



# Installation Quick Start

---

SUSE CaaS Platform 3



## Installation Quick Start

SUSE CaaS Platform 3

by Liam Proven, Christoph Wickert, Markus Napp, Sven Seeberg-Elverfeldt, and Jana Halačková

Publication Date: January 30, 2019

SUSE LLC

10 Canal Park Drive

Suite 200


Cambridge MA 02141

USA

<https://www.suse.com/documentation> 

Copyright © 2006– 2019 SUSE LLC and contributors. All rights reserved.

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or (at your option) version 1.3; with the Invariant Section being this copyright notice and license. A copy of the license version 1.2 is included in the section entitled “GNU Free Documentation License”.

For SUSE trademarks, see <http://www.suse.com/company/legal/> . All other third-party trademarks are the property of their respective owners. Trademark symbols (®, ™ etc.) denote trademarks of SUSE and its affiliates. Asterisks (\*) denote third-party trademarks.

All information found in this book has been compiled with utmost attention to detail. However, this does not guarantee complete accuracy. Neither SUSE LLC, its affiliates, the authors nor the translators shall be held liable for possible errors or the consequences thereof.

# Contents

## What is SUSE CaaS Platform 3? iv

- 1 SUSE CaaS Platform System Requirements 1**
  - 1.1 Cluster Size Requirements 1
  - 1.2 Minimum Node Specification 1
  - 1.3 Network Requirements 3
  - 1.4 Limitations 4
- 2 Installation 5**
  - 2.1 Installing The Administration Node 5
  - 2.2 Configuring the Administration Node 12
  - 2.3 Installing Cluster Nodes 15
    - Manual Installation 16 • Installation With AutoYaST 16
  - 2.4 Bootstrapping The Cluster 18
- 3 Configuration 21**
  - 3.1 Deployment Scenarios Requiring Second-stage Configuration 21
  - 3.2 Defaults 21

## What is SUSE CaaS Platform 3?

SUSE SUSE CaaS Platform is a Cloud-Native Computing Foundation (CNCF) certified Kubernetes distribution on top of SUSE MicroOS. SUSE MicroOS is a minimalist operating system based on SUSE Linux Enterprise 12 SP3, dedicated to hosting containers. SUSE MicroOS OS inherits the benefits of SUSE Linux Enterprise in the form of a smaller, simpler, and more robust operating system, optimized for large, clustered deployments. It also features an atomic, transactional update mechanism, making the system more resilient against software-update-related problems. SUSE CaaS Platform automates the orchestration and management of containerized applications and services with powerful Kubernetes capabilities, including:

- Workload scheduling places containers according to their needs while improving resource utilization.
- Service discovery and load balancing provides an IP address for your services, and distributes load behind the scenes.
- Application scaling up and down accommodates changing loads.
- Non-disruptive rollout/rollback of new applications and updates enables frequent changes without downtime.
- Health monitoring and management supports application self-healing and ensures application availability.

In addition, SUSE CaaS Platform simplifies the platform operator's experience, with everything you need to get up and running quickly, and to manage the environment effectively in production. It provides:

- A complete container execution environment, including a purpose-built container host operating system (SUSE MicroOS), container runtime, and container image registries.
- Enhanced datacenter integration features that enable you to plug Kubernetes into new or existing infrastructure, systems, and processes.
- Application ecosystem support with SUSE Linux Enterprise container base images, and access to tools and services offered by SUSE Ready for CaaS Platform partners and the Kubernetes community.
- End-to-End security, implemented holistically across the full stack.


- Advanced platform management that simplifies platform installation, configuration, re-configuration, monitoring, maintenance, updates, and recovery.
- Enterprise hardening including comprehensive interoperability testing, support for thousands of platforms, and world-class platform maintenance and technical support.

You can deploy SUSE CaaS Platform onto physical servers or use it on virtual machines. After deployment, it is immediately ready to run and provides a highly-scalable cluster.

While SUSE CaaS Platform inherits benefits of SUSE Linux Enterprise and uses tools and technologies well-known to system administrators such as `cloud-init` and Salt, the main innovation (compared to SUSE Linux Enterprise Server) comes with **transactional updates**. A transactional update is an update that can be installed when the system is running without any downtime. A transactional update can be rolled back, so if the update fails or is not compatible with your infrastructure, you can restore the previous system state.

SUSE CaaS Platform uses the Btrfs file system with the following characteristics:

- The root filesystem and its snapshots are read-only.
- Sub-volumes for data sharing are read-write.
- SUSE CaaS Platform introduces overlays of the `/etc` directories used by `cloud-init` and Salt.

For more information, including a list of the various components which make up SUSE CaaS Platform please refer to the Release Notes on <https://www.suse.com/releasenotes/> .

# 1 SUSE CaaS Platform System Requirements

Before you begin the installation, please make sure your system meets all the requirements listed below.

## 1.1 Cluster Size Requirements

SUSE CaaS Platform is a dedicated cluster operating system and only functions in a multi-node configuration. It requires a connected group of four or more physical or virtual machines.

The minimum supported cluster size is four nodes: a single administration node, one master node, and two worker nodes.



### Note: Test and Proof-of-Concept Clusters

It is possible to provision a three-node cluster with only a single worker node, but this is not a supported configuration for deployment.

For improved performance, multiple master nodes are supported, but there must always be an odd number. For cluster reliability, when using multiple master nodes, some form of DNS load-balancing should be used.

Any number of worker nodes may be added up to the maximum cluster size. For the current maximum supported number of nodes, please refer to the Release Notes on <https://www.suse.com/releasesnotes/>.

## 1.2 Minimum Node Specification

Each node in the cluster must meet the following minimum specifications. All these specifications must be adjusted according to the expected load and type of deployments.

(v)CPU

- 4 Core AMD64/Intel\* EM64T processor
- 32-bit processors are not supported

## Memory

- 8 GB

Although it may be possible to install SUSE CaaS Platform with less memory than recommended, there is a high risk that the operating system will run out of memory and subsequently causes a cluster failure.



### Note: Swap partitions

Kubernetes does not support swap.

For technical reasons, an administration node installed from an ISO image will have a small swap partition which will be disabled after installation. Nodes built using AutoYaST do not have a swap partition.

## Storage Size

- 40 GB for the root file system with Btrfs and enabled snapshots.



### Note: Cloud default root volume size

In some Public Cloud frameworks the default root volume size of the images is smaller than 40GB. You must resize the root volume before instance launch using the command line tools or the web interface for the framework of your choice.

## Storage Performance

- IOPS: 500 sequential IOPS
- Write Performance: 10MB/s



### Note: etcd Storage requirements

Storage performance requirements are tied closely to the [etcd hardware recommendations \(https://github.com/etcd-io/etcd/blob/master/Documentation/op-guide/hardware.md#disks\)](https://github.com/etcd-io/etcd/blob/master/Documentation/op-guide/hardware.md#disks)

## 1.3 Network Requirements

- All the nodes on the cluster must be on a the same network and be able to communicate directly with one another.

### Important: Reliable Networking

Please make sure all nodes can communicate without interruptions.

- All nodes in the cluster must be assigned static IP addresses. Using dynamically assigned IPs will break cluster functionality after update/reboot.
- The admin node and the Kubernetes API master must have valid Fully-Qualified Domain Names (FQDNs), which can be resolved both by all other nodes and from other networks which need to access the cluster.  
Admin node and Kubernetes API master node should be configured as CNAME records in the local DNS. This improves portability for disaster recovery.
- A DNS server to resolve host names. If you are using host names to specify nodes, please make sure you have reliable DNS resolution at all times, especially in combination with DHCP.

### Important: Unique Host Names

Host names must be unique. It is recommended to let the DHCP server provide not only IP addresses but also host names of the cluster nodes.

- On the same network, a separate computer with a Web browser is required in order to complete bootstrap of the cluster.
- We recommend that SUSE CaaS Platform is setup to run in two subnets in one network segment, also referred to as VPC or VNET. The administration node should run in a subnet that is not accessible to the outside world and should be connected to your network via VPN or other means. Consider a security group/firewall that only allows ingress traffic on ports 22 (SSH) and 443 (https) for the Administrative node from outside the VPC. All nodes must have access to the Internet through some route in order to connect to SUSE Customer Center and receive updates, or be otherwise configured to receive updates, for example through SMT.



Depending on the applications running in your cluster you may consider exposing the subnet for the cluster nodes to the outside world. Use a security group/firewall that only allows incoming traffic on ports served by your workload. For example, a containerized application providing the backend for REST based services with content served over https should only allow ingress traffic on port 443.

### Important: Unique Host Names

All nodes' host names must be unique. If two or more nodes have the same host name, bootstrap of the cluster will fail.

### Note: Clusters without Fully-Qualified Domain Names

For test purposes, IP addresses can be substituted for the FQDNs for the administration node and master node, but this is not supported for production deployment.

## 1.4 Limitations

- SUSE CaaS Platform 3 does not support remote installations with Virtual Network Computing (VNC)..
- SUSE CaaS Platform is a dedicated cluster-node operating system. Dual-booting with other operating systems is not supported. It must be the only operating system installed on each node.

## 2 Installation

### ! Important: Scope of this document

This document is meant to give you a quick way to deploy a default installation of SUSE CaaS Platform. The instructions will assume the absolute minimum of configuration.

If you require any specific configuration for your environment *at all*, this document is not for you and you must consult the full deployment guide.

This chapter gives a quick overview on how to run through a default installation of SUSE CaaS Platform onto dedicated physical or virtual machines.

In addition, SUSE CaaS Platform supports other deployment methods and scenarios, such as from pre-installed disk images and onto OpenStack or certain supported public-cloud providers.

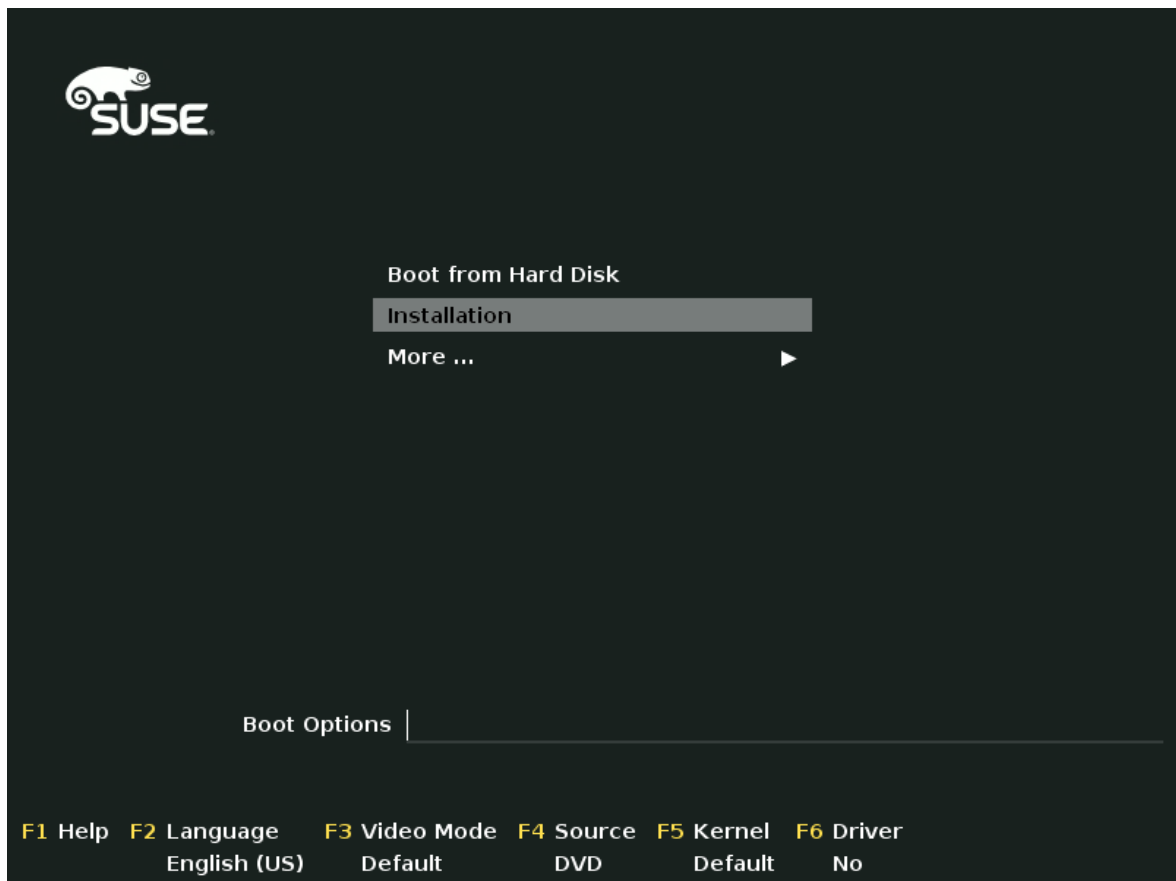
For these and other deployment scenarios, see the SUSE CaaS Platform Deployment Guide on <http://www.suse.com/documentation/>.

### 2.1 Installing The Administration Node

The following instructions assume that you are installing onto a dedicated physical or virtual machine, using either an ISO image or a physical removable medium, such as a DVD or bootable USB key. The hard disk should not contain any existing partitions.

1. Insert the bootable SUSE CaaS Platform installation medium, then reboot the computer to start the installation program. On machines with a traditional BIOS you will see the graphical boot screen shown below. On machines equipped with UEFI, a slightly different boot screen is used. Secure Boot on UEFI machines is supported.

Use **F2** to change the language for the installer. A corresponding keyboard layout is chosen automatically. For more information about changing boot options, see [https://www.suse.com/documentation/sles-12/book\\_sle\\_deployment/data/sec\\_i\\_yast2\\_s-tartup.html](https://www.suse.com/documentation/sles-12/book_sle_deployment/data/sec_i_yast2_s-tartup.html).



2. Select *Installation* on the boot screen, then press `Enter`. This boots the system and loads the SUSE CaaS Platform installer.
3. Configure the following mandatory settings on the *Installation Overview* screen.

### Tip: Help And Release Notes

From this point on, a brief help document and the Release Notes can be viewed from any screen during the installation process by selecting *Help* or *Release Notes* respectively.

#### Keyboard Layout

The *Keyboard Layout* is initialized with the language settings you have chosen on the boot screen. Change it here, if necessary.

#### Password for root User

Type a password for the system administrator account (called the root user) and confirm it.



## Warning: Do Not Forget The root Password

You should ensure that you will not lose the root password! After you entered it here, the password cannot be retrieved. See [https://www.suse.com/documentation/sles-12/book\\_sle\\_deployment/data/sec\\_i\\_yast2\\_user\\_root.html](https://www.suse.com/documentation/sles-12/book_sle_deployment/data/sec_i_yast2_user_root.html) for more information.

### Registration Code or SMT Server URL

Enter the *Registration Code or SMT Server URL*. SMT Server URLs must use https or http; other protocols are not supported.

### System Role

As this is the first node in the cluster, you must select "Administration Node (Dashboard)" as the *System Role* for your system.

### NTP Servers

Enter host names or IP addresses of one or more *NTP Servers* for the node, separated by white spaces or colons. While a single time server is sufficient, clusters should use at least three for optimal precision and reliability.

For more information about NTP, refer to [https://www.suse.com/documentation/sles-12/book\\_sle\\_admin/data/cha\\_netz\\_xntp.html](https://www.suse.com/documentation/sles-12/book_sle_admin/data/cha_netz_xntp.html)

Optionally, you can customize the following settings. If you do not make any changes, defaults are used. A brief summary of the settings is displayed below the respective settings option.

### Partitioning

Review the partition setup proposed by the system and change it if necessary. You have the following options:

#### Select a hard disk

Select a disk to which install SUSE CaaS Platform to with the recommended partitioning scheme.

#### *Custom Partitioning (for Experts)*

Opens the *Expert Partitioner* described in [https://www.suse.com/documentation/sles-12/book\\_sle\\_deployment/data/sec\\_yast2\\_i\\_y2\\_part\\_expert.html](https://www.suse.com/documentation/sles-12/book_sle_deployment/data/sec_yast2_i_y2_part_expert.html) .



## Warning: For Experts Only

As the name suggests, the *Expert Partitioner* is for experts only. Custom partitioning schemes that do not meet the requirements of SUSE CaaS Platform are not supported.

### REQUIREMENTS FOR CUSTOM PARTITIONING SCHEMES

- SUSE CaaS Platform exclusively supports the file system types Btrfs and OverlayFS. A read-only Btrfs file system is used for the root file system which enables transactional updates.
- For snapshots, partitions must have a capacity of at least 11 GB.
- Depending on the number and size of your containers, you will need sufficient space under the `/var` mount point.

To accept the proposed setup without any changes, choose *Next* to proceed.

### Booting

This section shows the boot loader configuration. Changing the defaults is only recommended if really needed. Refer to [https://www.suse.com/documentation/sles-12/book\\_sle\\_admin/data/cha\\_grub2.html](https://www.suse.com/documentation/sles-12/book_sle_admin/data/cha_grub2.html) for details.

### Network Configuration

If the network could not be configured automatically while starting the installation system, you have to manually configure the *Network Settings*. Please make sure at least one network interface is connected to the Internet in order to register your product.

By default, the installer requests a host name from a DHCP server. If you set a custom name in the *Hostname/DNS* tab, make sure that it is unique.

For more information on configuring network connections, refer to [https://www.suse.com/documentation/sles-12/book\\_sle\\_admin/data/sec\\_basic-net\\_yast.html](https://www.suse.com/documentation/sles-12/book_sle_admin/data/sec_basic-net_yast.html).

## Important: Reliable Networking

Please make sure all nodes are on the same network and can communicate without interruptions. If you are using host names to specify nodes, please make sure you have reliable DNS resolution at all times, especially in combination with DHCP.

### Kdump

Kdump saves the memory image (“core dump”) to the file system in case the kernel crashes. This enables you to find the cause of the crash by debugging the dump file. For more information, see [https://www.suse.com/documentation/sles-12/book\\_sle\\_tuning/data/cha\\_tuning\\_kdump\\_basic.html](https://www.suse.com/documentation/sles-12/book_sle_tuning/data/cha_tuning_kdump_basic.html).

## Warning: Kdump With Large Amounts Of RAM

If you have a system with large amounts of RAM or a small hard drive, core dumps may not be able to fit on the disk. If the installer warns you about this, there are two options:

1. Enter the *Expert Partitioner* and increase the size of the root partition so that it can accommodate the size of the core dump. In this case, you will need to decrease the size of the data partition accordingly. Remember to keep all other partitioning parameters (e.g. the root file system, the mount point of the data partition) when doing these changes.
2. Disable kdump completely.

### System Information

View detailed hardware information by clicking *System Information*. In this screen you can also change *Kernel Settings*. For more information, see [https://www.suse.com/documentation/sles-12/book\\_sle\\_tuning/data/cha\\_tuning\\_io.html](https://www.suse.com/documentation/sles-12/book_sle_tuning/data/cha_tuning_io.html).

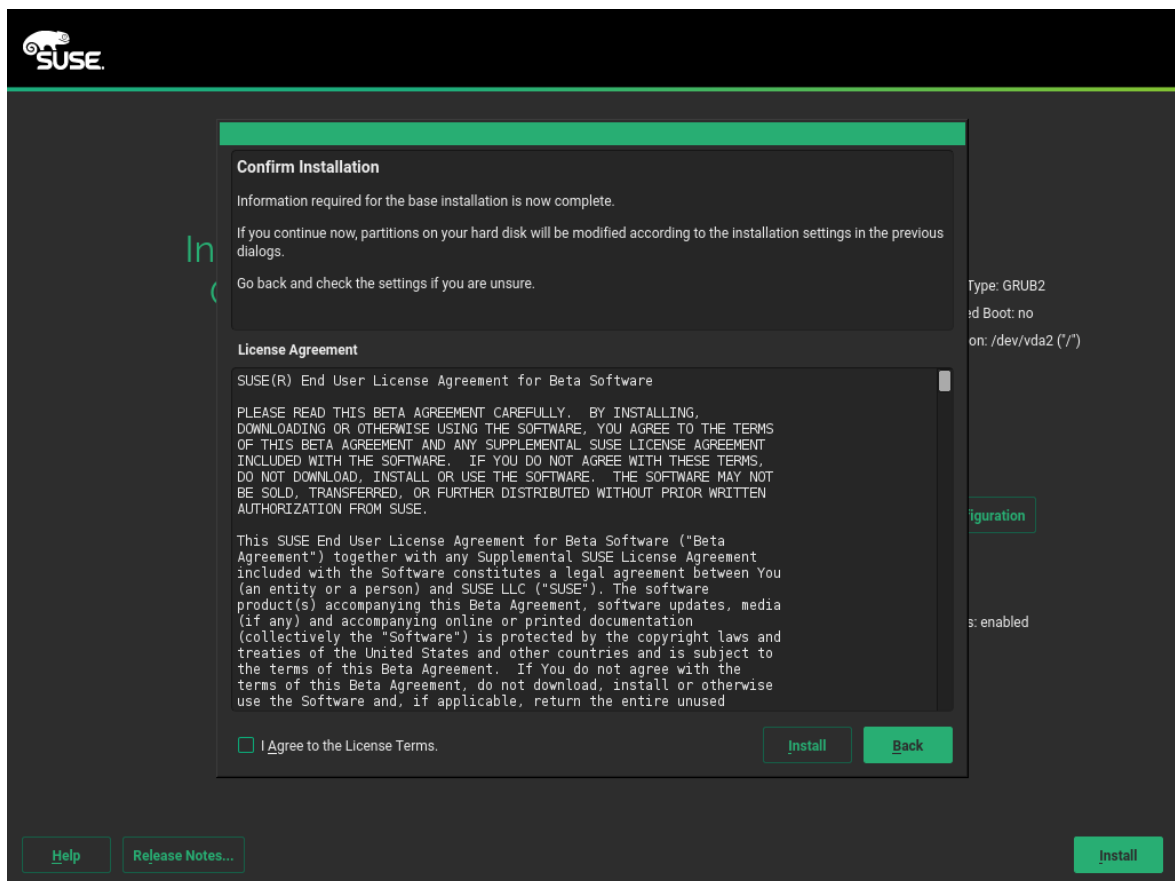
Proceed with *Next*.



## Tip: Installing Product Patches At Installation Time

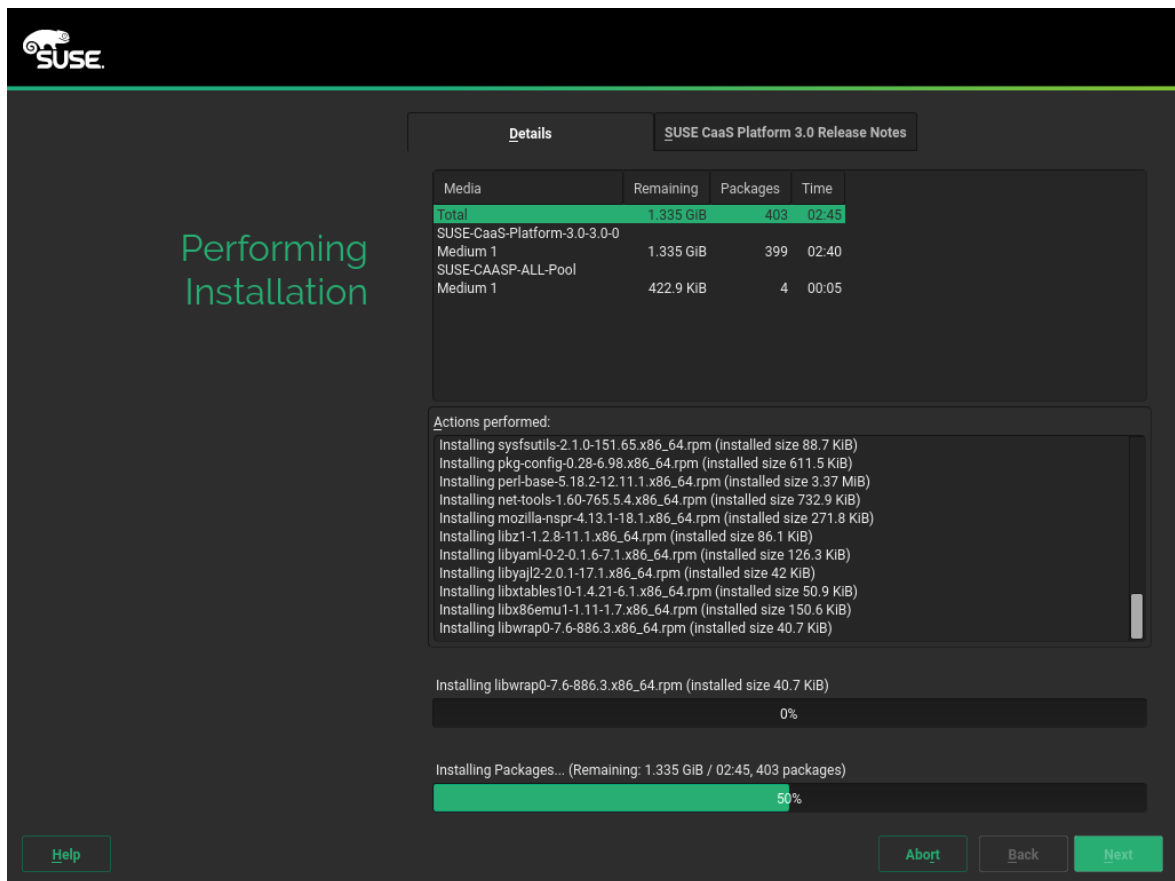
If SUSE CaaS Platform has been successfully registered at the SUSE Customer Center, you are asked whether to install the latest available online updates during the installation. If choosing *Yes*, the system will be installed with the most current packages without having to apply the updates after installation. Activating this option is recommended.

4. After you have finalized the system configuration on the *Installation Overview* screen, click *Install*. Up to this point no changes have been made to your system. After you click *Install* a second time, the installation process starts.



5. During the installation, the progress is shown in detail on the *Details* tab.



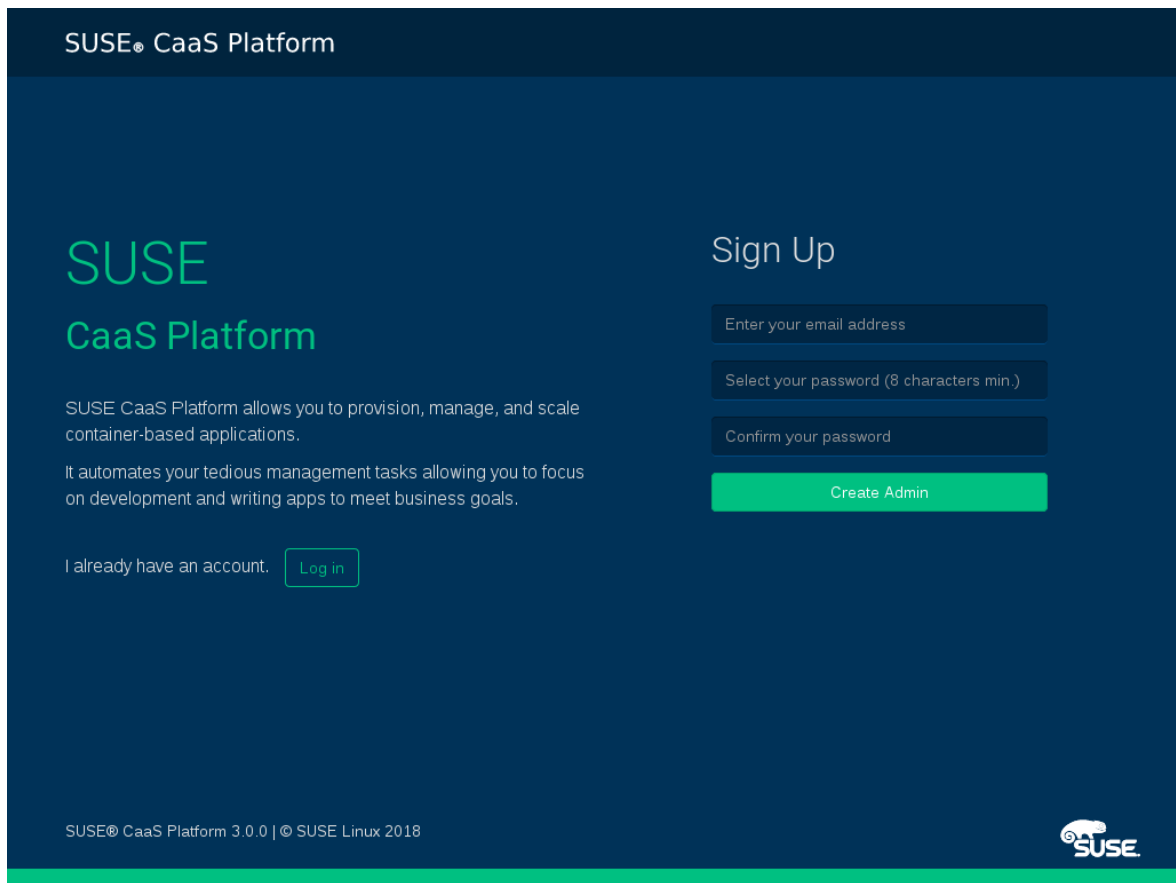


6. After the installation routine has finished, the computer is rebooted into the installed system.

## 2.2 Configuring the Administration Node

In order to finalize the configuration of the administration node, a few additional steps are required.

1. After the administration node has finished booting and you see the login prompt, point a web browser to <https://caasp-admin.example.com>, where `caasp-admin.example.com` is the host name or IP address of the administration node. The IP address is shown above the login prompt.
2. To create an Administrator account, click *Create an account* and provide an e-mail address and a password. Confirm the password and click *Create Admin*. You will be logged into the dashboard automatically.



3. Fill in the values for the minimum required options.

SUSE® CaaS Platform Logout

📘 Signed in successfully. ✕

## Initial CaaS Platform Configuration

Generic settings

**Internal Dashboard Location**

 📘

Cluster services

Install Tiller (Helm's server component)

Overlay network settings Show

Proxy settings Enable **Disable**

SUSE registry mirror Enable **Disable**

Cloud provider integration 📘 Enable **Disable**

Container runtime 📘

The choice of container runtime is completely transparent to end-users of the cluster. Neither Kubernetes manifests nor container images need to be changed.

**Choose the runtime**

**Docker open source engine**

*Docker open source engine (default) is a production-ready runtime, fully supported by SUSE.*

System wide certificate Show

**Next**

SUSE® CaaS Platform 3.0.0 | © SUSE Linux 2018

### Dashboard Location

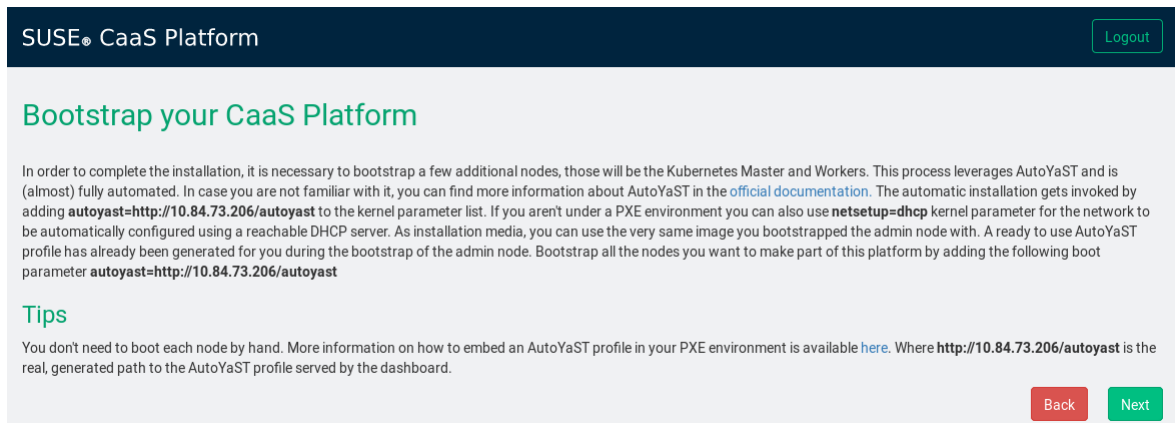
Host name or IP of the node running this web interface.

### Proxy Settings

If enabled, you can set proxy servers for HTTP and HTTPS. You may also configure exceptions and choose whether to apply the settings only to the container engine or to all processes running on the cluster nodes.

4. Click *Next*.

5. You will be shown an information screen about AutoYaST.



The screenshot shows a web interface for the SUSE CaaS Platform. At the top, there is a dark blue header with the text "SUSE® CaaS Platform" on the left and a "Logout" button on the right. Below the header, the main content area has a light gray background. The title "Bootstrap your CaaS Platform" is displayed in a green font. The main text explains that to complete the installation, additional nodes (Kubernetes Master and Workers) need to be bootstrapped using AutoYaST. It provides instructions on how to add the `autoyast=http://10.84.73.206/autoyast` kernel parameter to the boot process, either manually or via DHCP. A "Tips" section follows, advising that users do not need to boot each node by hand and providing a link for more information. At the bottom right of the content area, there are two buttons: "Back" (red) and "Next" (green).

This is now the time for you to install the master/worker nodes for the cluster.

## ! Important: Leave The Admin Node Configuration Open

If you quit out of the configuration process at this time by closing the browser window/tab, you will lose all input you made to the settings and have to restart the configuration from the beginning.

We recommend you leave the browser window/tab open while you perform the installation of the other nodes in a new window.

6. Click *Next* to proceed and install some cluster nodes as described in *Section 2.3, "Installing Cluster Nodes"*.

## 2.3 Installing Cluster Nodes

Cluster nodes can be installed manually, with AutoYaST, or are provisioned automatically on public cloud service. Manual installation is only feasible for a small number of workers. For larger numbers outside of public cloud service, AutoYaST is recommended.

You can start the setup via PXE. For the full procedure, refer to the SUSE Linux Enterprise 12 Deployment Guide: [https://www.suse.com/documentation/sles-12/singlehtml/book\\_sle\\_deployment/book\\_sle\\_deployment.html#cha.deployment.prep\\_boot](https://www.suse.com/documentation/sles-12/singlehtml/book_sle_deployment/book_sle_deployment.html#cha.deployment.prep_boot).

You can directly use the `initrd` and `linux` files from your install media, or install the package `tftpboot-installation-CAASP-3.0` on the TFTP server. The package provides the required `initrd` and `linux` files in the `/srv/tftpboot/` directory. You need to modify the paths used in the SUSE Linux Enterprise 12 Deployment Guide to correctly point to the files provided by the package.

Before you can set up a cluster node, you have to install and bootstrap an administration node to run the administration dashboard. Refer to [Section 2.1, "Installing The Administration Node"](#) for information on how to install the administration node.

### 2.3.1 Manual Installation

1. Follow the steps as described in [Section 2.1, "Installing The Administration Node"](#).
2. In step 3, select `Cluster Node` as *System Role* and enter the host name or IP address of the *Administration Node*.

#### Important: Reliable Networking

Please make sure all nodes are on the same network and can communicate without interruptions. If you are using host names to specify nodes, please make sure you have reliable DNS resolution at all times, especially in combination with DHCP.

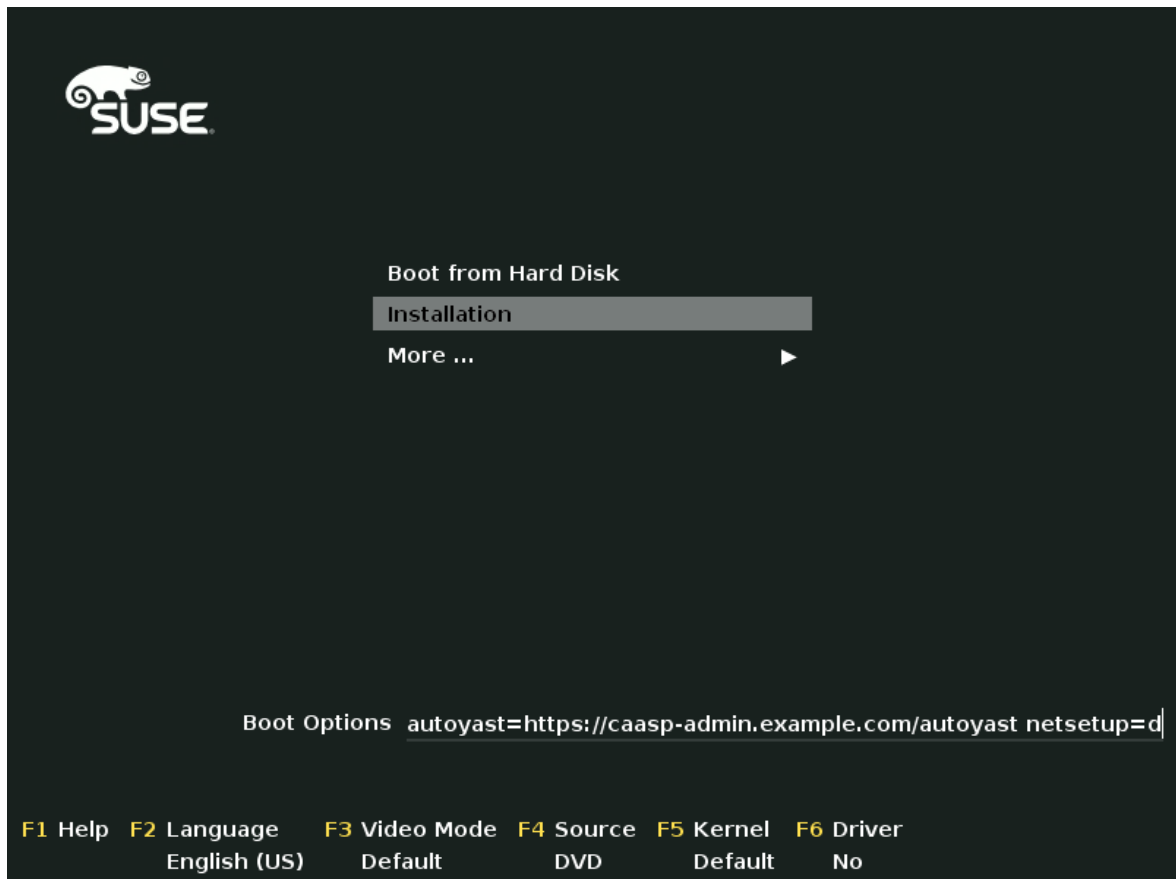
3. After you have finalized the system configuration on the *Installation Overview* screen, click *Install*. Up to this point no changes have been made to your system. After you click *Install* a second time, the installation process starts.  
After a reboot, the worker should show up in the dashboard and can be added to your cluster.

### 2.3.2 Installation With AutoYaST

Please make sure you have finished the initial configuration described in [Section 2.2, "Configuring the Administration Node"](#) before installing cluster nodes with AutoYaST.

1. Insert the SUSE CaaS Platform DVD into the drive, then reboot the computer to start the installation program.

2.



Select *Installation* on the boot screen. Enter the necessary *Boot Options* for AutoYaST and networking, namely:

#### autoyast

Path to the AutoYaST file. For more information, refer to [https://www.suse.com/documentation/sles-12/book\\_autoyast/data/invoking\\_autoinst.html#commandline\\_ay](https://www.suse.com/documentation/sles-12/book_autoyast/data/invoking_autoinst.html#commandline_ay)

#### ifcfg

Network configuration. If you are using dhcp, you can simply enter `ifcfg=eth0=dhcp`. Make sure to replace `eth0` with the actual name of the interface that you want to use DHCP for. For manual configuration, refer to [https://www.suse.com/documentation/sles-12/book\\_autoyast/data/ay\\_adv\\_network.html](https://www.suse.com/documentation/sles-12/book_autoyast/data/ay_adv_network.html).

#### hostname

The host name for the node, if not provided by DHCP. If you manually specify a host name, make sure it is unique.

Press `Enter`. This boots the system and loads the SUSE CaaS Platform installer.

3. The rest of the installation will run fully automated. After a reboot, the worker should show up in the dashboard and can be added to your cluster.

## 2.4 Bootstrapping The Cluster

To complete the installation of your SUSE CaaS Platform cluster, it is necessary to bootstrap at least three additional nodes; those will be the Kubernetes master and workers.

1. Return to your admin node; with the AutoYaST instructions screen open from before.
2. Click *Next*.
3. On the screen *Select nodes and roles*, you will see a list of `salt-minion` IDs under *Pending Nodes*. These are internal IDs for the master/worker nodes you have just set up and which have automatically registered with the admin node in the background.
4. *Accept* individual nodes into the cluster or click *Accept All Nodes*.
5. Assign the roles of the added nodes.  
By clicking on *Select remaining nodes*, all nodes without a selected role will be assigned the `Worker` role.

## Select nodes and roles

4 nodes found ▼ Select remaining nodes

Please choose an odd number of *Master* nodes, then click "Select remaining nodes" to assign the remaining nodes to the *Worker* role.

ID	Hostname	Role
31599b0cd7c84fe0aa5a9f752b3d10bb	worker-1	Master <b>Worker</b> Unused
761c3c7478d44d548472fe3b7106ffb9	master	<b>Master</b> Worker Unused
98720f66d7be4d94935b2076ae3e5b38	worker-0	Master <b>Worker</b> Unused
3fb0be4f193b4124b22ef2ffdb974b22	worker-2	Master <b>Worker</b> Unused

Back Next

---

Pending Nodes ▼ Accept All Nodes

You currently have no nodes to be accepted for bootstrapping.

## ! Important: Minimum Cluster Size

You must designate at least 1 master node and 2 worker nodes..

## 💡 Tip: Assign Unused Nodes Later

Nodes that you do not wish to designate for a role now, can later be assigned one on the Velum status page.

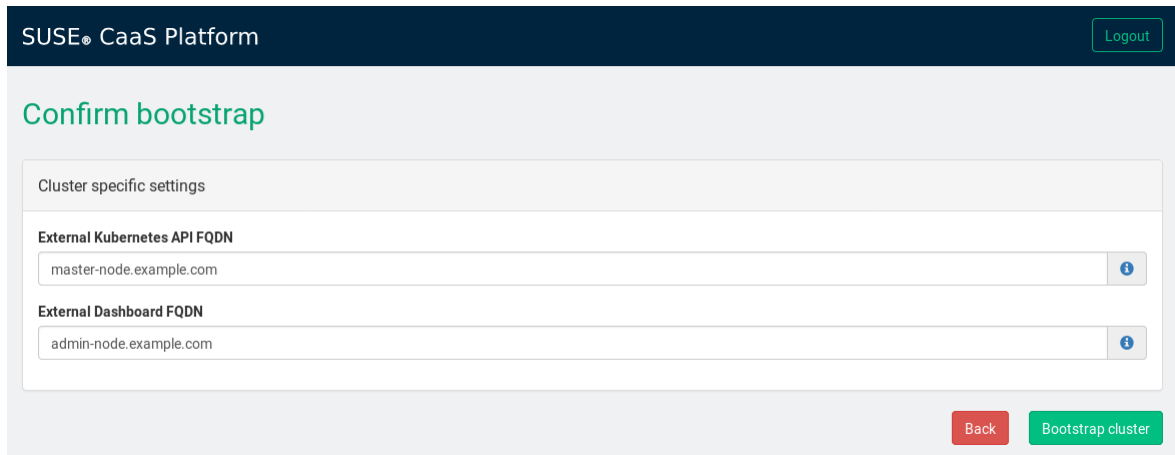
- Once you have assigned all desired nodes a role, click *Next*.
- The last step is to configure the external FQDNs for dashboard and Kubernetes API. These values will determine where the nodes in the cluster will attempt to communicate.

## 📝 Note: Master Node Loadbalancer FQDN

If you are planning a larger cluster with multiple master nodes, they must all be accessible from a single host name. If not, the functionality of Velum will degrade if the original master node is removed.



Therefore, you should ensure that there is some form of load-balancing or reverse proxy configured at the location you enter here.



SUSE® CaaS Platform Logout

## Confirm bootstrap

Cluster specific settings

**External Kubernetes API FQDN**  
master-node.example.com ⓘ

**External Dashboard FQDN**  
admin-node.example.com ⓘ

Back Bootstrap cluster

### External Kubernetes API FQDN

Name used to reach the node running the Kubernetes API server.

In a simple deployment with a single master node, this will be the name of the node that was selected as the master node during bootstrapping of the cluster.

### External Dashboard FQDN

Name used to reach the admin node running Velum.

8. Click on *Bootstrap cluster* to finalize the initial setup and start the bootstrapping process. The status overview will be shown while the nodes are bootstrapped for their respective roles in the background.

## 3 Configuration

SUSE CaaS Platform is typically configured in two stages: first, during the installation process, and subsequently, using `cloud-init` and/or Velum settings page. During the initial installation, SUSE CaaS Platform comes as preconfigured as possible.

### 3.1 Deployment Scenarios Requiring Second-stage Configuration

Second-stage configuration is typically used for large-scale clusters built using pre-installed virtual machine disk images. Because these images are pre-installed, the installation program does not run, and so it is not possible to choose options interactively. Therefore, they must be specified in configuration files.

For smaller clusters of just a few nodes, `cloud-init` is generally not necessary, and so it is not described here in the *Installation Quick Start Guide*. For information on using `cloud-init`, refer to the Configuration chapter in the *Deployment Guide*.

The following section describes the default values for the system configuration. If you require different settings than these, follow the *Deployment Guide* to build your cluster.

### 3.2 Defaults

The following are the defaults for the first stage of configuration:

#### timezone

This is set to *UTC* by default. It is not recommended to change this, but it is possible using `cloud-init`; see the *Deployment Guide*.

#### keyboard

This is set to *us* by default. If desired, choose an alternative layout from the list on the first screen of the installation program.

#### locale

This is set to *en\_US.utf8* by default, but can be changed by `cloud-init`.