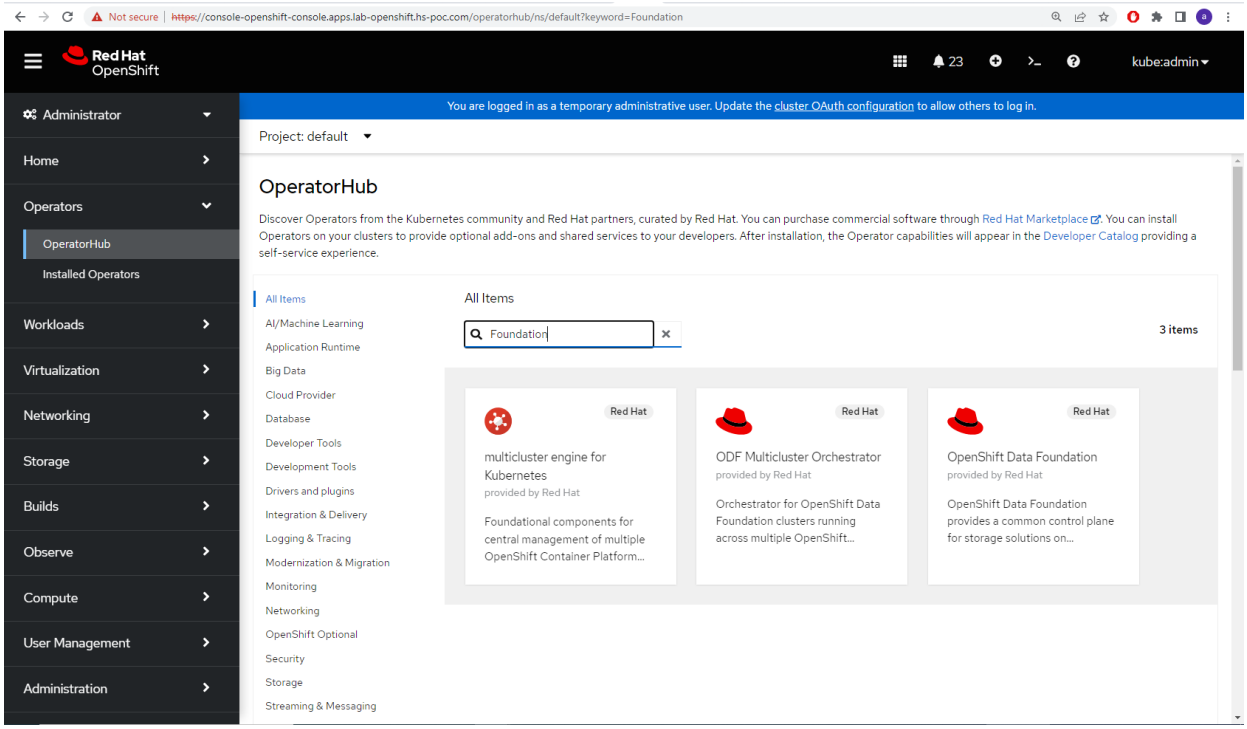

Figure 29 Red Hat OpenShift Data Foundation

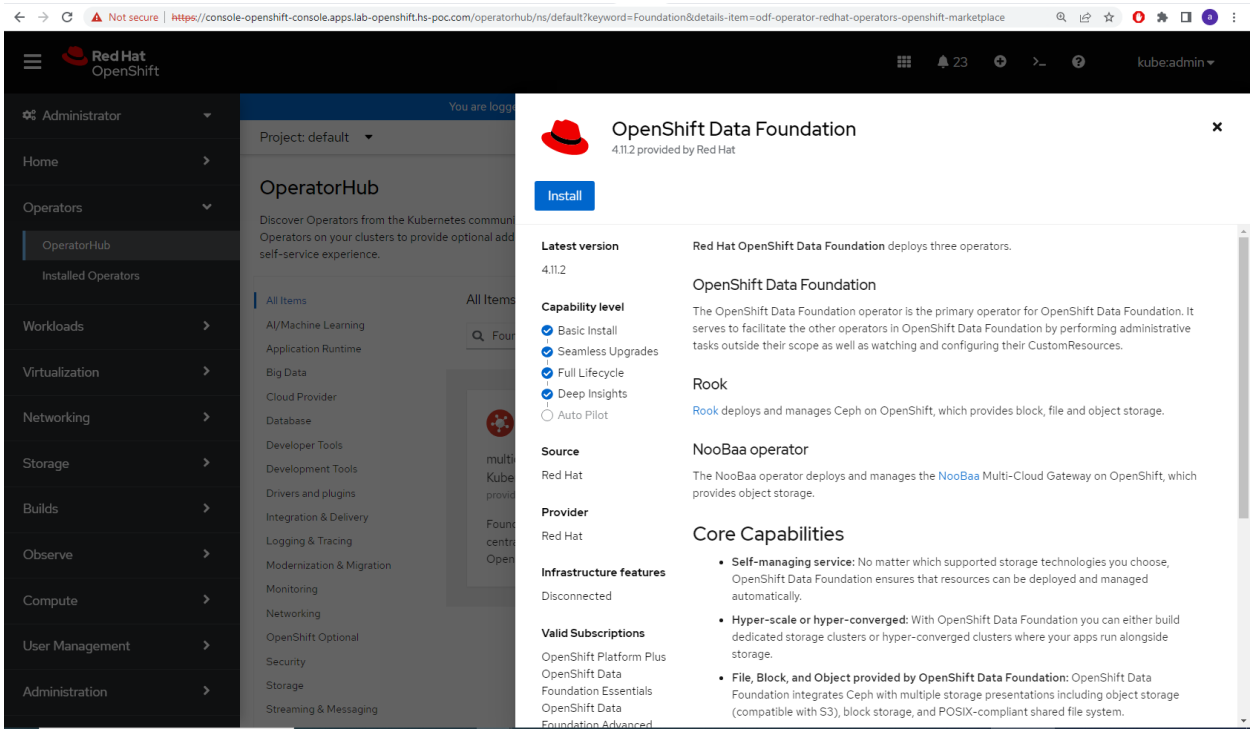

Figure 30 Installation of Red Hat OpenShift Data Foundation 4.11.2

Step11: Once the installation of Red Hat OpenShift Data Foundation is complete, we will create the storage system by choosing the nodes that are involved in the cluster along with the storage drives. Please note that Red Hat Data foundation uses Ceph as the underlying storage technology.

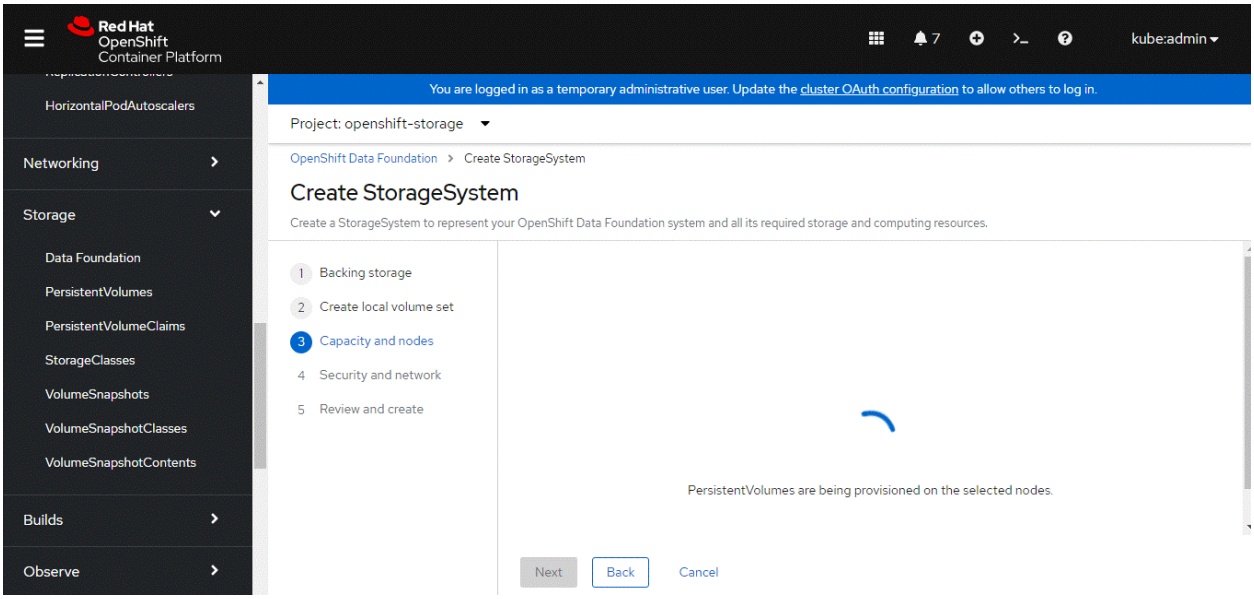

Figure 31 Choose the drives involved in the storage cluster

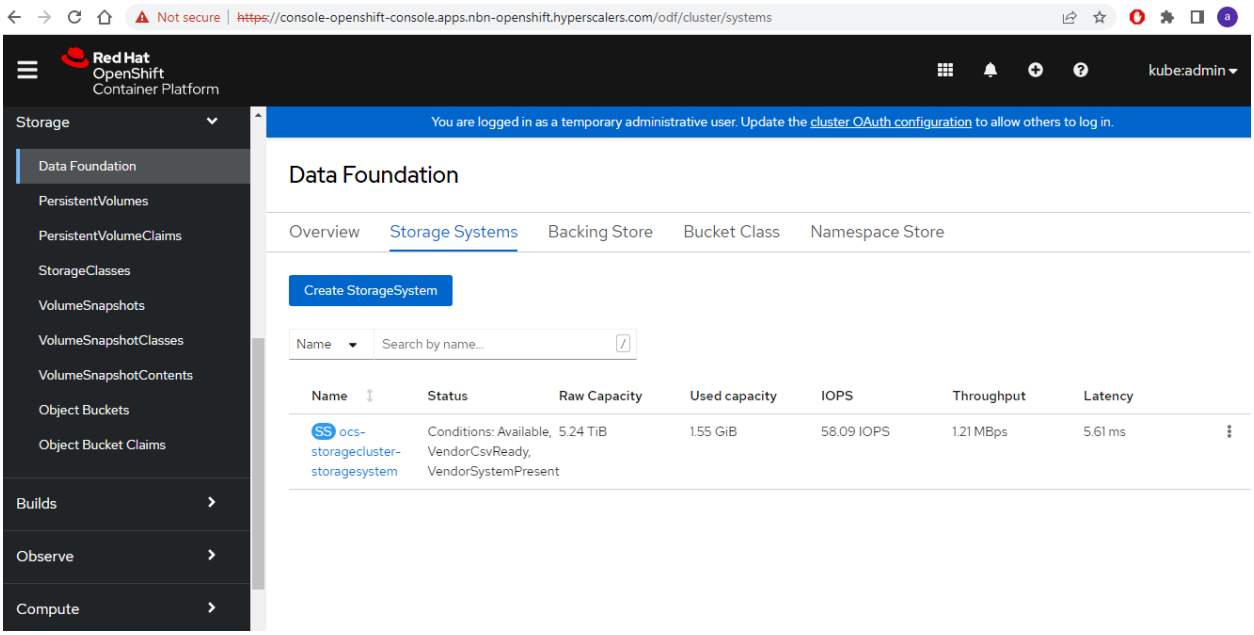

Figure 32 Red Hat OpenShift Storage System

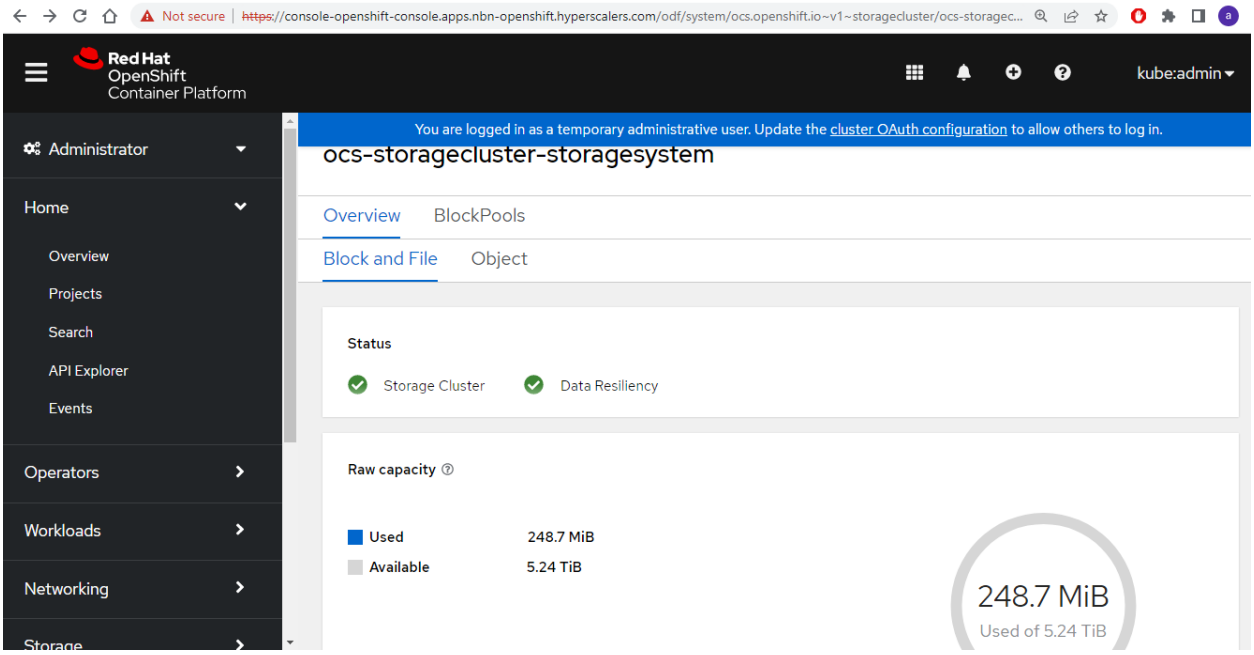

Figure 33 Block, File and Object health status

Step12: Finally, for future maintenance and upgrade of the storage system in the Red Hat OpenShift can be done from Installed Operators -> OpenShift Data Foundation -> Subscription

Note that the automatic updates and upgrades are to be disabled in a production system to enable a controlled development and operation of the compute cluster and the storage cluster.

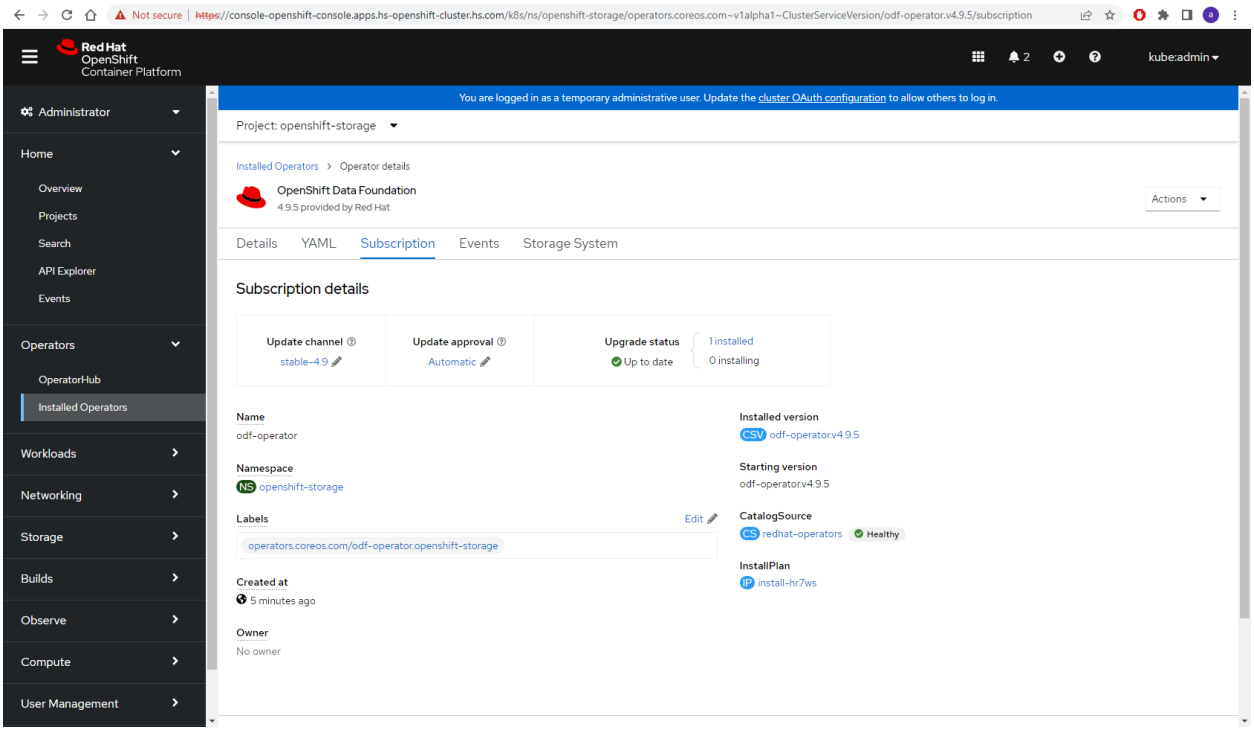

Figure 34 Red Hat Data Foundation version details

4 CONFIGURE THE APPLIANCE

Once the base project deployment is complete, we need to configure the appliance to deploy the workloads in the form of containers.

Install Operators:

Red Hat OpenShift platform provides several operators that are pre-engineered to perform specific applications like Elastic Search, Kafka, SSL certificate management etc. Please note that a subscription to individual services may be required to use the applications in Red Hat OpenShift environment.

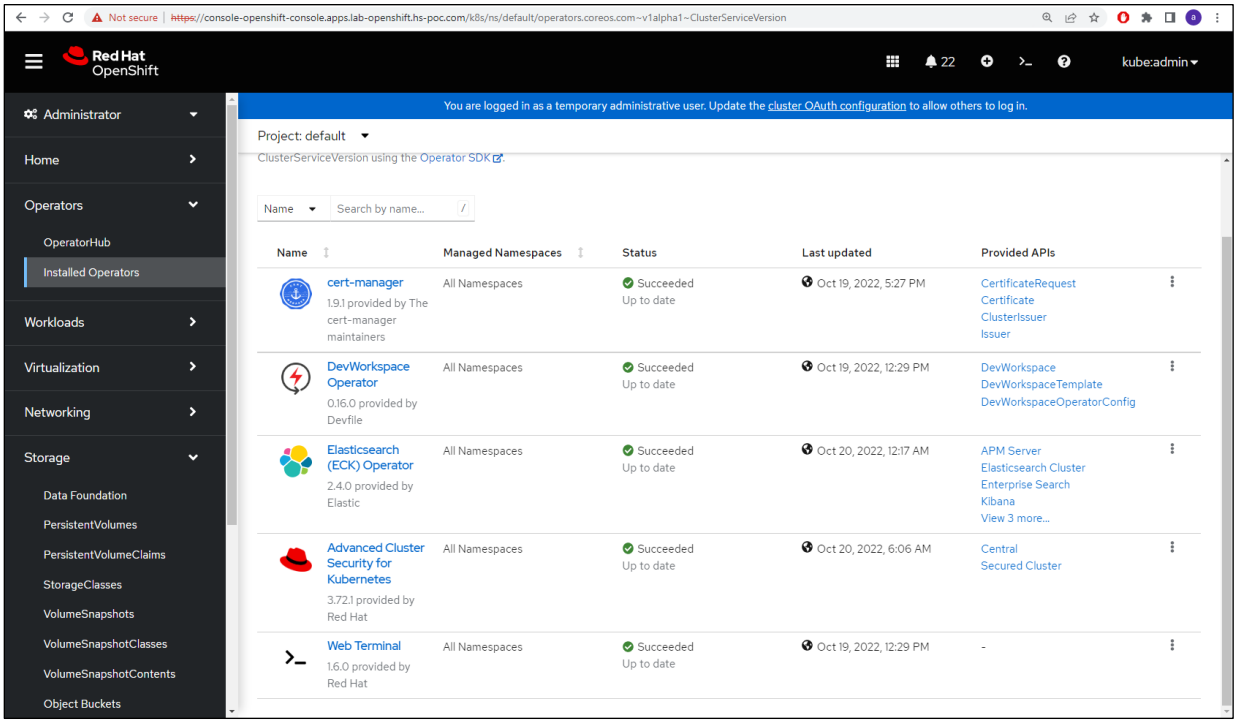

Figure 35 Red Hat OpenShift Operators

Persistent volumes, Claims and Storage classes:

In a containerized environment, the storage is classified as Storage class which is consumed by the persistent volumes using the persistent volume claims. Each of the persistent volume claim is associated with an application which decides the size of the storage allocation to them. As mentioned earlier, Ceph is the underlying storage technology for the Red Hat OpenShift that can provide a block, filesystem, and object storage classes for the application to use from. Below are some of the screenshots that show the storage platform provided by the Red Hat Data Foundation.

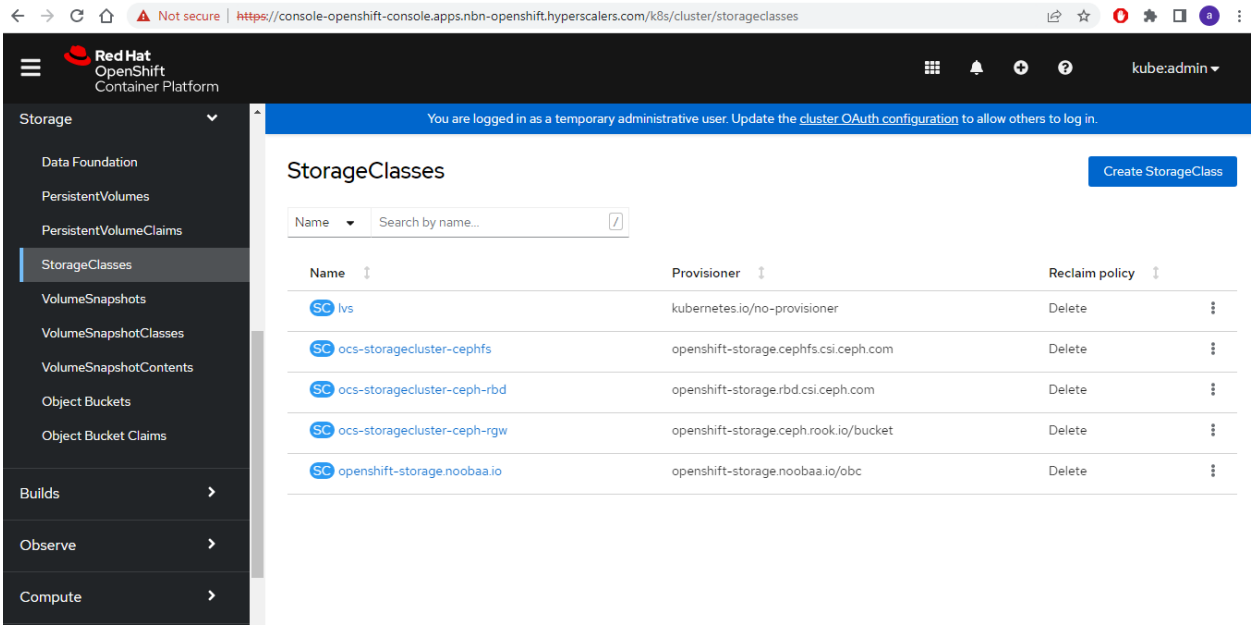

Figure 36 OpenShift Block, File and Object storage classes

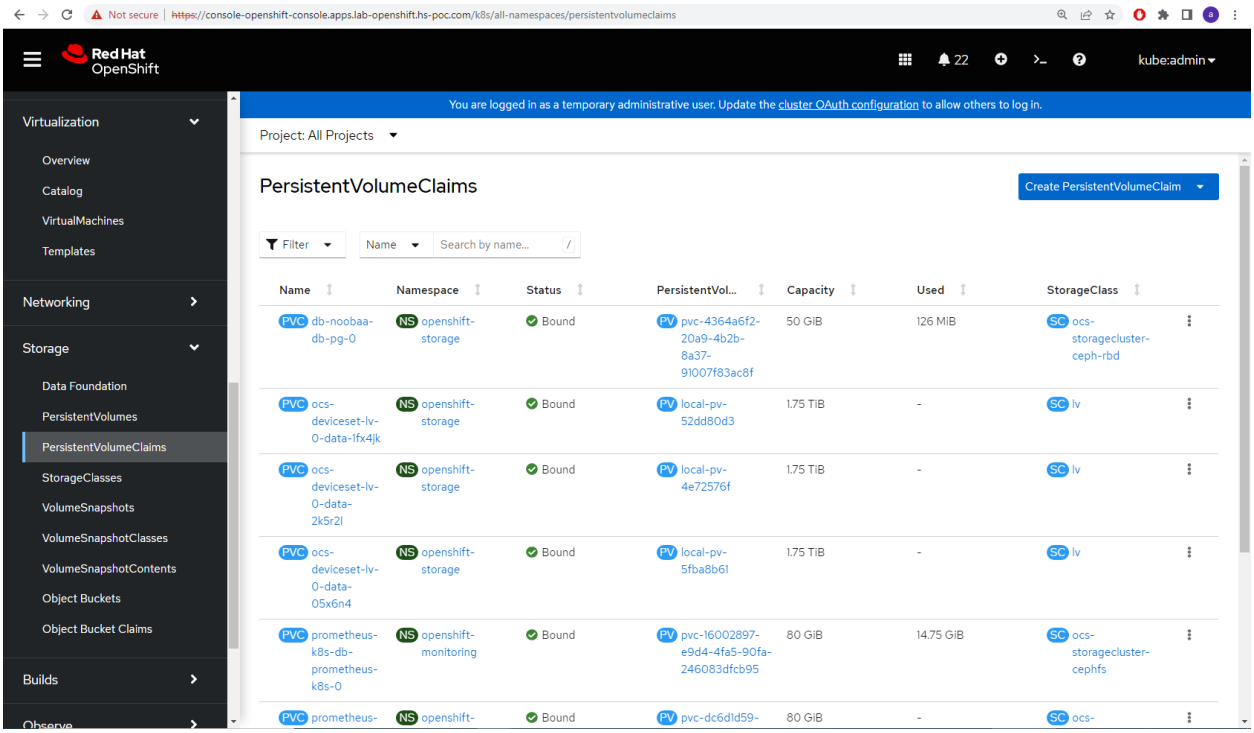

Figure 37 Red Hat Data Foundation Persistent Volume Claims

Pods:

Pods are the basic unit of the containerized resource where an application resides. Pods can be replicated for high availability and the route to the application is defined in the pods. Every pod is associated with a namespace or project and they can be accessed via SSH.

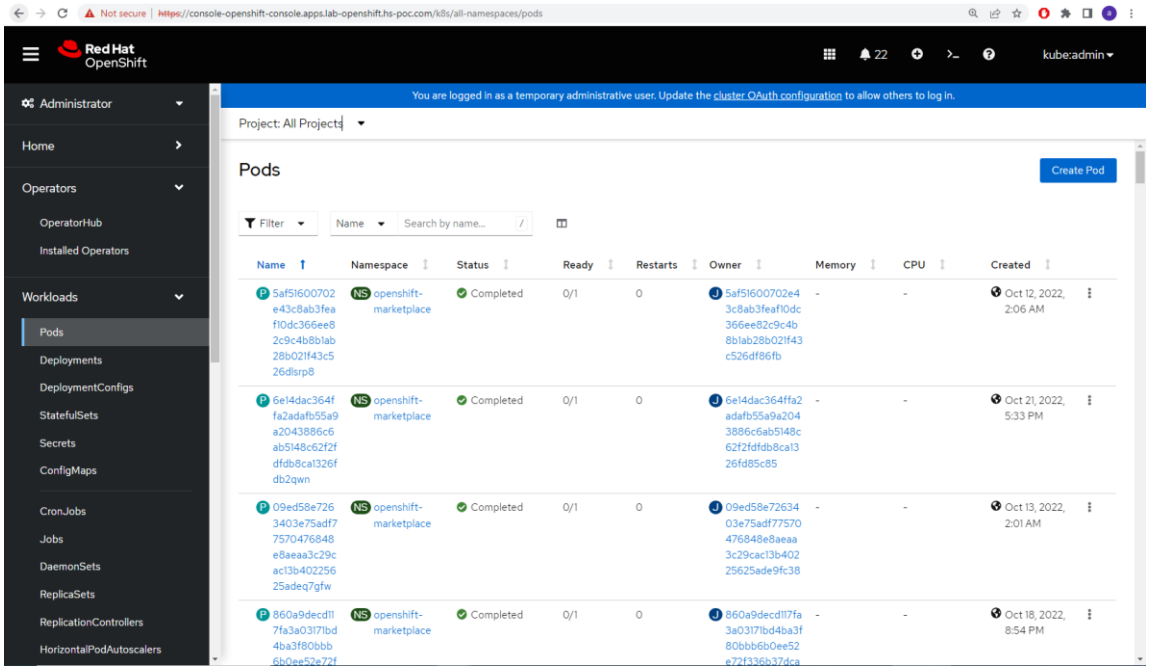

Figure 38 Pods in Red Hat OpenShift

Virtualization:

Red Hat Virtualisation is enabled on the OpenShift using the Virtualization Operator. This enables the user to have the containers alongside the virtual machines to provide adaptability for application that are built for Kubernetes containers and those apps built for virtual machines. Hence, the same physical resource can used to deploy containerised as well as virtual machine based workloads.

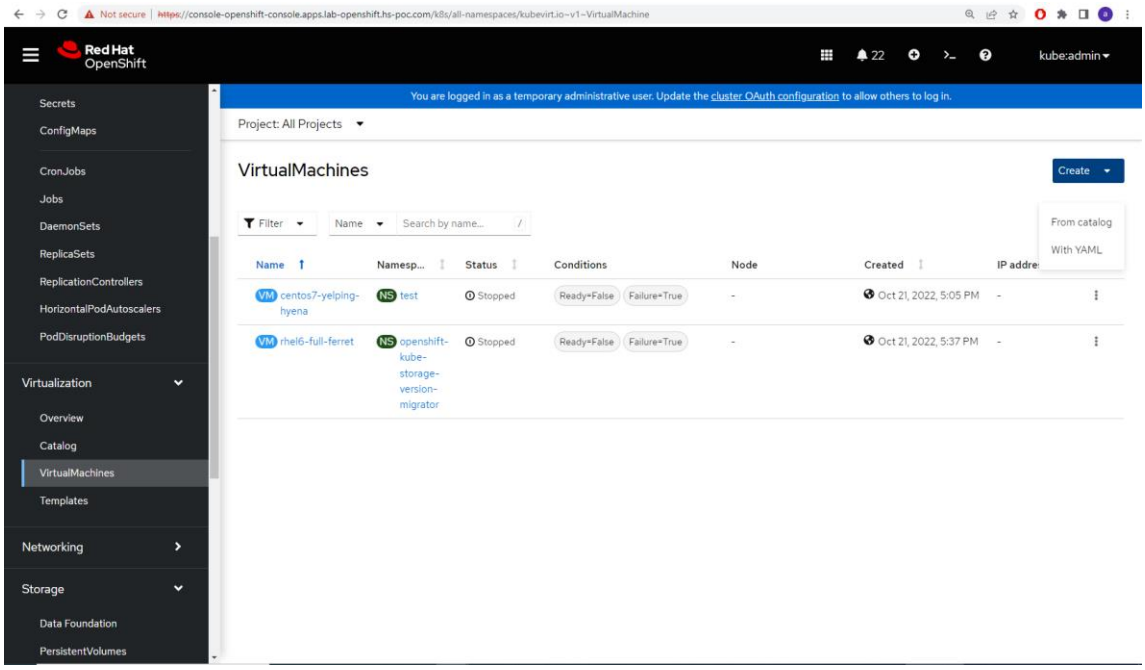

Figure 39 Virtualization in Red Hat OpenShift

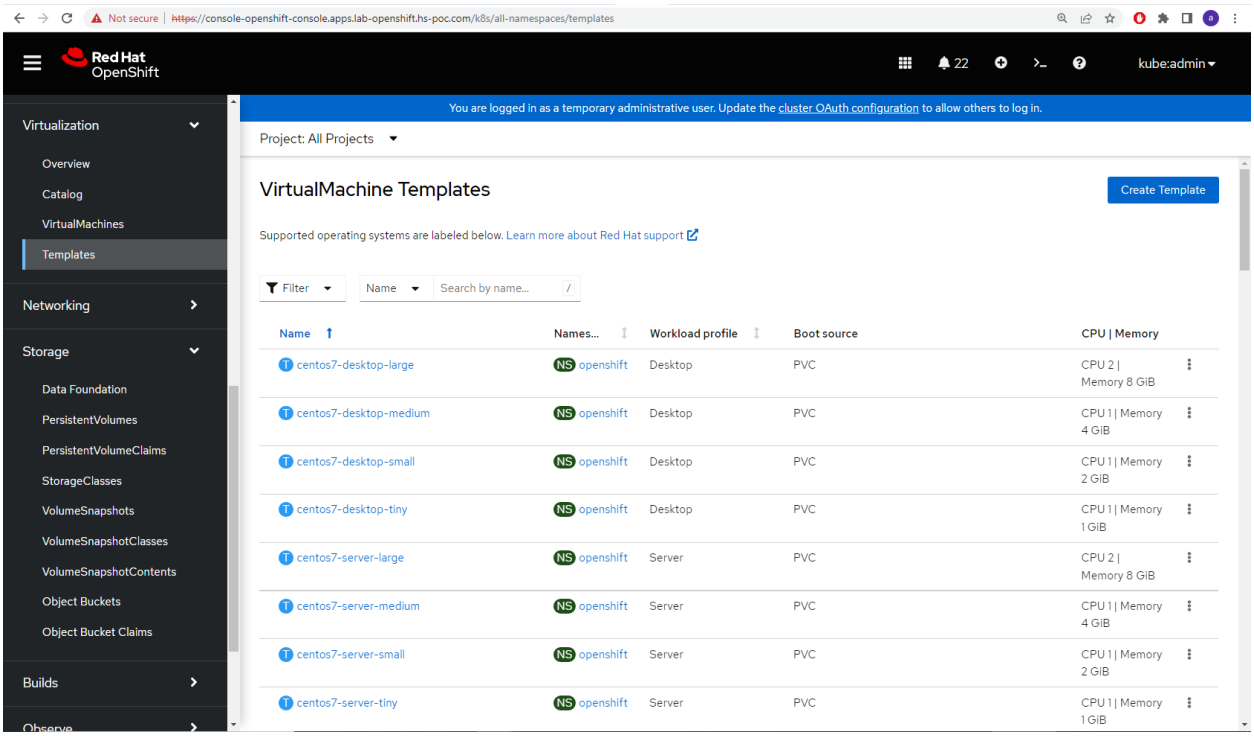

Figure 40 Virtual Machine templates in Red Hat OpenShift

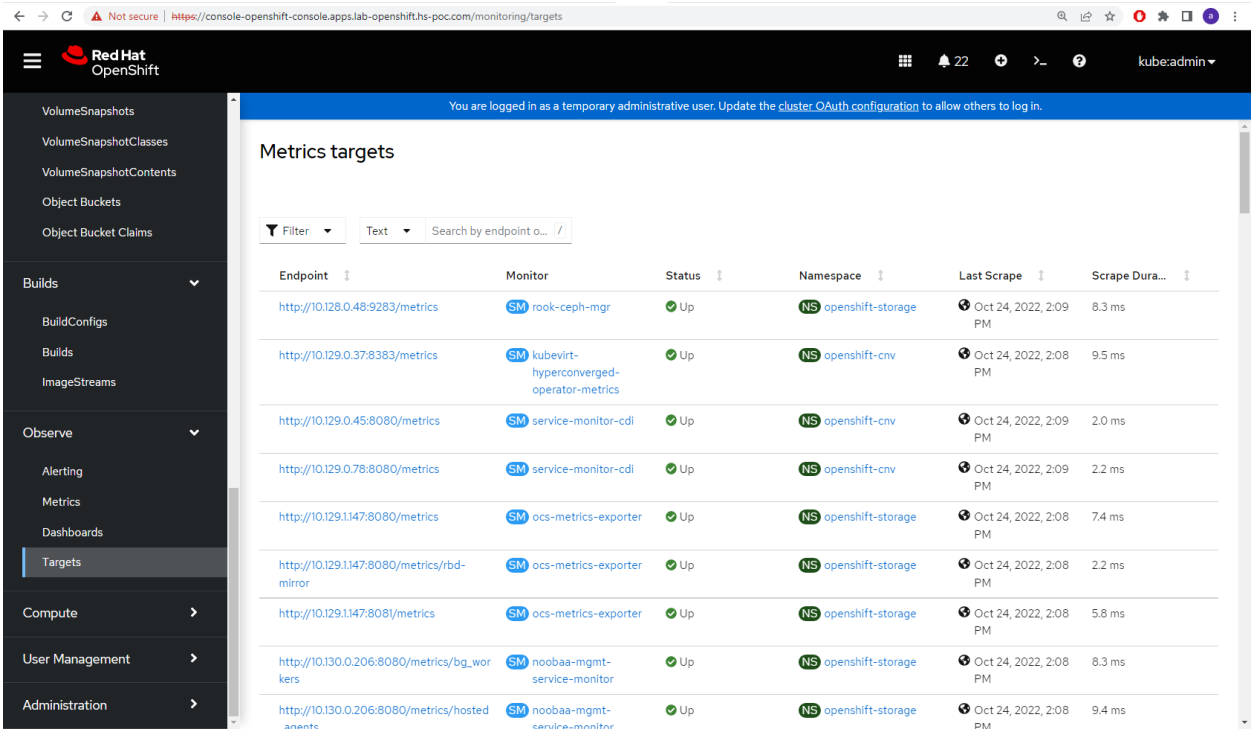

Figure 41 Access Endpoints to Monitoring and Health Metrics

Cluster Upgrades and Subscription:

Red Hat OpenShift provides a 60-day free evaluation to setup and support your cluster deployment and then we can attach a license/subscription based on purchase period for the support and services. Cluster upgrade can run on the live environment without affecting the workload but it is highly recommended to understand the requirements of your workloads before upgrading as it could cause inter-operability issues within the apps in the container.

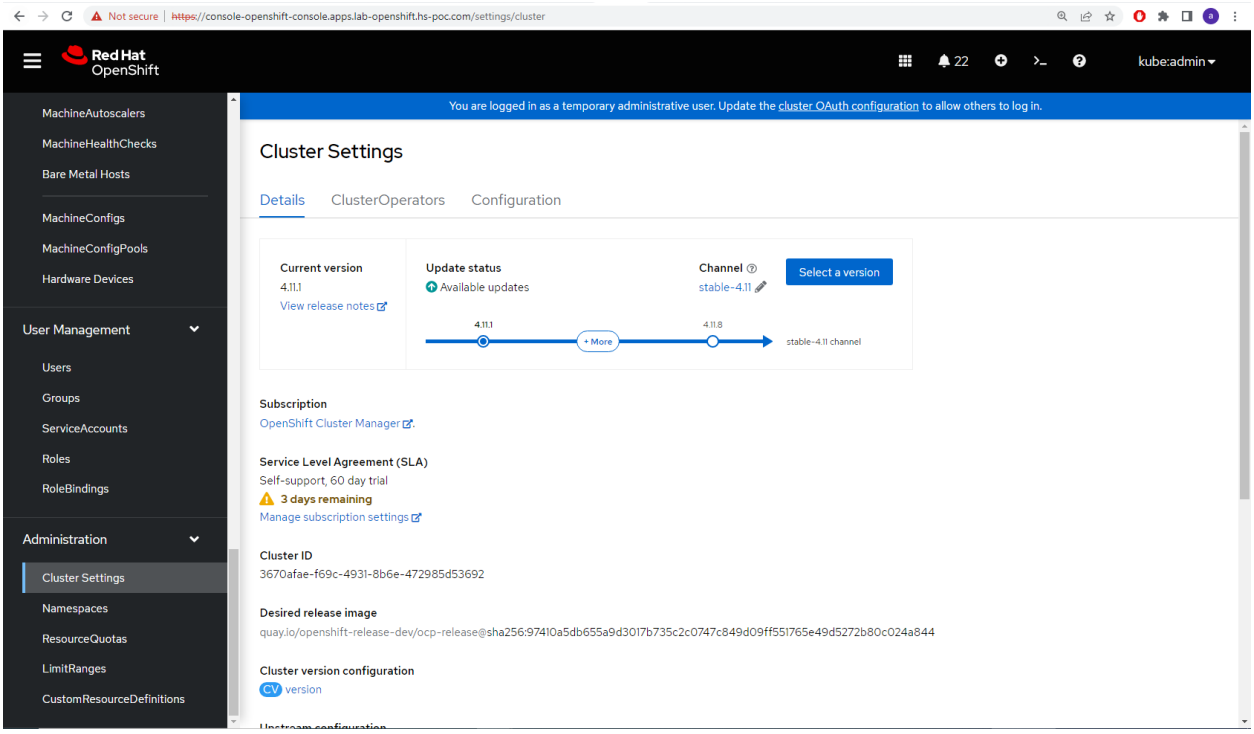

Figure 42 Red Hat OpenShift Cluster maintenance and upgrade

5 TESTING THE APPLIANCE

To test the appliance, a new project is created in the Developer section of the Red Hat OpenShift as shown in the screenshot below. Any container-based application or helm charts in the Red Hat container catalogue, samples or git-hub can be deployed either using the URL or the yaml templates.

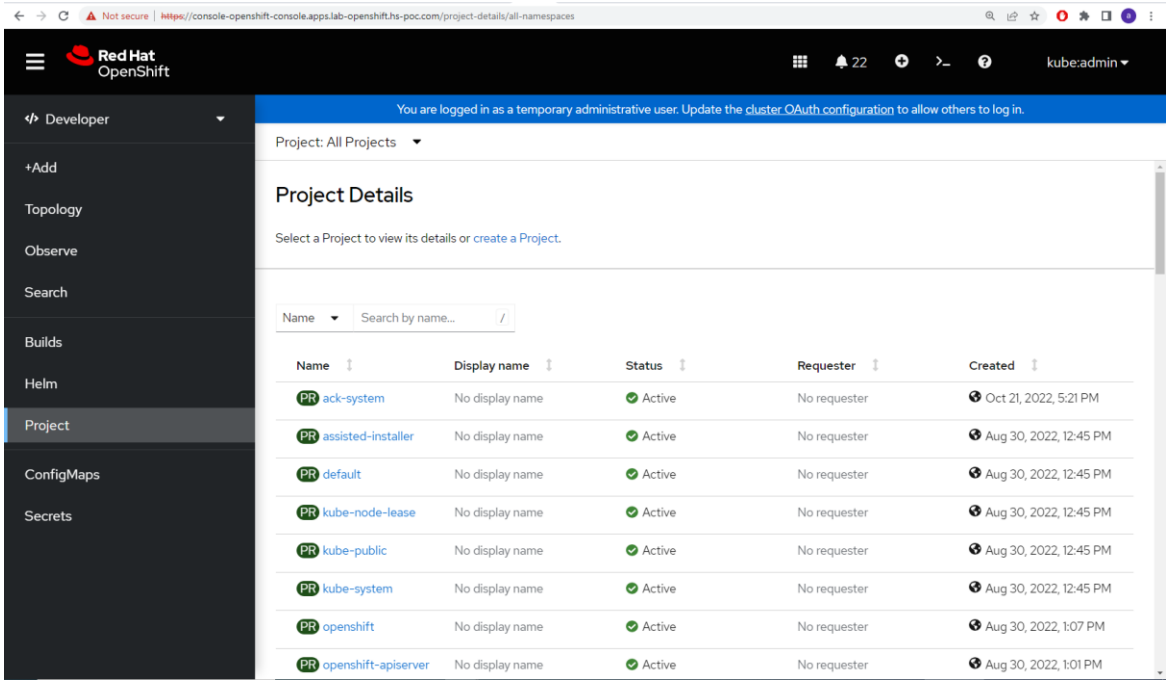

Figure 43 Projects in Red Hat OpenShift

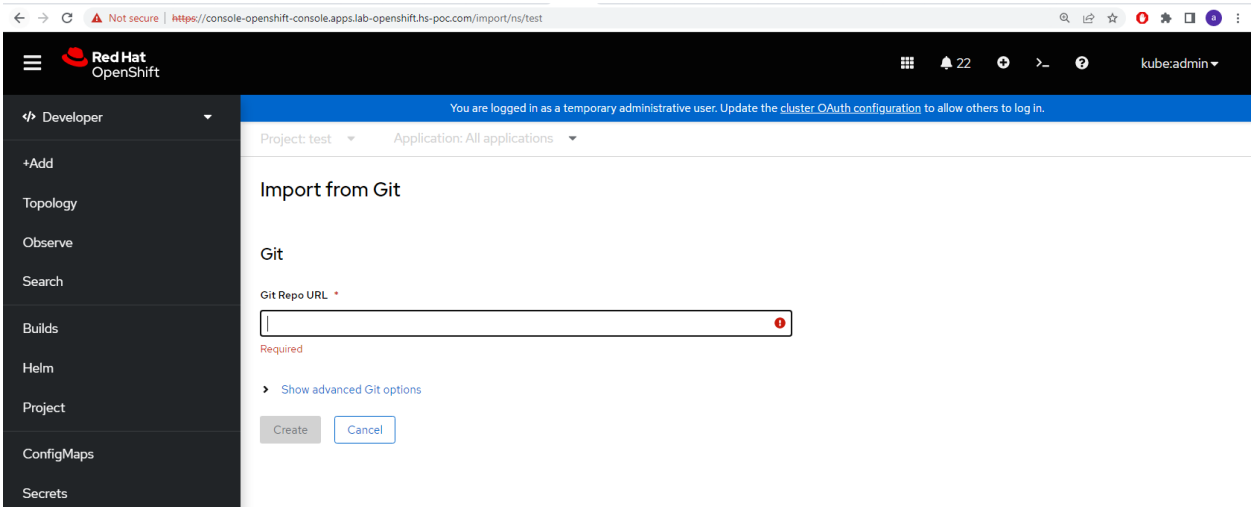

Figure 44 GitHub deployment of apps

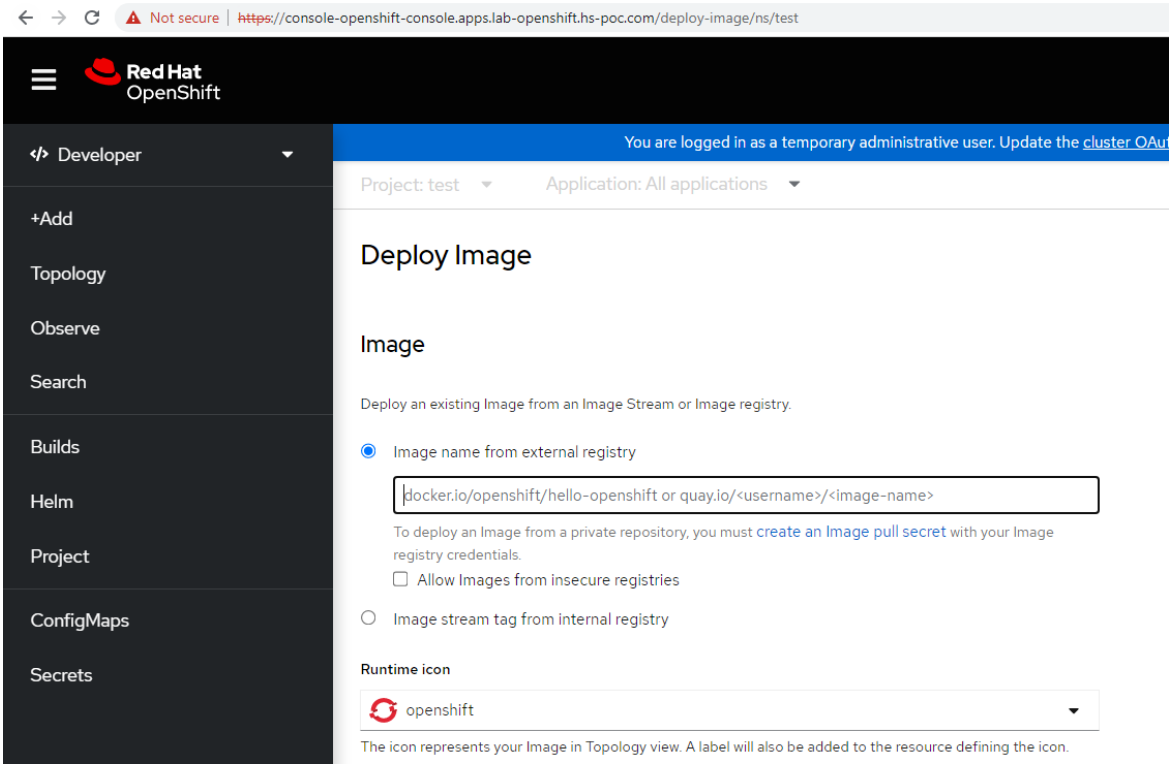

Figure 45 Deploy docker based images

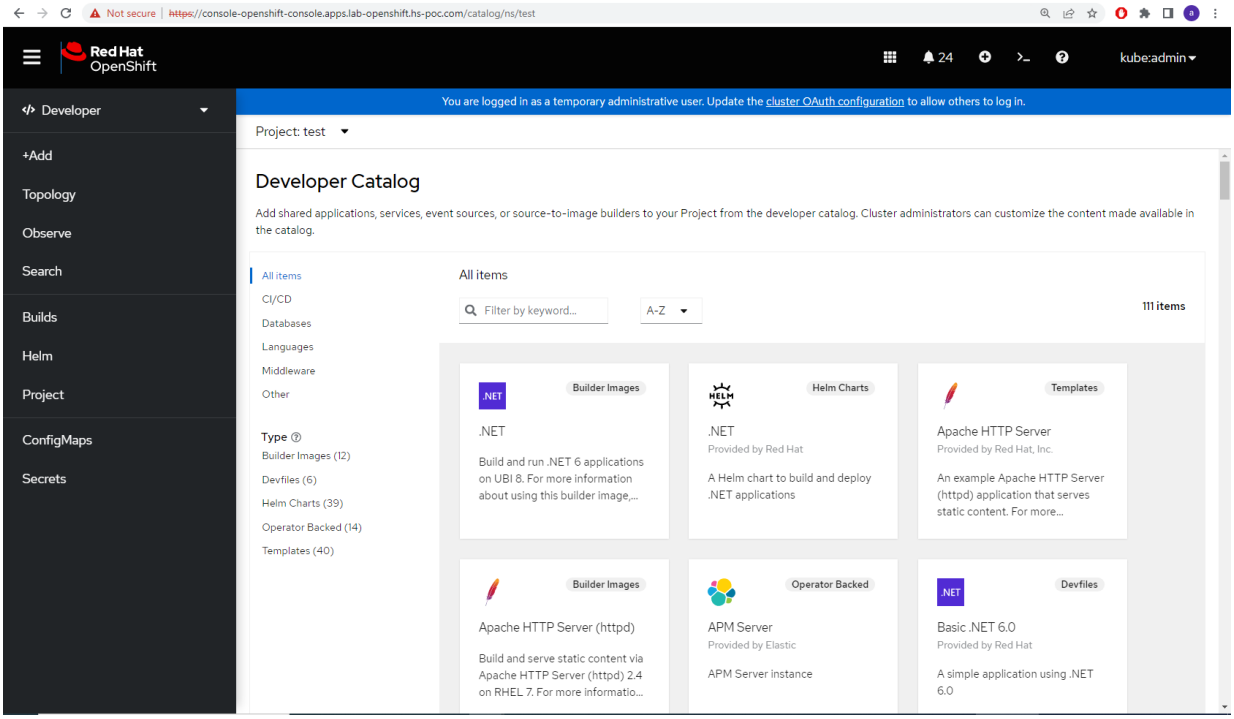

Figure 46 Red Hat OpenShift Catalogue

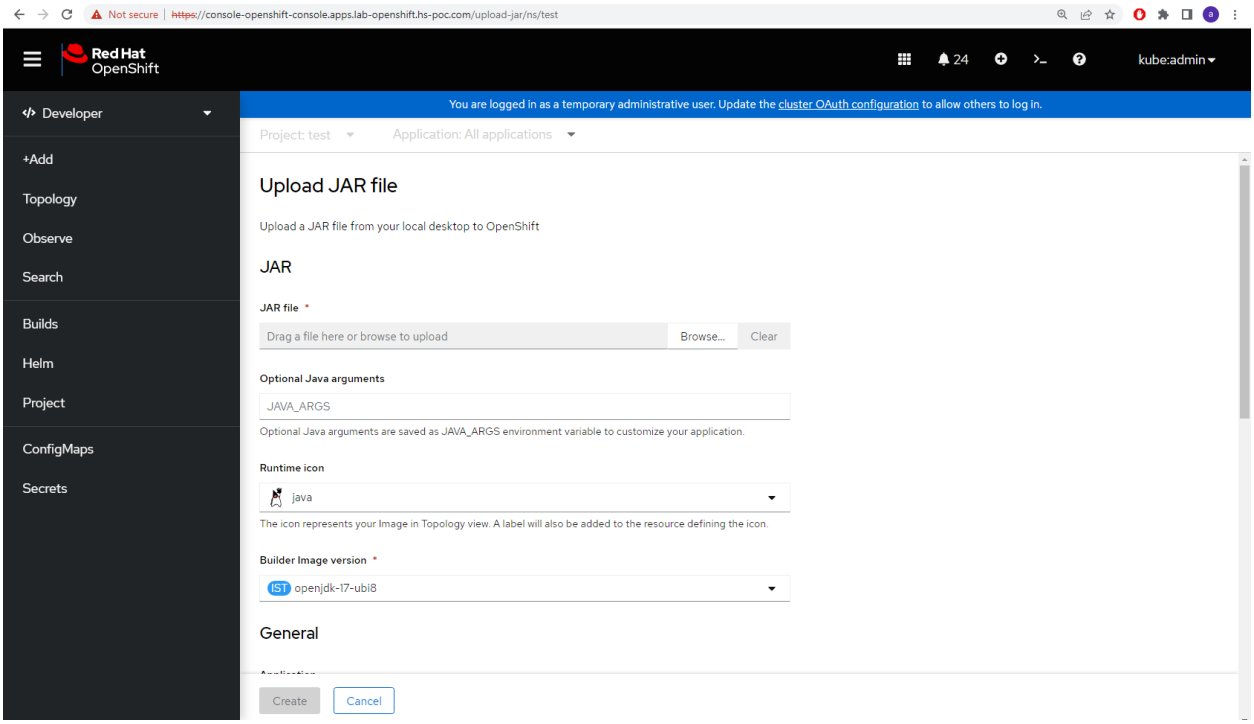

Figure 47 JAVA based image deploy

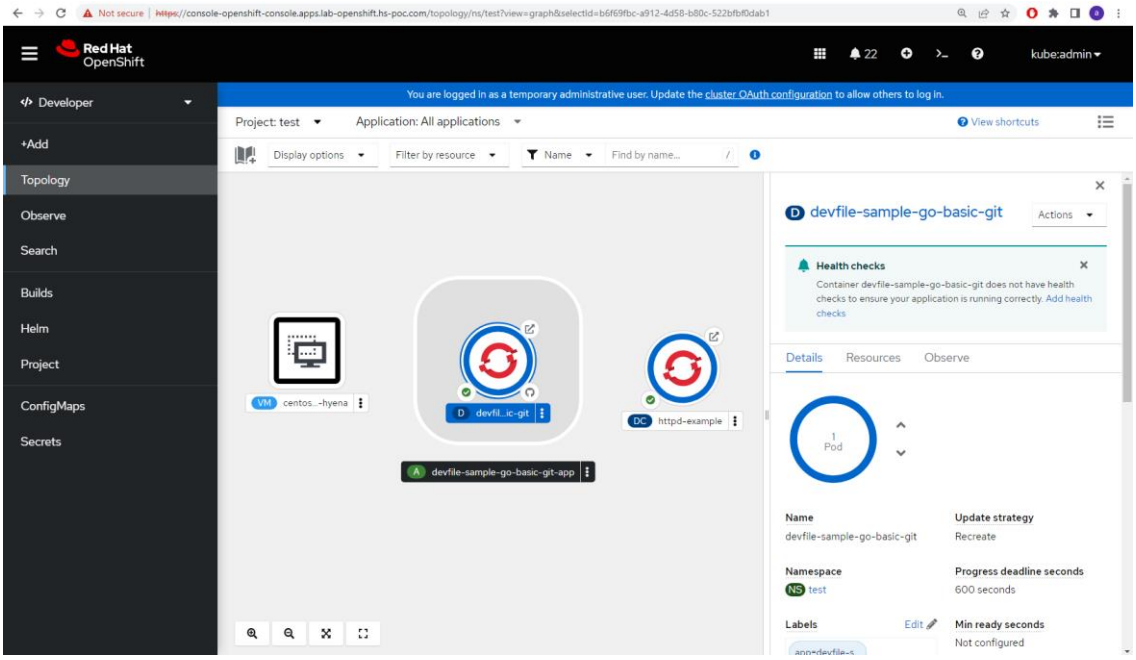

Figure 48 Sample application in Red Hat OpenShift containers

Highly sophisticated confluent Kafka as well as community supported open-source Kafka application can be deployed in the Red Hat OpenShift. They incorporate multiple applications like zookeeper, brokers, Hadoop, etc to inter-communicate and coordinate with each other to enable data streaming, message queuing and analytics.

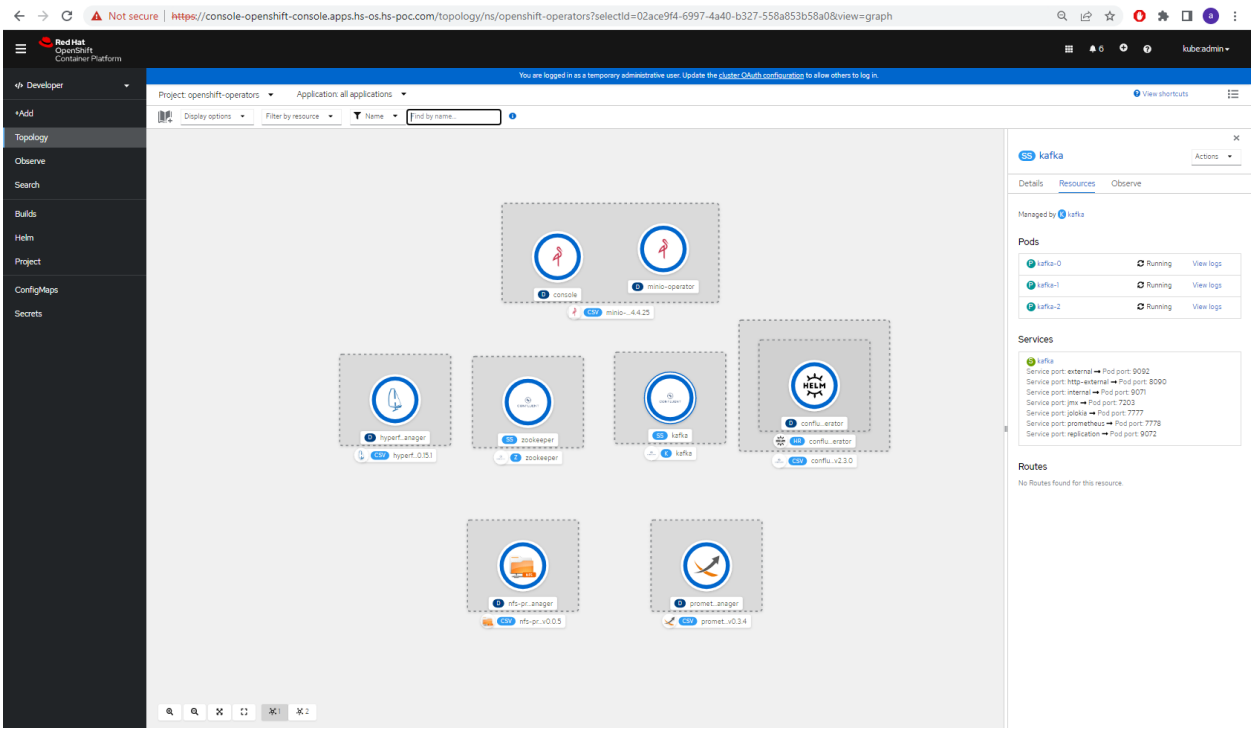

Figure 49 Kafka, MinIO object storage and Prometheus monitoring deployed in Red Hat OpenShift

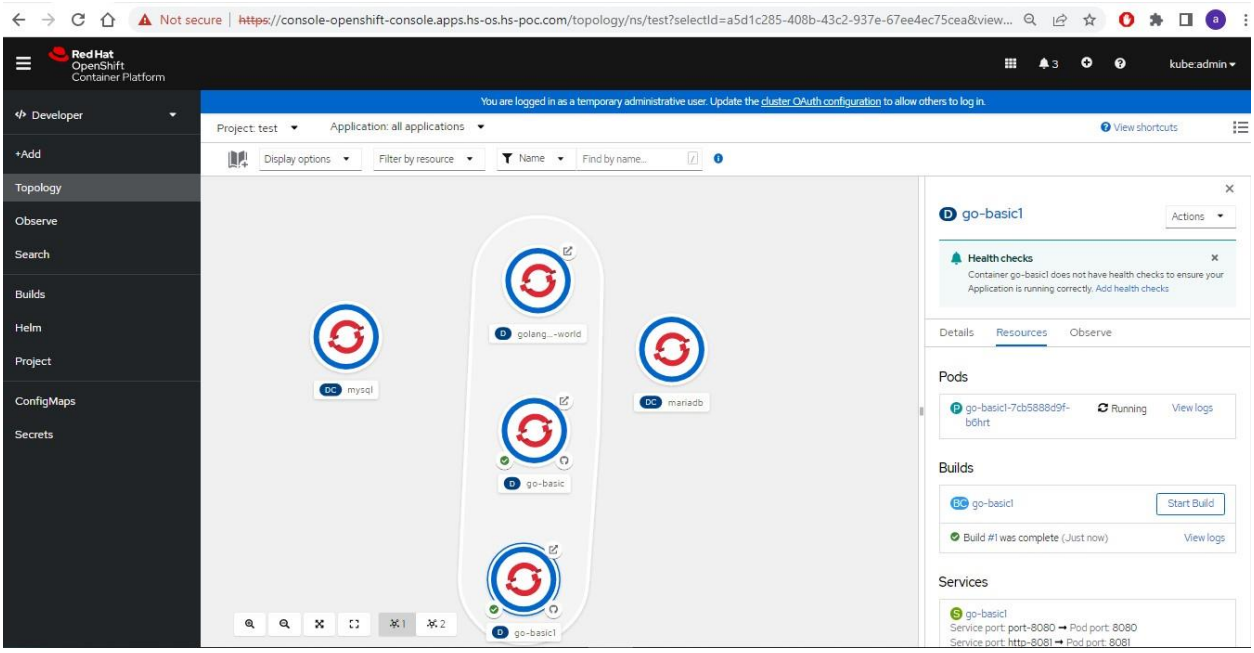

Figure 50 Go-Lang basic sample application and databases deployed in Red Hat OpenShift

Figure 51 Go-basic Test Application endpoint Routing

OpenShift command line terminal

```
bash-4.4 ~ $ oc get pods --namespace=openshift-storage
NAME                                                              READY   STATUS      RESTARTS        AGE
cluster-cleanup-job-worker3-kgg6g                                 0/1     Completed   0               7d21h
csi-addons-controller-manager-79c5cd7b5d-b8gt5                    2/2     Running     0               18h
csi-cephfsplugin-fhfch                                            3/3     Running     0               39h
csi-cephfsplugin-provisioner-5994bc45c9-td2l9                     6/6     Running     7               39h
csi-cephfsplugin-provisioner-5994bc45c9-z422k                     6/6     Running     0               39h
csi-cephfsplugin-snkds                                            3/3     Running     3               39h
csi-cephfsplugin-vl4df                                            3/3     Running     3               39h
csi-rbdplugin-8mh6c                                               4/4     Running     4               39h
csi-rbdplugin-lxfcg                                               4/4     Running     4               39h
csi-rbdplugin-provisioner-689b7b69fc-54d95                        7/7     Running     8               39h
csi-rbdplugin-provisioner-689b7b69fc-xc6cq                        7/7     Running     0               18h
csi-rbdplugin-vsdcn                                               4/4     Running     0               39h
noobaa-operator-7d6fb87df5-q9d64                                  1/1     Running     0               39h
ocs-metrics-exporter-859c6677d5-bs7jx                             1/1     Running     1               39h
ocs-operator-57555fd664-fqr44                                     1/1     Running     1               39h
odf-console-bff9bfc89-rd7pr                                       1/1     Running     1               39h
odf-operator-controller-manager-6d8d699db6-qmkg4                  2/2     Running     0               18h
rook-ceph-crashcollector-worker1-6dc594dbdd-6642p                 1/1     Running     0               39h
rook-ceph-crashcollector-worker2-8589c7bb85-nprhv                 1/1     Running     0               18h
rook-ceph-crashcollector-worker3-768b865d9b-l7h2w                 1/1     Running     0               16h
rook-ceph-mds-ocs-storagecluster-cephfilesystem-a-57ffdd74nkzxv   2/2     Running     0               39h
rook-ceph-mds-ocs-storagecluster-cephfilesystem-b-576d6b5fjsms9   2/2     Running     0               16h
rook-ceph-mgr-a-6f66f44898-64p5x                                  2/2     Running     0               16h
rook-ceph-mon-a-7d8dd8bc8b-g8chf                                  2/2     Running     0               39h
rook-ceph-mon-b-55f9479b77-crl5f                                  2/2     Running     2               39h
rook-ceph-mon-d-b7fc6f84b-qzwxp                                   2/2     Running     0               18h
rook-ceph-operator-6f8d6c5d5c-jq4r7                               1/1     Running     3 (129m ago)    39h
rook-ceph-rgw-ocs-storagecluster-cephobjectstore-a-7dbbf87xstmb   1/2     Running     197 (2m12s ago) 18h
bash-4.4 ~ $
```

Figure 52 OpenShift CLI listing Pods deployed in Red Hat OpenShift

6 ADDENDUM

Nodes involved in the Red Hat OpenShift cluster has the following Field Replaceable Units (FRU)

Hyper Scalers FRU Dump

Date

Modify: 2022-06-21 01:43:20.000000000 +0000

Update

Version: 3A06

Firmware Revision: 3.64

System

Product Name: QuantaPlex T43Z-2U 20S5ZCU0050

Chassis

Serial Number: QTFCYB213001E

Motherboard

Version: 31S5ZMB0010

Serial Number: QTFCYB2130014

Memory

Locator: CPU0_A0

Locator: CPU0_B0

Locator: CPU0_C0

Locator: CPU0_D0

Locator: CPU0_E0

Locator: CPU0_F0

Locator: CPU0_G0

Locator: CPU0_H0

Locator: CPU1_A0

Locator: CPU1_B0

Locator: CPU1_C0

Locator: CPU1_D0

Locator: CPU1_E0

Locator: CPU1_F0

Locator: CPU1_G0

Locator: CPU1_H0

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Manufacturer: Samsung

Module Manufacturer ID: Bank 1, Hex 0xCE

Memory Subsystem Controller Manufacturer ID: Unknown

Serial Number: 2214-509D0DF1

Serial Number: 2214-509D0D9F

Serial Number: 2214-509D0DF2

Serial Number: 2214-509F3CC2

Serial Number: 2214-509D0F20

Serial Number: 2214-509D0F33

Serial Number: 2214-509D0CD9

Serial Number: 2214-509D0E10

Serial Number: 2214-509D0CC9

Serial Number: 2214-509D0F44

Serial Number: 2214-509D0E0E

Serial Number: 2214-509D0E23

Serial Number: 2214-509F3C9A

Serial Number: 2214-509D0CA1

Serial Number: 2214-509D0F35

Serial Number: 2214-509D15E2

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Part Number: M393A4K40DB3-CWE

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Size: 32 GB

Non-Volatile Size: None

Volatile Size: 32 GB

Cache Size: None

Logical Size: None

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Speed: 3200 MT/s

Configured Memory Speed: 2666 MT/s

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Type Detail: Synchronous Registered (Buffered)

Processor

Version: Intel(R) Xeon(R) Silver 4316 CPU @ 2.30GHz

Version: Intel(R) Xeon(R) Silver 4316 CPU @ 2.30GHz

PCI Storage Card

Firmware

PCI Network Card

---

vendor: Mellanox Technologies

vendor: Mellanox Technologies

product: MT2894 Family [ConnectX-6 Lx]

product: MT2894 Family [ConnectX-6 Lx]

serial: b8:ce:f6:d6:28:fe

serial: b8:ce:f6:d6:28:ff

serial: 12:25:9f:d7:55:d1

Hard Disks

NAME HCTL TYPE VENDOR MODEL REV TRAN

sda 14:0:0:0 disk TOSHIBA TOSHIBA_USB_DRV PMAP usb

NAME SERIAL SIZE

sda 07088843CE1C4C56 15G

nvme0n1 S4EUNX0R755640L 232.9G

nvme2n1 S64GNE0RA00500 1.8T

nvme1n1 S4EUNX0RA39540R 232.9G

nvme5n1 S64GNE0RA00498 1.8T

nvme4n1 S64GNE0RA00493 1.8T

nvme3n1 S64GNE0RA00504 1.8T

GPU

---

7 COPYRIGHT AND LICENSING

8 REFERENCES

1. Hyperscalers, "About HS," [Online]. Available: https://www.hyperscalers.com/about-us-hyperscalers.
2. Chraibi, J. (2020) *A guide to enterprise kubernetes with OpenShift*, *Develop, deploy, and evolve with Red Hat Hybrid Cloud*. Available at: https://cloud.redhat.com/blog/enterprise-kubernetes-with-openshift-part-one?extIdCarryOver=true&sc_cid=701f2000001Css5AAC (Accessed: November 2, 2022).
3. *Features and benefits of Red Hat openshift* (no date) *Features and benefits of Red Hat OpenShift*. Available at: https://www.redhat.com/en/technologies/cloud-computing/openshift/features (Accessed: November 2, 2022).
4. Fruhlinger, J. (2022) *What is hyperconvergence?*, *Network World*. Network World. Available at: https://www.networkworld.com/article/3207567/what-is-hyperconvergence.html (Accessed: November 2, 2022).
5. *Ceph (software)* (2022) *Wikipedia*. Wikimedia Foundation. Available at: https://en.wikipedia.org/wiki/Ceph_(software) (Accessed: November 2, 2022).
6. *Containerization (computing)* (2022) *Wikipedia*. Wikimedia Foundation. Available at: https://en.wikipedia.org/wiki/Containerization_(computing) (Accessed: November 2, 2022).
7. *Pods and services* (no date) *Pods and Services - Core Concepts | Architecture | OpenShift Online Pro*. Available at: https://docs.openshift.com/online/pro/architecture/core_concepts/pods_and_services.html#:~:text=OpenShift%20Online%20leverages%20the%20Kubernetes,or%20virtual)%20to%20a%20container (Accessed: November 2, 2022).
8. *Kubernetes Components* (2022) *Kubernetes*. Available at: https://kubernetes.io/docs/concepts/overview/components/ (Accessed: November 2, 2022).
9. *Domain name system* (2022) *Wikipedia*. Wikimedia Foundation. Available at: https://en.wikipedia.org/wiki/Domain_Name_System (Accessed: November 2, 2022).
10. *Intro to ceph* (no date) *Intro to Ceph - Ceph Documentation*. Available at: https://docs.ceph.com/en/quincy/start/intro/ (Accessed: November 2, 2022).
11. *Chapter 7. infrastructure requirements red hat openshift Data Foundation 4.9* (no date) *Red Hat Customer Portal*. Available at: https://access.redhat.com/documentation/en-us/red_hat_openshift_data_foundation/4.9/html/planning_your_deployment/infrastructure-requirements_rhodf (Accessed: November 2, 2022).

Index