

SUSE CaaS Platform Integration with SUSE Enterprise Storage on HPE® Platforms

Written by:
Bryan Gartner, SUSE®



Table of Contents

Table of Contents	2
Use Case Overview	3
Target Audience	3
Integration	3
Solution Prerequisites	4
Integration Steps	6
Summary	9
Resources	9

Use Case Overview

Often, the initial deployments of newly containerized microservices utilize simple, stateless APIs and functions. Yet, over time, as previous monolithic applications are migrated towards containers or as more complex service topologies are created, the need for stateful containers arises. Due to the inherently resilient nature of container orchestration, having persistent volume stores that can be mounted wherever a container is launched becomes a requirement. This can come either in the form of dynamically created volumes or, for long-term retention, persistent volumes. To address this, one of the best integrations between Kubernetes-based container-as-a-service (CaaS) infrastructure, such as SUSE CaaS Platform, is with a Ceph-based software-defined storage infrastructure, such as SUSE Enterprise Storage.

Target Audience

This document is intended for IT decision makers, architects, system administrators and technicians who are implementing the combination of a flexible, Kubernetes-based CaaS platform and Ceph-based storage for persistent volumes. You should be familiar with all of the traditional IT infrastructure pillars—networking, computing and storage—along with each pillar’s specific, local use cases for sizing, scaling and limitations within their environments.

Integration

As mentioned in the Use Case Overview, requiring both dynamic and static volumes is key to running stateful containerized microservices. One agile method to accomplish that is to integrate a Kubernetes-based CaaS, such as SUSE CaaS Platform, with a Ceph-based, software-defined storage solution, such as SUSE Enterprise Storage. Coupling this with the inherent value-add and multiple HPE system platform product lines yields agile, cost-effective and performant options. These three solutions have the following key attributes, as described below.

Kubernetes-based Container-as-a-Service

SUSE CaaS Platform is an enterprise-class container management solution that enables IT and DevOps professionals to easily deploy, manage and scale container-based applications and services. It includes Kubernetes to automate lifecycle management of modern applications, as well as surrounding technologies that enrich Kubernetes and make the platform itself easy to operate. As a result, enterprises that use SUSE CaaS Platform can reduce application delivery cycle times and improve business agility. SUSE is focused on delivering an exceptional operator experience with SUSE CaaS Platform. With deep competencies in infrastructure, systems, process integration, platform security, lifecycle management and enterprise-grade support, SUSE aims to ensure that IT operations teams can deliver the power of Kubernetes to their users quickly, securely and efficiently. With SUSE CaaS Platform you can:

- Achieve faster time to value with an enterprise-ready container management platform that is built from industry-leading technologies and delivered as a complete package, with everything you need to quickly offer container services.
- Simplify management and control of your container platform with efficient installation, easy scaling and update automation.
- Maximize return on your investment with a flexible container services solution for today and tomorrow.

Ceph-based Persistent Storage

SUSE Enterprise Storage is an intelligent software-defined storage solution, powered by Ceph technology, that enables you to transform your enterprise storage infrastructure. It provides IT organizations with a simple-to-manage, agile

infrastructure for increased speed of delivery, durability and reliability. SUSE Enterprise Storage enables you to accelerate innovation, reduce costs and alleviate proprietary hardware lock-in by transforming your enterprise storage infrastructure with a truly open and unified intelligent software-defined storage solution. Let SUSE and our partners provide you with the most flexible enterprise-grade storage solution, offering:

- A single truly unified software-defined storage cluster that includes applications with object, block and file system storage. It provides ubiquitous and universal access to your legacy and modern applications, as well as automated durability of your data, with high availability and disaster recovery options.
- Unlimited scalability with a distributed storage cluster designed to scale to thousands of nodes and multi-hundred petabyte environments and beyond to meet your growing data requirements.
- Utilize industry-standard, off-the-shelf hardware that is at minimum 30 percent less expensive than average capacity-optimized solutions to drive significant CAPEX savings.
- A highly redundant storage infrastructure design that maximizes application availability with no single points of failure.
- Self-healing capabilities that minimize storage administration involvement and optimize data placement, enabling rapid reconstruction of redundancy and maximizing system resiliency and availability.
- Automated re-balancing and optimized data placement, with an easy-to-manage intelligent solution that continuously monitors data utilization and infrastructure without manual intervention and without increasing IT staff.

HPE Platform Portfolio

Utilizing the software solutions described above with the server platforms offered by HPE provides many alternatives for scale, cost-effectiveness and performance options that could align with local IT staff platform preferences. The [HPE server portfolio offers these options and features:](#)

- Density-optimized: high performance and efficiency for big data and the most demanding workloads
- Mission-critical: systems of intelligence to fuel your digital transformation in a world where time and data are the new currency and business continuity is expected
- Composable: fully adaptable and ready for Hybrid-IT to future-proof your data center for today's workloads and tomorrow's disruptors
- Internet of Things (IoT): realize the potential of the IoT to provide compute at the network edge
- Cloud: high-capacity, mass-compute open infrastructure with security and software to match
- Virtualized use cases

These software solutions have been tested across several HPE platforms, with reference configurations available. You can find these on both the [SUSE](#) and [HPE](#) alliance partner websites.

Note: Any [SUSE YES](#) certified HPE platform can be used for the physical nodes of either solution's deployment, as long as the certification refers to the major version of the underlying SUSE operating system required by the solution release.

Solution Prerequisites

Given the use case, the following software-defined solutions are assumed to be in place and running:

- SUSE CaaS Platform
- SUSE Enterprise Storage

While it is possible to implement the integration with the basic, minimal configurations of each solution, we recommend following the additional best practices noted below:

- Network and Network services

Routable Networks

- An intranet or client-access network to access both solutions.
- Each solution should have a "public" network that is routable to the other one and to any client-facing network, yet can be intranet, on-premises only. The integration access happens across this network.
- An internet-facing network, if needed, to access cluster resources from beyond an organization's internal networks.

Potentially Non-routed Networks

- A Baseboard Management Control (BMC) network, used to connect and control HPE iLO interfaces.
- SUSE CaaS Platform has a pair of internal cluster networks that can be non-routable outside the cluster, but they need to be sized according to the number of workload containers to be deployed.
- SUSE Enterprise Storage also has a private network, which is used for backend data replication and should have capacity planned to address recovery and the number of replicas.

Services

As with most software-defined, cluster-based solutions, the following network services should be accessible:

- One or more valid Network Time Protocol (NTP) services, with all the cluster nodes configured to use this.
- One or more valid Domain Name Service (DNS) services, with all the cluster nodes registered.
- Ensure access to software security updates and fixes by registering nodes to the SUSE Customer Center or by creating a local Subscription Management Tool service.
- A reliable Load Balancer (LB), to provide high-availability access to core Kubernetes Master API functions. This could be implemented with any number of technologies, including simple HAProxy software configurations.

System Platform

For each of the software-defined solutions, ensure that each respective node not only meets the minimum requirements for that role, but also has the resource capacity to accommodate the expected number and type of workloads. These minimum guidelines are detailed in the respective solution deployment guides. For both solutions, performance and log monitoring capabilities can be deployed to help access the utilization metrics over time.

These prerequisites are summarized in the following logical deployment diagram:

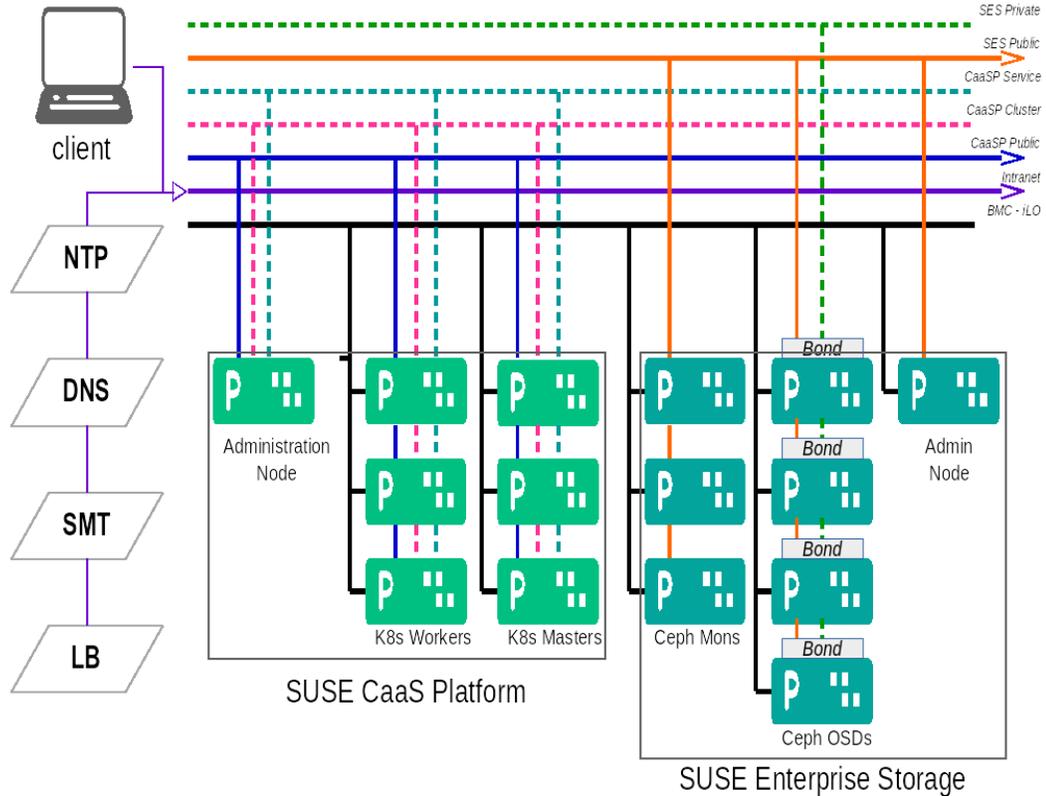


Figure 1: Logical Deployment Diagram. Deployment Logical View.

Integration Steps

The remainder of this document is meant as a companion guide to the official network, system and software product deployment documentation, citing specific settings as needed for this reference implementation. Default settings are assumed to be in use, unless otherwise cited, to accomplish the respective best practices and design decisions herein.

Note: In the process steps designated below, environment variables are used to highlight the values you can elect to change to match your local environment. Use caution if you are simply cutting-and-pasting from this document into a live environment without validating those values. Entries enclosed in < > also need to be adjusted to match your environment.

TIP: The prompt designated in the steps below help to illustrate where to run each command and the respective user access.

To create the necessary integration from a Kubernetes-based SUSE CaaS Platform cluster and a Ceph-based storage cluster such as SUSE Enterprise Storage, first start on the SUSE Enterprise Storage.

- Login to the respective Admin Node.
- Collect the list of SUSE Enterprise Storage Monitor Node IP Addresses from `/etc/ceph/ceph.conf`.
- Create and validate a dedicated storage pool (e.g., "caasp-pool") for rbd-based applications using the following commands. Adjust the placement group size (e.g., "512") as desired.

```
root@ses-admin # export CephPool="caasp-pool" CephPGSize="512"
root@ses-admin # ceph osd pool create ${CephPool} ${CephPGSize}
root@ses-admin # ceph osd pool ls
root@ses-admin # ceph osd pool application enable ${CephPool} rbd
root@ses-admin # ceph osd pool application get ${CephPool}
```

- Capture and encode the admin key, which will be used later.

```
root@ses-admin # ceph auth get-key client.admin | base64
```

- Create a specific user (e.g., "caasp") with the necessary capabilities to utilize the "caasp-pool." Then capture and encode user the key, which will be used later.

```
root@ses-admin # export CephUser="caasp"
root@ses-admin # ceph auth get-or-create client.${CephUser} mon 'allow r' osd 'allow class-read
object_prefix rbd_children, allow rwx pool=${CephPool}' -o ceph.client.${CephUser}.keyring
root@ses-admin # ceph auth get-key client.${CephUser} | base64
```

From the SUSE CaaS Platform, login to the respective Administration Node.

- Create the file *ceph-secret-admin.yaml*.

```
# file - ceph-secret-admin.yaml
apiVersion: v1
kind: Secret
metadata:
  name: ceph-secret-admin
  namespace: default
type: "kubernetes.io/rbd"
data:
  key: <insert output string from client.admin encode step>
```

- Apply and validate deployment of the file *ceph-secret-admin.yaml*.

```
root@caasp-admin # kubectl apply -f ceph-secret-admin.yaml
root@caasp-admin # kubectl get secrets
```

- Create the file *ceph-secret-cephuser.yaml*.

```
# file - ceph-secret-cephuser.yaml
apiVersion: v1
kind: Secret
metadata:
  name: ceph-secret-CephUser
  namespace: default
type: "kubernetes.io/rbd"
data:
  key: <insert output string from client.${CephUser} encode step>
```

- Apply and validate the file *ceph-secret-cephuser.yaml*.

```
root@caasp-admin # sed -i s/CephUser/`echo ${CephUser}`/g ceph-secret-cephuser.yaml
root@caasp-admin # kubectl apply -f ceph-secret-caasp.yaml
root@caasp-admin # kubectl get secrets
```

- Create the file `ses-rbd-storage-class.yaml`, using the SUSE Enterprise Storage Monitor Node IP addresses collected earlier.

```
# file - ses-rbd-storage-class.yaml
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  name: ses-rbd-sc
provisioner: kubernetes.io/rbd
parameters:
  monitors: <IPAddressMon1>:6789, <IPAddressMon2>:6789, <IPAddressMon3>:6789
  adminId: admin
  adminSecretName: ceph-secret-admin
  adminSecretNamespace: default
  pool: CephPool
  userId: CephUser
  userSecretName: ceph-secret-CephUser
```

- Apply and validate the storage volume claim file `ses-rbd-storage-class.yaml`.

```
root@caasp-admin # sed -i s/CephPool/`echo ${CephPool}`/g ses-rbd-storage-class.yaml
root@caasp-admin # sed -i s/CephUser/`echo ${CephUser}`/g ses-rbd-storage-class.yaml
root@caasp-admin # kubectl apply -f ses-rbd-storage-class.yaml
root@caasp-admin # kubectl get sc <-n default>
```

- Create the volume claim file `ses-rbd-persistent-volume-claim.yaml`, referencing the previous storage class.

```
# file - ses-rbd-persistent-volume-claim.yaml
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: ses-rbd-pvc
spec:
  storageClassName: ses-rbd-sc
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 1Gi # adjust size as needed
```

- Apply and validate the persistent volume claim file `ses-rbd-persistent-volume-claim.yaml`.

```
root@caasp-admin # kubectl apply -f ses-rbd-persistent-volume-claim.yaml
root@caasp-admin # kubectl get pvc <-n default>
```

Then, to validate a working integration, create a simple container, such as the busybox, using the persistent volume claim. Refer to the "Creating Pods with Persistent Volumes" section of the [SUSE CaaS Platform Administration Guide](#).

As needed, both SUSE CaaS Platform and SUSE Enterprise Storage can be scaled with additional resource nodes to provide more capacity, capabilities and performance. Such operations are described in the respective administration guides.

Summary

At this point, you should have a working integration between SUSE CaaS Platform and SUSE Enterprise Storage. Volumes for containerized workloads can now be utilized as needed and made available to wherever in the Kubernetes cluster the workload is deployed. Beyond providing the Ceph-based RBD persistent storage for stateful containers, this integration can also be utilized as the backing store for monitoring and performance logs, for tools such as Prometheus and for visualization tools, such as Grafana, to access as well.

Resources

HPE Platforms

- [SUSE Linux Enterprise Server Certification and Support](#), [SUSE YES Certified Bulletin Search](#)

SUSE CaaS Platform

- Product Documentation: <https://www.suse.com/documentation/suse-caasp-3/>

SUSE Enterprise Storage

- Product Documentation: <https://www.suse.com/documentation/suse-enterprise-storage-5/>