

# SUSE CaaS Platform 3 with HPE Synergy

Written by:

Bryan Gartner, SUSE®



# Table of Contents

- Table of Contents .....2
- 1. Executive Summary .....3
- 2. Target Audience.....3
- 3. Solution Overview .....3
- 4. Solution Components.....4
  - 4.1 Facility .....4
  - 4.2 Network.....4
  - 4.3 Computing.....4
  - 4.4 Storage .....6
  - 4.5 Software .....6
- 5. Solution Details .....9
  - 5.1 Deployment Flow ..... 10
  - 5.2 Additional Considerations ..... 12
- 6. Conclusion ..... 13
- 7. Resources and Additional Links ..... 13
  - Table 7.1: Bill of Materials - HPE Synergy* ..... 13
  - Table 7.2: Bill of Materials - Software Number*..... 13

## 1. Executive Summary

This reference configuration is designed to help organizations plan and install an on premises container-as-a-service (CaaS) platform infrastructure to deploy, orchestrate and manage containerized workloads. It provides an implementation overview for a production-ready deployment, with design considerations, implementation suggestions and best practices.

For most enterprise-level businesses, the demand for an on-premise, infrastructure to deploy and orchestrate containers is increasing rapidly, allowing the hosting of both production and other phases of workloads development. As with any software-defined infrastructure, all the classic IT disciplines around networking, computing and storage are involved and need to be effectively planned and integrated.

Implementing an intelligent, software-defined infrastructure consisting of SUSE CaaS Platform technology and HPE® Synergy system hardware enables you to transform your enterprise. The inherent composability of this solution delivers CaaS to help reduce costs, while still providing the flexibility to keep up with your future demands. Leveraging this tested approach, you will have the confidence to deploy a working solution in an agile manner and be able to maintain and scale it as needed over time.

## 2. Target Audience

This document is intended for IT decision makers, architects, system administrators and technicians who are implementing a flexible, software-defined CaaS platform. You should be familiar with the traditional IT infrastructure pillars—networking, computing and storage—along with the pillars' local use cases for sizing, scaling and limitations within each pillars' environments.

## 3. Solution Overview

Even on the journey to a full [Cloud Native Landscape](#), the classic IT pillars are still in play. General requirements include the need for a small, purpose-built operating system with a container runtime engine and a container orchestration platform to distribute workloads across a target, clustered instance. The dominant technology for container orchestration is Kubernetes. With its large community of developers and a plethora of features and capabilities, Kubernetes has become the defacto standard and is included across most CaaS platforms. With all of these attributes in place, both a developer and operations teams can effectively deploy, manage and deliver functionality to their end users in a resilient and agile manner.

**Note:** As a further reference, the [National Institute of Standards and Technology's \(NIST\) Definition of Microservices, Application Containers and System Virtual Machines](#) describes the important characteristics of application containers.

While the footprint of each containerized microservice is typically much smaller than a similar bare-metal or virtual machine workload, all of the IT pillars must still accommodate the sheer scale of such high-volume workloads. Core networking technologies, including high-speed, scalable network switches and network interfaces, must provide the foundational bandwidth for compute resource nodes, the workload and non-direct attached storage resources. Physical resource nodes or even virtual machines running atop a hypervisor can deliver the CaaS functionality.

A composable infrastructure such as HPE Synergy provides an ideal target platform, enabling you to start small and scale out over time with inherent consistency. This platform addresses both network and storage resources with industry-

leading x86-based servers, offering balanced performance and efficiency. Finally, it is the CaaS software—the SUSE CaaS Platform—that provides the user interface for setup, configuration, maintenance and long-term operation of the core software-defined infrastructure requirements, bonding them into a cohesive service offering. These components are described in the following sections.

## 4. Solution Components

### 4.1 Facility

While beyond the scope of this document, the heating, ventilation and air conditioning (HVAC) requirements of hosting such an infrastructure solution should be carefully considered and planned. To aid in determining the power requirements for system deployment, use the online or downloadable version of [HPE Power Advisor](#). Using this tool, you can plan for the needs of your solution and order the correct Power Distribution Unit (PDU) to account for the local power conditions and connections in the final installation location.

### 4.2 Network

Networking and the associated services are the technology components that typically require the most advanced planning. Connectivity, capacity and bandwidth requirements for a CaaS infrastructure have a fair amount of complexity, especially when integrated with an existing IT infrastructure and accounting for the magnitude of microservices being deployed.

With the inherent features and composability of HPE Synergy, much of the configuration complexity for network interconnects can easily be codified into a template for each target resource node. This includes the mapping of network interfaces to network fabrics within a collection of managed HPE Synergy frames. The baseline network bandwidth provided by HPE Synergy, for each resource node and collectively within the frame, is sufficient for a CaaS deployment. While beyond the scope of this document, the only remaining consideration is for capacity and bandwidth for the interconnecting top-of-rack switches. In addition, access to the HPE iLO for each resource node can be effectively coordinated and managed through the user interface of HPE Synergy Composer.

### 4.3 Computing

HPE Synergy, the first Composable Infrastructure, empowers IT to create and deliver new value easily and continuously. This single infrastructure reduces operational complexity for traditional workloads 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 any configuration for any application. It also enables a broad range of workloads—from bare metal, to virtual machines, to containers, to operational models such as hybrid cloud and DevOps. HPE Synergy enables IT to rapidly react to new business demands with the following components:

#### HPE Synergy 12000 Frame

- Is uniquely architected as composable infrastructure (CI) to match the powerful infrastructure-as-code capabilities of the HPE intelligent software architecture. Flexible access to compute, storage and fabric resources enables customers to maximize their resource usage and to repurpose underutilized resources. Linking together multiple HPE Synergy Frames efficiently scales the infrastructure with a dedicated, single view of the entire management network.

- Allows for creating multiple composable domains in the infrastructure to efficiently deliver available resources to the business. HPE Synergy Frames reduce complexity by using intelligent auto-discovery to find all available resources and accelerate workload deployments. This drives IT efficiency as the business grows and delivers balanced performance across resources to increase solution effectiveness.

#### HPE Synergy Composer

- Provides enterprise-level management to compose and deploy system resources to your application needs. This management appliance uses software-defined intelligence with embedded HPE OneView to aggregate compute, storage and fabric resources in a manner that scales to your application needs, instead of restricting you to the fixed ratios 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.
- Includes HPE OneView, a comprehensive unifying management interface designed for converged infrastructure management. A unifying platform increases the productivity of every member of the internal IT team across servers, storage and networking. By streamlining processes, incorporating best practices and creating a new holistic way to work, HPE OneView provides organizations with more efficiency. It is designed for open integration with existing tools and processes to extend these efficiencies.
- Is instrumental for the deployment and management of HPE servers and enclosure networking. HPE Synergy Composer 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 server hardware blades and rack servers, networking interconnects and computing resources.
- Offers a uniform console 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 eliminates the need for the administrator to enter and maintain the same configuration data more than once and keeps all versions up to date. It encapsulates and abstracts many underlying tools behind the integrated resource model, enabling the administrator to operate with new levels of simplicity, speed and agility to provision, monitor and maintain the solution.

#### HPE Synergy ImageStreamer

- Can quickly provision an operating environment across a large number infrastructure blocks or nodes. HPE Synergy ImageStreamer can deploy and update many systems quickly (possibly as fast as you can reboot servers) to rapidly expand or change environments. This management capability is implemented using redundant physical appliances, enabling production environments to maintain high availability in operations. These management appliances are automatically set up with active-active storage to control and protect your image repository. Your image content might contain an OS or even a complete application stack. Your images can be swiftly applied to multiple compute nodes to optimize your IT service deliveries.

**Note:** At this point, HPE Synergy ImageStreamer does not yet support Btrfs-based operating system node deployments, which is the basis for SUSE CaaS Platform, so this technology is not utilized or detailed in this solution document.

#### HPE Synergy 480 Gen10 Compute Module

- Delivers an efficient and flexible two-socket workhorse to support most demanding workloads. HPE Synergy powered by the Intel® Xeon® Scalable Family of processors, up to 3TB DDR4, more storage capacity and controllers, and a variety of GPU options within a composable architecture. HPE Synergy 480 Gen10 Compute Module is the ideal platform for general-purpose enterprise workload performance now and in the future.
- For this the implementation, HPE Synergy 480 Gen10 servers were utilized for all node roles. Example configurations are included in the “Resources and Additional Links” section.

**Note:** Any [SUSE YES](#) certified HPE platform, such as the HPE Synergy 480 Gen10, can be used for the physical nodes of this deployment, as long as the certification refers to the major version of the underlying SUSE operating system required by the SUSE CaaS Platform release.

## 4.4 Storage

Each of the resource nodes is expected to have some local, direct attached storage, which is predominantly used for the node's operating system. For this deployment, a pair of disk drives, configured as a RAID1 volume, via the HPE Synergy Composer templates for the operating system helps to provide fewer points of failure.

This storage configuration is also used for the ephemeral container images running on any given node. These are simply copies of the reference container image version, obtained from a centralized location such as a registry, which is utilized during the workload's runtime. After the workload is terminated or moved, this disk space is reclaimed and becomes available for subsequent usages.

**TIP:** This local storage element can become a limiter for both the number and size of container workloads running on a given node; it can also become a performance limiter. As such, having performant local storage devices (such as solid-state drives) is encouraged. Monitoring of workloads and scale can also help to determine if capacity or performance constraints are becoming an issue.

For any stateful microservice deployments, additional storage technologies should be integrated to house persistent volumes. Both static and dynamic volume claims can be utilized via the modular Container-Storage-Interface (CSI) technologies. A commonly utilized approach for these non-cloud-native workloads that require persistent storage is a Ceph-based, software-defined storage infrastructure. This solution can provide block and file or other protocol integration options with a Kubernetes-based CaaS platform.

**Note:** Such integrations, with solutions such SUSE Enterprise Storag, are detailed in other reference documents.

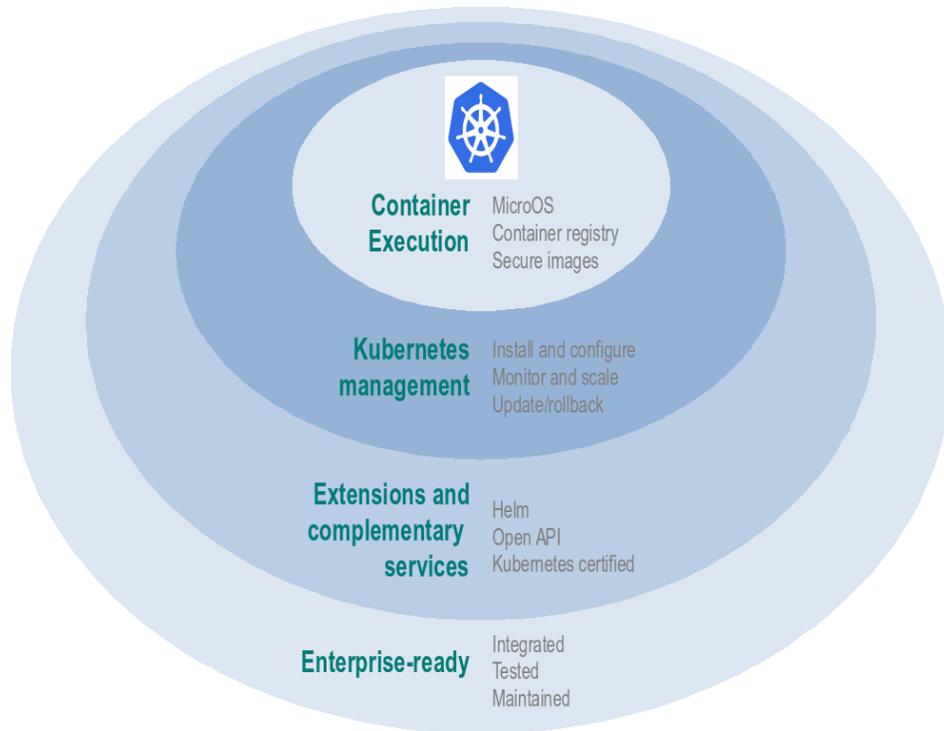
## 4.5 Software

SUSE CaaS Platform is an enterprise-class container management solution that enables IT and DevOps professionals to more easily deploy, manage and scale container-based applications and services. It includes Kubernetes to automate the lifecycle management of modern applications and 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, 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.



**Figure 4.1:** SUSE CaaS Platform Features.

#### Host Operating System

- Typically a small footprint operating system installation, having just enough functionality to support the container runtime engine, leaving as many CPU, memory and I/O resources available for the containerized workloads as possible.
- SUSE currently delivers this as MicroOS, a read-mostly operating system based on SUSE Linux Enterprise Server built for microservices. This is complemented by a distributed key-value store provided by etcd to retain persistent configuration data. In addition, MicroOS provides a snapshot-driven, transactional-update methodology to perform atomic upgrades.

#### Container Runtime Engine(s)

- Comprised of both a format for and service to run containerized applications on top of the host operating system.
- SUSE provides support for Docker® Community Edition Engine, a format for application containers.
- SUSE also offers a technical preview of CRI-O, an implementation of Container Runtime Interface (CRI), designed specifically for Kubernetes as a lightweight alternative, using Open Container Initiative (OCI) images.

## Container Networking

- An intra-cluster service and overlay network used for container and orchestration communication.
- SUSE currently utilizes the Container Network Interface (CNI) with the Flannel plugin and a configuration management, web interface to set up and deploy these networks.

## Container Orchestration

- A service to manage deployments of containerized workloads, known as Kubernetes or K8s, which is the current, de facto standard open source implementation for large-scale container orchestration.
- SUSE currently delivers and supports a Cloud-Native Computing Foundation (CNCF) certified Kubernetes distribution. It includes a role-based access control technology to limit access to resources, functions and services as desired.

In addition to these high-level architectural components, SUSE CaaS Platform provides and relies upon the following types of nodes and roles:

## Administration Node

- Provides a cluster infrastructure management system. Each service is ran as a container on this host. It includes configuration management, plus a web-based dashboard to manage other node types within the cluster.

## Master Node(s)

- Oversees Kubernetes container workload orchestration services across the cluster and manages the Kubernetes Worker Nodes.

## Worker Node(s)

- Is where the user-defined containerized workloads and services run in Kubernetes pods. A pod is the smallest and simplest unit that you create and deploy as a running process for Kubernetes, encapsulating one or more container images and resources.

## Client System

- One or more existing systems, with your choice of operating system, used to access the cluster and various services provided from a command line, via kubectl and helm and a web browser.

**Note:** Refer to the "Architectural Overview" section of the [SUSE CaaS Platform Deployment Guide](#) for more details.

The following networking requirements must be in place for a successful deployment:

- Cluster network
  - Choose a subnet range that will span the total number of cluster nodes. This range can also be segmented or secured for access to specific node roles as desired.
  - All of the cluster node types must be able to communicate on the same network, with this primary network interface card. A client system with similar network access is also required for command line and web browser interaction with the cluster, especially during setup.
  - Higher-speed network interface cards (minimum of 10 GigE and above) and switching are preferred. This is because the number of containerized workloads can be high and the cards share this infrastructure capacity, both from an external and intra-cluster perspective.
- Internal networks

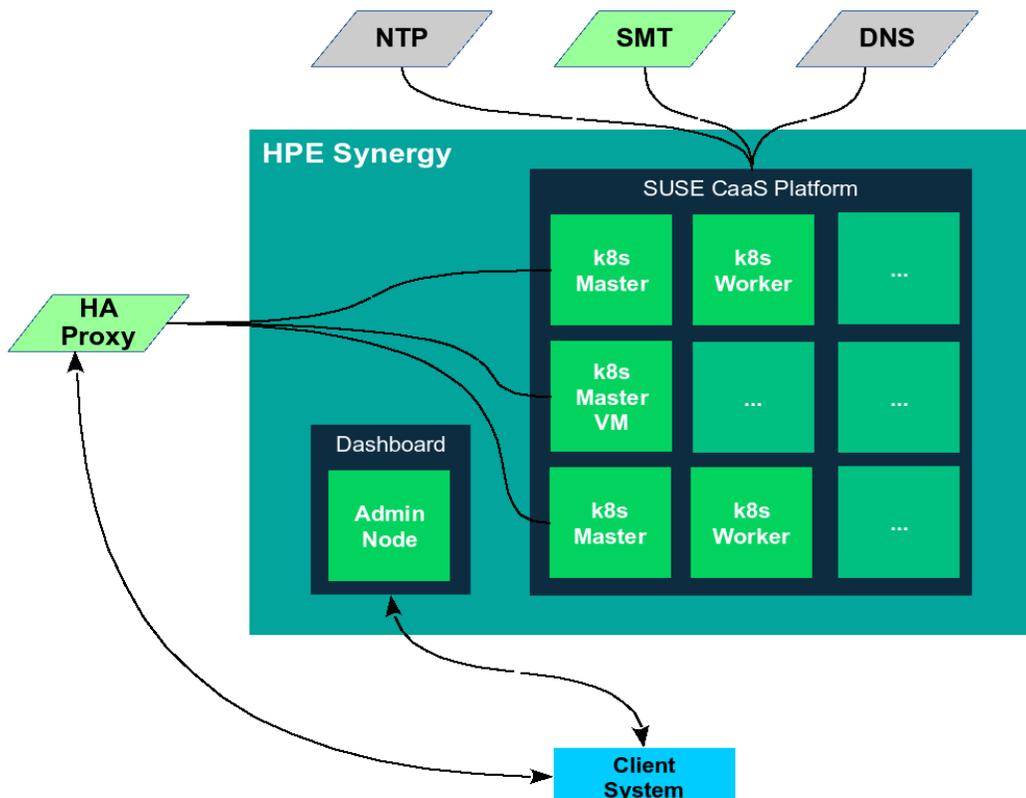
- Known as the Overlay and Service networks, these are used by Kubernetes and the underlying Flannel network plug-in to manage the internal cluster and container connections. These are implemented with bridges to the main cluster network.
- Core Network Infrastructure Components/Services
  - Domain Name Service (DNS): an external network-accessible service to map IP addresses to hostnames.
  - Network Time Protocol (NTP): an external network-accessible service to obtain and synchronize system times to aid in timestamp consistency.
  - Software Update Service: access to a network-based repository for software update packages. This can be accessed directly from each node via registration to the [SUSE Customer Center](#) or from local servers running a [SUSE Subscription Management Tool](#) (SMT) instance. As each node is deployed, it can be pointed to the respective update service; and the update notification and applications of updates can be managed by the configuration management web interface.

**Note:** Refer to the "Networking Requirements" section of the [SUSE CaaS Platform Deployment Guide](#) for more details and port specifics.

**Important:** These internal network ranges should be planned prior to deployment. They are usually non-routable network ranges and cannot be changed without redeploying the entire cluster.

## 5. Solution Details

This document focuses on a new SUSE CaaS Platform deployment, which could be scaled over time. A starting point could be a single hypervisor host with virtual machines set up for each of the needed roles to create a basic, minimal cluster to simply evaluate functionality. Over time, more physical nodes could be added to augment the cluster or to



replace some of the initial, virtual machine-based roles. To provide a production-ready cluster and take advantage of the HPE Synergy platform and its composable features, the following figure shows the target logical cluster deployment:

**Figure 5.1:** *Deployment Logical View.*

## 5.1 Deployment Flow

The remainder of this section is 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.

Given the detailed information contained in the [SUSE CaaS Platform Deployment Guide](#), only the following additional, incremental configurations and modifications are described below:

- Pre-Installation checklist. Obtain the following software media and documentation artifacts:
  - Download the CaaS Platform x86\_64 install media (DVD1) from the [SUSE CaaS Platform site](#).
  - To ensure access to technical support and software updates, utilize either trial or purchased subscriptions for all resource nodes. The bill of materials section in the “Resources and Additional Links” section outlines the type and quantity of subscriptions needed.
  - Obtain and preview the [SUSE CaaS Platform documentation](#) focusing on the:
    - Installation Quick Start
    - Deployment Guide
- Create an HPE Synergy Composer template to address:
  - Minimum CPU/Memory/Disk/Networking requirements of the SUSE CaaS Platform being applied to any available HPE Synergy resource nodes. See the [SUSE CaaS Platform Deployment Guide](#) for this information.
  - Consistent and up-to-date versions of BIOS/uEFI/device firmware, to reduce potential troubleshooting issues later.
  - The BIOS/uEFI settings are reset to defaults for a known baseline, consistent state or with desired values.
  - Use of either local or shared direct-attached storage elements, in a RAID1 configuration, to be used as the operating system installation target.
  - Network interfaces, including HPE iLO, configured as desired to match the target network topology and connectivity.
- Verify that the following considerations for various network service configurations are ready to access:
  - Ensure that you have access to a valid, reliable, external NTP service; this is a critical requirement for all nodes.
  - Set up external DNS A records for all nodes. Decide on subnet ranges and configure the switch ports accordingly to match the nodes in use.
  - Ensure that you have access to software security updates and fixes by registering nodes to the [SUSE Customer Center](#) or by creating a local [Subscription Management Tool service](#).
- Administration Node Installation

**Note:** *For the sake of consistency in this document, the installation process used across all nodes, whether virtual or physical, was done from ISO images. Other options are available as noted in the [SUSE CaaS Platform Deployment Guide](#).*

- Install the SUSE CaaS Platform administration node either as a virtual machine or as the first physical resource node.

- Apply the HPE Synergy Composer template to the target resource node in the HPE Synergy frame.
- If a physical node, configure the following HPE iLO virtual media image for use: SUSE CaaS Platform ISO image (bootable).
- Follow the "Installing the Administration Node" steps described in the [SUSE CaaS Platform Deployment Guide](#).
- When the installation is complete and the system reboots, use a client system to access the Velum Dashboard web-interface at the Fully-Qualified Domain Name (FQDN) of the Administration Node.
- Continue the setup described in the "Configuring the Administration Node" section of the [SUSE CaaS Platform Deployment Guide](#). Ensure that the following items are addressed:
  - On the home page, select "Create Admin" to set up an account with a valid email address and password. Then log in and do the following:
    - Check the "Install Tiler (Helm's server component)" box in "Cluster Services"; this will be used extensively later.
    - Ensure that the Overlay and Service network settings match the desired values, if the default values are not satisfactory.
    - Select the desired container runtime. For this deployment, the Docker open source engine was used.
  - On the "Bootstrap your CaaS Platform" page:
    - Refer to the location of the AutoYast file. If you would like to automate other resource node installations.
- Master and Worker Nodes Installation

**Note:** To provide more resiliency and failover capabilities, it is recommended to install three Master Nodes, plus two Worker Nodes. Additional nodes, for master or worker roles, can be added to the cluster or swapped into the cluster at any point.

- Begin the installation of the remaining resource nodes; at least three are required for a minimal cluster.
- As with the installation of the Administration Node, apply the respective HPE Synergy Composer template and use the HPE iLO virtual media.
- Complete the installation steps as described in the "Installing Master and Worker Nodes" section of the [SUSE CaaS Platform Deployment Guide](#).
- HAProxy Setup

**Note:** As noted in the [SUSE CaaS Platform Deployment Guide](#), HAProxy should be configured to provide load balancing for the Kubernetes API server (port 6443) and for DEX (OIDC Connect, port 32000). Also, ensure that the "stats" configuration stanza does not conflict with any other services (default of port 80) on the target node.

- Either a load-balancer or HAProxy is the suggested approach to make the cluster accessible to client systems and, more specifically, to some of the core Kubernetes Master Node API functions. It does this via a virtual IP address, which then sends the call to any active masters. While not required for a single master cluster, setting this up in advance allows for expansion and substitutions later on.
- For HAProxy, this process can be run on any host or virtual machine that has access to the cluster network. The steps to deploy this service are described in the "Configuring Load Balancing with HAProxy" section of the [SUSE Linux Enterprise High Availability Extension 12 SP4 Administration Guide](#).

- Bootstrap the Cluster

**Note:** If using a SUSE Linux Enterprise Server 12 or newer release host as the client, both the `kubectl` and `helm` commands can be found [in SUSE Package Hub](#). When all the cluster nodes are installed and rebooted, use the client system again to log in and access the Velum Dashboard web-interface at the FQDN of the Administration Node to continue the cluster formation.

- There should be a corresponding list of nodes in the "Pending Nodes" section, so select "Accept All Nodes."
- Designate the "Master" and "Worker" to the respective nodes, then click "Next."
- Enter the Kubernetes Master LB/HAProxy virtual IP FQDN setting for the "External Kubernetes API FQDN."
- Enter the FQDN of the Administration Node in "External Dashboard FQDN." Then select "Bootstrap Cluster."
- When this process completes, you should have a fully functional SUSE CaaS Platform cluster. You can validate this by logging into the Administration Node and running:
  - `root@caasp-admin# kubectl cluster-info`
  - `root@caasp-admin# kubectl get nodes`
  - `root@caasp-admin# kubectl get pods -n kube-system`
- You can also validate this by logging into a client system:
  - Using a web browser, log into the Velum Dashboard web-interface with the admin credentials at the FQDN of the Administration Node.
  - Download the `kubeconfig` file and put a copy in the default location of `~/.kube/config`,
  - Ensure that the client system has `kubectl` installed. Then run the same set of `kubectl` commands from the previous section.
- To further validate the functionality of the cluster, you can find a representative set of the upstream Kubernetes conformance test suite at Heptio / Sonobuoy (<https://github.com/heptio/sonobuoy>). This can be easily run via a browser on the client system that has access to the admin `kubeconfig` file, as noted in the "Getting Started" section of this site.

After completing the steps described in this document, you should have a working CaaS platform that is scalable through the addition of more resource nodes, as needed. SUSE CaaS Platform provides a complete suite of the necessary software and processes that leverage the composability aspect of HPE Synergy to create a production-worthy and agile platform for the deployment of containerized microservice workloads.

## 5.2 Additional Considerations

Review the information in the tables below to understand the administration aspects of the cluster. In addition, review the following sections of the [SUSE CaaS Platform Administration Guide](#):

- Changing or scaling the cluster node count by adding and removing resource nodes
- Updating the cluster node's operating system and core packages
- Setting up resource metrics gathering and visualization
- Addressing security by creating users and managing fine-grained, role-based access controls
- Creating and utilizing a local container image registry
- Accessing log locations and creating collection options

## 6. Conclusion

After understanding and working through the steps described in this document, you should have a working container-as-a-service platform that is scalable through the addition of even more resource nodes, as needed. SUSE CaaS Platform provides a complete suite of the necessary software and processes which leverages the composability aspect of HPE Synergy to create a production-worthy and agile platform for deployment of containerized microservice workloads.

## 7. Resources and Additional Links

### HPE Synergy

- [HPE Synergy 480 Gen10 System](#)

**Table 7.1:** Bill of Materials - HPE Synergy

Role	Quantity	Product Number	Product Name	Description
Frame	1	797740-B22	HPE Synergy 12000 Frame Configure-to-order Frame with 2x FLM 6x Power Supplies 10x Fans	
	1	804353-B22	HPE Synergy Composer	add second module for HA configuration
Administration Node	1	871946-B21	HPE Synergy 480 Gen10 3104 1P 16GB-R S100i SATA Entry Compute Module	
Master Node	1	871946-B21	HPE Synergy 480 Gen10 3104 1P 16GB-R S100i SATA Entry Compute Module	recommend 3 for HA configuration
Worker Node	1	871946-B21	HPE Synergy 480 Gen10 3104 1P 16GB-R S100i SATA Entry Compute Module	minimum of 2, but can scale to <100

### CaaS Platform

- [Product Documentation](#)

**Table 7.2:** Bill of Materials - Software Number

Role	Quantity	Product Number	Product Name	Description
Software	1	Q9K53AAE	SUSE CaaS Platform, 1-2 Sockets 3yr 24x7	for each node of the deployment