



Deploy K3s on SUSE Linux Enterprise Micro and Cisco UCS C220, C240, and C240 SD

Contents

Purpose of this document	3
Introduction	3
Solution overview	5
Prerequisites	5
Configure Cisco UCS C240 SD through the Cisco IMC.....	5
Install SLE Micro.....	15
Install K3s.....	34
K3s integration into the workload management tool	38
Conclusion	46
For more information.....	46

Purpose of this document

This document provides a high-level procedure for deploying the K3s lightweight Kubernetes distribution on Cisco UCS® C220, C240, and C240 SD Rack Servers in space-constrained locations. The focus will be on areas where the deployment deviates from default installations. Everything that is not specified in this document can be configured based on the default settings for your local environment.

Introduction

During the past few years, organizations have been participating in a radical transformation of the way that modern applications are built, deployed, and operated. Monolithic applications are being broken down into microservices and serverless functions to ease development exponentially, facilitate lifecycle management, increase the speed at which new features are deployed, and improve the availability of services offered.

More and more mission-critical workloads have become containerized. According to various Gartner and IDC estimates, between 35 and 50 percent of an enterprise's application sprawl is now containerized—and not just the application front ends or the dashboards, but mission-critical workloads such as revenue-generating data analytics pipelines, middleware, and core business logic.

Not only are workloads and applications changing, but the locations at which data is generated, accessed, and partially processed are changing from the data center to a highly distributed world. Hybrid cloud, edge, the Internet of Things (IoT), and similar technologies are becoming the default for more and more companies, and IT departments must find ways to deploy, manage, and support containerized workloads at nearly every place: in the data center, at the shop floor, in vehicles, and in the public cloud.

This document provides a sample configuration for deploying a container platform on a single server that provides all the capabilities of the data center while fitting into a shortened network rack at the shop floor or edge location: the Cisco UCS C240 SD Rack Server. For the operating system, the solution uses SUSE Enterprise Linux (SLE) Micro: an optimized container option based on the proven enterprise-class Linux distribution. The lightweight Kubernetes service K3s, which is optimized to run on a single server, eliminates the need to install multiple servers.

About Cisco Unified Computing System

The solution uses a Cisco UCS C240 M5SX Rack Server with solid-state disks (SSDs) and hard-disk drives (HDDs). The configuration can be used with any Cisco UCS C-Series Rack Server.

Cisco UCS C240 M5 Rack Server overview

The Cisco UCS C240 M5 Rack Server is an enterprise-class server in a 2-rack-unit (2RU) form factor. It is designed to deliver exceptional performance, expandability, and efficiency for storage and I/O-intensive infrastructure workloads. These workloads include big data analytics, virtualization, and graphics-intensive and bare-metal applications.

The Cisco UCS C240 M5 server provides:

- Support for a 2RU 2-socket server using Intel® Xeon® Scalable processors
- Support for 2666-MHz DDR4 DIMMs and 128-GB DIMMs
- Increased storage density with 24 front-pluggable 2.5-inch small-form-factor (SFF) drive bays, or 12 front-pluggable 3.5-inch large-form-factor (LFF) drive bays and 2 rear 2.5-inch SFF drive bays

-
- Non-Volatile Memory Express (NVMe) PCI Express (PCIe) SSD support (for up to 2 drives on the standard chassis SKU or up to 10 drives on the NVMe-optimized SKU)
 - Cisco® 12-Gbps SAS RAID modular controller and Cisco 12-Gbps SAS host bus adapter (HBA) controller
 - 2 Flexible Flash (FlexFlash) Secure Digital (SD) card slots or 2 modular M.2 SATA slots
 - 10-Gbps embedded Intel x550 10GBASE-T LAN-on-motherboard (LOM) port
 - 1 modular LOM (mLOM) slot
 - 6 PCIe Generation 3 (Gen 3) slots
 - Up to 2 hot-pluggable redundant power supplies

The Cisco UCS C240 M5 server can be deployed as a standalone device or as part of a managed Cisco Unified Computing System™ (Cisco UCS) environment. Cisco UCS unifies computing, networking, management, virtualization, and storage access into a single integrated architecture that can enable end-to-end server visibility, management, and control in both bare-metal and virtualized environments. With a Cisco UCS managed deployment, the Cisco UCS C240 M5 takes advantage of our standards-based unified computing innovations to significantly reduce customers' total cost of ownership (TCO) and increase business agility.

About SUSE Linux Enterprise Micro

SUSE Linux Enterprise, or SLE, Micro is an ultra-reliable, lightweight operating system purpose-built for containerized and virtualized workloads. It uses the enterprise-hardened security and compliance components of SUSE Linux Enterprise and merges them with a modern, immutable, developer-friendly OS platform.

About K3s lightweight Kubernetes

K3s is packaged as a single binary about 50 MB in size. Bundled in that single binary is everything needed to run Kubernetes anywhere, including low-powered IoT and edge-based devices. The binary includes:

- The container runtime
- Any essential host utilities, such as iptables, socat and, du

The only OS dependencies are the Linux kernel itself and proper dev, proc, and sysfs mounts (these are included automatically in all modern Linux distributions).

K3s bundles these Kubernetes components:

- kube-apiserver
- kube-controller-manager
- kube-scheduler
- kubelet
- kube-proxy

Solution overview

The Cisco Integrated Management Console (IMC) is the basic option for managing Cisco UCS C-Series Rack Servers over the network and installing the operating system using the virtual console. SUSE Rancher is used to manage the Kubernetes installations in the data center, in branch offices, at the edge, and in the public cloud (Figure 1).

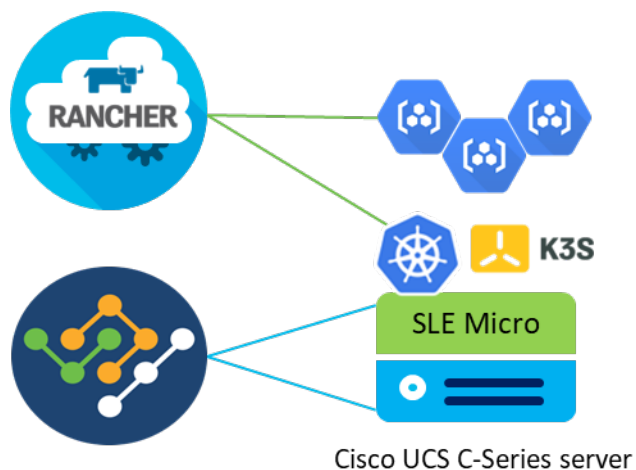


Figure 1.
Solution overview with Cisco IMC and SUSE Rancher

Prerequisites

The following items need to be preconfigured before you begin the setup and configuration of a K3s system on a Cisco UCS C240 SD server:

- Linux host with kubectl client binary installed and access to the Internet to download required software packages
- One Cisco UCS C240 SD racked and cabled
- Domain Host Configuration Protocol (DHCP) server to provide an IP address to the Cisco IMC
- Monitor, keyboard, and mouse for initial IMC configuration

Configure Cisco UCS C240 SD through the Cisco IMC

Use the procedure described in this section to prepare the Cisco UCS C-Series server for the SLE Micro installation. The main focus here is the configuration of the storage and the network.

The configuration steps presented here are for Cisco UCS servers that are not connected to a pair of Cisco UCS fabric interconnects (Cisco UCS managed mode) or to the Cisco Intersight™ platform. All configuration steps are performed on the local Cisco IMC.

Perform initial setup

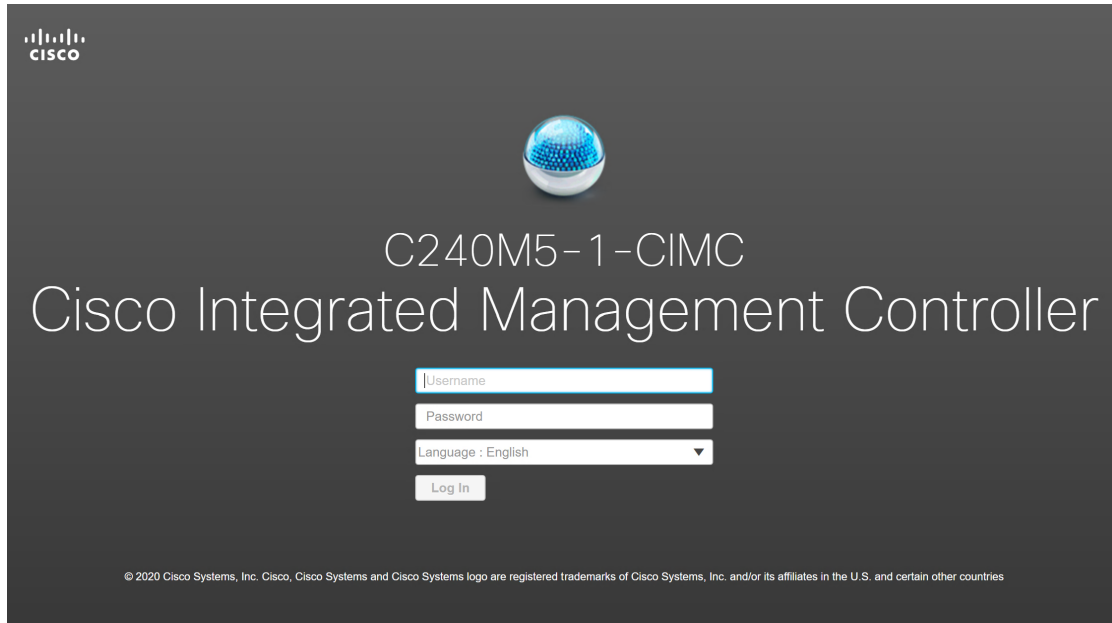
Hardware installation details and the initial server setup process are documented in the used server's installation documentation. For Cisco the UCS C240 SD M5, the document can be found here:

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/c240sdm5/install/c240sdm5/C240M5_chapter_01.html.

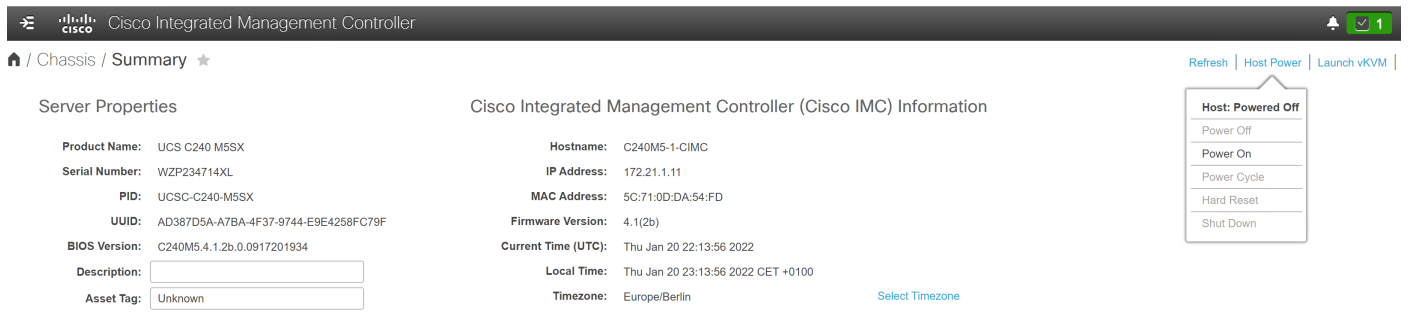
Configure the storage and network

Follow these steps to configure the storage and network for SLE Micro and K3s:

1. Open the IMC page in a web browser.



2. Enter the username and password and click Log In.
3. Some of the configuration steps require the server to be powered on. At the top right side of the window, click Host Power > Power On and then click OK in the pop-up window. Wait 30 seconds before proceeding to the next step.

The image shows the "Summary" page of the Cisco IMC. The page title is "Cisco Integrated Management Controller (Cisco IMC) Information". On the left, there are "Server Properties" including Product Name (UCS C240 M5SX), Serial Number (WZP234714XL), PID (UCSC-C240-M5SX), UUID (AD387D5A-A7BA-4F37-9744-E9E4258FC79F), BIOS Version (C240M5.4.1.2b.0.0917201934), Description (empty field), and Asset Tag (Unknown). On the right, there are "Cisco Integrated Management Controller (Cisco IMC) Information" including Hostname (C240M5-1-CIMC), IP Address (172.21.1.11), MAC Address (5C:71:0D:DA:54:FD), Firmware Version (4.1(2b)), Current Time (UTC) (Thu Jan 20 22:13:56 2022), Local Time (Thu Jan 20 23:13:56 2022 CET +0100), and Timezone (Europe/Berlin). At the top right, there are navigation links: "Refresh", "Host Power", and "Launch vKVM". A pop-up window titled "Host: Powered Off" is visible, showing options: "Power Off", "Power On", "Power Cycle", "Hard Reset", and "Shut Down".

4. In the left menu, click Storage > Cisco 12G Modular Raid Controller.

Server Properties

Product Name: UCS C240 M5SX
 Serial Number: WZP234714XL
 PID: UCSC-C240-M5SX
 UUID: AD387D5A-A7BA-4F37-9744-E9E4258FC79F
 BIOS Version: C240M5.4.1.2b.0.0917201934
 Description:
 Asset Tag:

Cisco Integrated Management Controller (Cisco IMC) Information

Hostname: C240M5-1-CIMC
 IP Address: 172.21.1.11
 MAC Address: 5C:71:0D:DA:54:FD
 Firmware Version: 4.1(2b)
 Current Time (UTC): Thu Jan 20 13:00:33 2022
 Local Time: Thu Jan 20 14:00:33 2022 CET +0100
 Timezone: Europe/Berlin [Select Timezone](#)

Chassis Status

- Power State: ● On
- Overall Server Status: ✔ Good
- Temperature: ✔ Good
- Overall DIMM Status: ✔ Good
- Power Supplies: ✔ Good
- Fans: ✔ Good
- Locator LED: ● Off
- Overall Storage Status: ✔ Good

Server Utilization

(%)

100
90
80
70
60
50
40
30
20
10
0

Server

- Overall Utilization (%)
- CPU Utilization (%)
- Memory Utilization (%)
- IO Utilization (%)

5. Click Create Virtual Drive from Unused Physical Drives.

Cisco Integrated Management Controller

Home / ... / Cisco 12G Modular Raid Controller with 4GB cache (max 26 drives) / **Controller Info** ★

Controller Info | Physical Drive Info | **Virtual Drive Info** | Battery Backup Unit | Storage Log

[Create Virtual Drive from Unused Physical Drives](#) | [Create Virtual Drive from an Existing Virtual Drive Group](#) | [Import Foreign Config](#) | [Clear Foreign Config](#) | [Clear Boot Drive](#) | [Get Storage Firmware Log](#) | [Enable Drive Security](#) | [Disable Drive Security](#) | [Clear Cache](#) | [Clear all Configuration](#) | [Set Factory Defaults](#) | [Switch to Remote Key Management](#) | [Switch to Local Key Management](#)

- Select RAID Level 1, pick two SSD drives for the operating system, and click the >> button.
- Under Virtual Drive Properties, enter **boot** as the name and then click Create Virtual Drive.

RAID Level: 1 Enable Full Disk Encryption:

Create Drive Groups

Physical Drives selected 0 / total 14

ID	Size(MB)	Model	Interface	Type
<input type="checkbox"/>	1	952720 MB	SEAGA...	HDD SAS
<input type="checkbox"/>	2	952720 MB	SEAGA...	HDD SAS
<input type="checkbox"/>	3	952720 MB	SEAGA...	HDD SAS
<input type="checkbox"/>	4	952720 MB	SEAGA...	HDD SAS
<input type="checkbox"/>	5	952720 MB	SEAGA...	HDD SAS

Drive Groups

Name
<input type="checkbox"/> DG [23,24]

Virtual Drive Properties

Name:

Access Policy:

Read Policy:

Cache Policy:

Disk Cache Policy:

Write Policy:

Strip Size (MB):

Size: MB

- Click Virtual Drive Info, select the boot virtual drive in the list, and click Set as Boot Drive.

Controller Info Physical Drive Info **Virtual Drive Info** Battery Backup Unit Storage Log

Virtual Drives

VD-0

Virtual Drive Number	Name	Status	Health	Size	RAID Level	Boot Drive
<input checked="" type="checkbox"/> 0	boot	Optimal	Good	962109 MB	RAID 1	false

- In the pop-up window, click OK to confirm the selection.

Are you sure you want to make the virtual drive as boot device- 0?

10. In the left menu, click Networking > Adapter Card MLOM.

The screenshot shows the Cisco Integrated Management Controller (IMC) interface. The breadcrumb path is: / ... / Cisco 12G Modular Raid Controller with 4GB cache (max 26 drives) (MRAID) / Controller Info. The left navigation menu is expanded to 'Networking' > 'Adapter Card MLOM'. The main content area has tabs for 'Controller Info', 'Physical Drive Info', 'Virtual Drive Info', 'Battery Backup Unit', and 'Storage Log'. Below the tabs are links for creating virtual drives, clearing boot drives, and enabling drive security. The 'Health/Status' section shows: Composite Health: Good, Controller Status: Optimal, RAID Chip Temperature: 55, and Storage Firmware Log Status: Not Downloaded. The 'Settings' section shows: Predictive Fail Poll Interval: 300 sec, Rebuild Rate: 30%, Patrol Read Rate: 30%, Consistency Check Rate: 30%, and Reconstruction Rate: 30%.

11. Click the vNICs tab.

The screenshot shows the Cisco IMC Controller interface for the vNICs configuration page. The breadcrumb path is: / ... / Adapter Card MLOM / General. The left navigation menu is expanded to 'Networking' > 'Adapter Card MLOM' > 'vNICs'. The main content area has tabs for 'General', 'External Ethernet Interfaces', 'vNICs', and 'vHBAs'. Below the tabs are links for 'Export vNIC', 'Import vNIC', 'Reset', and 'Reset To Defaults'. The 'Adapter Card Properties' section shows: PCI-Slot: MLOM, Vendor: Cisco Systems Inc, Product Name: UCS VIC 1457, Product ID: UCSC-MLOM-C25Q-04, Serial Number: FCH240274VQ, Version ID: V05, Hardware Revision: 5, Cisco IMC Management Enabled: no, and Configuration Pending: no. The 'Firmware' section shows: Running Version: 5.1(2e), Backup Version: 5.1(1.38), Startup Version: 5.1(2e), Bootloader Version: 5.0(3e), and Status: Fwupdate never issued. The 'Settings' section shows: ICSI Boot Capable: True, CDN Capable: True, usNIC Capable: True, Port Channel Capable: True, and a Description field. The 'Enable FIP Mode', 'Enable LLDP', and 'Port Channel' checkboxes are checked, while 'Enable VNTAG Mode' is unchecked.

- The factory default configuration comes with two virtual network interface cards (vNICs) defined: one assigned to port 0, and one assigned to port 1. Both vNICs are configured to allow any kind of traffic, with or without a VLAN tag. VLAN IDs must be managed at the operating system level.

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input type="checkbox"/> eth0	VIC-MLO...	3C:57:31:28:BF:5A	1500	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth1	VIC-MLO...	3C:57:31:28:BF:5B	1500	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A

A useful feature of the Cisco virtual interface card (VIC) is the capability to define multiple virtual network adapters to be presented to the operating system, with each configured for specific uses. For example, you can configure administration traffic with a maximum transmission unit (MTU) of 1500 to be compatible with all communication partners, and you can configure the network for storage traffic with MTU 9000 for the best throughput. This sample configuration uses this approach, creating two vNICs for administration traffic, two vNICs for default user traffic, and two vNICs for data traffic to the storage location. For high availability, the two network devices per traffic type will be combined in a bond on the operating system layer.

- Click the first vNIC in the list on the left side and change the default VLAN to 211. Click Save Changes.

vNIC Properties

General

Name: eth0

CDN: VIC-MLOM-eth0

MTU: 1500 (1500 - 9000)

Uplink Port: 0

MAC Address: Auto 3C:57:31:28:BF:5A (0 - 6)

Trust Host CoS:

PCI Order: 0 (0 - 3)

Default VLAN: 211 (0 - 4095)

VLAN Mode: Trunk

Rate Limit: OFF

Channel Number: N/A (1 - 1000)

PCI Link: 0 (0 - 1)

Enable NVGRE:

Enable VXLAN:

Geneve Offload:

Advanced Filter:

Port Profile: N/A

Enable PXE Boot:

Enable VMQ:

Enable Multi Queue:

No. of Sub vNICs: 64 (1 - 64)

Enable aRFS:

Enable Uplink Failover:

Failback Timeout: N/A (0 - 600)

► Ethernet Interrupt

► Ethernet Receive Queue

Save Changes Reset Values

14. Click the second vNIC in the list and change the default VLAN to 211 as well. Click Save Changes.

15. Click vNICs in the left menu and select the first vNIC in the table in the right pane. Click Clone vNIC.

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input checked="" type="checkbox"/>	eth0	VIC-MLO...	3C:57:31:28:BF:5A	1500	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth1	VIC-MLO...	3C:57:31:28:BF:5B	1500	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A

16. In the pop-up window, enter a name. Here, we used eth2. Change the default VLAN to 210. Click Add vNIC. Then confirm the creation of the vNIC by clicking OK in the next pop-up window.

17. Select the second vNIC in the table and click Clone vNIC.

Cisco Integrated Management Controller

Adapter Card MLOM / vNICs

General External Ethernet Interfaces **vNICs** vHBAs

vNICs

eth0
eth1
eth2

Host Ethernet Interfaces

Add vNIC Clone vNIC Delete vNICs

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input type="checkbox"/>	eth0	VIC-MLO...	3C:57:31:28:BF:5A	1500	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input checked="" type="checkbox"/>	eth1	VIC-MLO...	3C:57:31:28:BF:5B	1500	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth2	VIC-MLO...	3C:57:31:28:BF:5E	1500	0	0	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A

18. In the pop-up window, enter a name. Here, we used eth3. Change the default VLAN to 210. Click Add vNIC. Then confirm the creation of the vNIC by clicking OK in the next pop-up window.

Add vNIC

General

Name: eth3

CDN: []

MTU: 1500 (1500 - 9000)

Uplink Port: 1

MAC Address: Auto

Class of Service: 0 (0 - 6)

Trust Host CoS: []

PCI Order: 5 (0 - 5)

Default VLAN: None

VLAN Mode: Trunk

Rate Limit: OFF

Channel Number: [] (1 - 1000)

PCI Link: 0 (0 - 1)

Enable NVGRE: []

Enable VXLAN: []

Geneve Offload: []

Advanced Filter: []

Port Profile: []

Enable PXE Boot: []

Enable VMQ: []

Enable Multi Queue: []

No. of Sub vNICs: 64 (1 - 64)

Enable aRFS: []

Enable Uplink Failover: []

Failback Timeout: N/A (0 - 600)

Ethernet Interrupt

Ethernet Receive Queue

Ethernet Transmit Queue

Add vNIC Reset Values Close

19. Select the first vNIC in the table in the right pane. Click Clone vNIC.

Cisco Integrated Management Controller

Adapter Card MLOM / vNICs

General External Ethernet Interfaces **vNICs** vHBAs

vNICs

eth0
eth1
eth2
eth3

Host Ethernet Interfaces

Add vNIC Clone vNIC Delete vNICs

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input checked="" type="checkbox"/>	eth0	VIC-MLO...	3C:57:31:28:BF:5A	1500	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth1	VIC-MLO...	3C:57:31:28:BF:5B	1500	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth2	VIC-MLO...	3C:57:31:28:BF:5E	1500	0	0	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth3	VIC-MLO...	3C:57:31:28:BF:5F	1500	0	1	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A

20. In the pop-up window, enter a name. Here, we used **eth4**. Change the MTU to 9000 and change the default VLAN to 212. Click Add vNIC. Then confirm the creation of the vNIC by clicking OK in the next pop-up window.

Add vNIC

General

Name: eth4

CDN: []

MTU: 9000 (1500 - 9000)

Uplink Port: 0

MAC Address: Auto

Class of Service: 0 (0 - 6)

Trust Host CoS:

PCI Order: 6 (0 - 6)

Default VLAN: None 212

VLAN Mode: Trunk

Rate Limit: OFF (1 - 25000)

Channel Number: [] (1 - 1000)

PCI Link: 0 (0 - 1)

Enable NVGRE:

Enable VXLAN:

Geneve Offload:

Advanced Filter:

Port Profile: []

Enable PXE Boot:

Enable VMQ:

Enable Multi Queue:

No. of Sub vNICs: 64 (1 - 64)

Enable aRFS:

Enable Uplink Failover:

Failback Timeout: N/A (0 - 600)

▶ Ethernet Interrupt

▶ Ethernet Receive Queue

▶ Ethernet Transmit Queue

Add vNIC **Reset Values** **Close**

21. Select the second vNIC in the table and click Clone vNIC.

Cisco Integrated Management Controller

Home / ... / Adapter Card MLOM / vNICs

General External Ethernet Interfaces **vNICs** vHBAs

vNICs

eth0

eth1

eth2

eth3

eth4

Host Ethernet Interfaces

Add vNIC Clone vNIC Delete vNICs

	Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input type="checkbox"/>	eth0	VIC-MLO...	3C:57:31:28:BF:5A	1500	0	0	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input checked="" type="checkbox"/>	eth1	VIC-MLO...	3C:57:31:28:BF:5B	1500	0	1	0	NONE	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth2	VIC-MLO...	3C:57:31:28:BF:5E	1500	0	0	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth3	VIC-MLO...	3C:57:31:28:BF:5F	1500	0	1	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/>	eth4	VIC-MLO...	3C:57:31:28:BF:60	9000	0	0	0	212	TRUNK	disabled	disabled	N/A	N/A	N/A

22. In the pop-up window, enter a name. Here, we used **eth5**. Change the MTU to 9000 and change the default VLAN to 212. Click Add vNIC. Then confirm the creation of the vNIC by clicking OK in the next pop-up window.

The screenshot shows the 'Add vNIC' configuration window with the following settings:

- Name: eth5
- CDN: (empty)
- MTU: 9000 (range: 1500 - 9000)
- Uplink Port: 1
- MAC Address: Auto
- Class of Service: 0 (range: 0 - 6)
- Trust Host CoS: (unchecked)
- PCI Order: 7 (range: 0 - 7)
- Default VLAN: 212
- VLAN Mode: Trunk
- Rate Limit: OFF
- Channel Number: (empty) (range: 1 - 1000)
- PCI Link: 0 (range: 0 - 1)
- Enable NVGRE: (unchecked)
- Enable VXLAN: (unchecked)
- Geneve Offload: (unchecked)
- Advanced Filter: (unchecked)
- Port Profile: (empty)
- Enable PXE Boot: (unchecked)
- Enable VMQ: (unchecked)
- Enable Multi Queue: (unchecked)
- No. of Sub vNICs: 64 (range: 1 - 64)
- Enable aRFS: (unchecked)
- Enable Uplink Failover: (unchecked)
- Fallback Timeout: N/A (range: 0 - 600)

Buttons at the bottom right: Add vNIC (blue), Reset Values, Close.

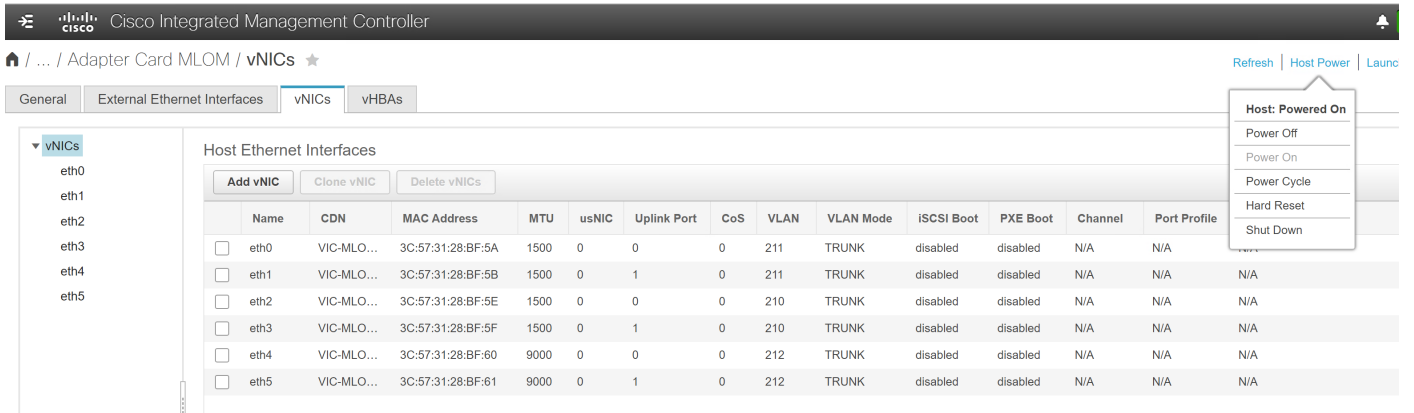
23. The vNICs tab shows the final list of vNICs. For every VLAN ID, there are two vNICs: one on uplink port 0, and one on uplink port 1.

You will use this list later to validate the bond configuration at the operating system layer.


The screenshot shows the 'vNICs' configuration page in the Cisco IMC. The 'vNICs' tab is selected, and the 'Host Ethernet Interfaces' table is displayed. The table contains the following data:

Name	CDN	MAC Address	MTU	usNIC	Uplink Port	CoS	VLAN	VLAN Mode	iSCSI Boot	PXE Boot	Channel	Port Profile	Uplink Failover
<input type="checkbox"/> eth0	VIC-MLOM-eth0	3C:57:31:28:BF:5A	1500	0	0	0	211	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth1	VIC-MLOM-eth1	3C:57:31:28:BF:5B	1500	0	1	0	211	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth2	VIC-MLOM-eth2	3C:57:31:28:BF:5E	1500	0	0	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth3	VIC-MLOM-eth3	3C:57:31:28:BF:5F	1500	0	1	0	210	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth4	VIC-MLOM-eth4	3C:57:31:28:BF:60	9000	0	0	0	212	TRUNK	disabled	disabled	N/A	N/A	N/A
<input type="checkbox"/> eth5	VIC-MLOM-eth5	3C:57:31:28:BF:61	9000	0	1	0	212	TRUNK	disabled	disabled	N/A	N/A	N/A

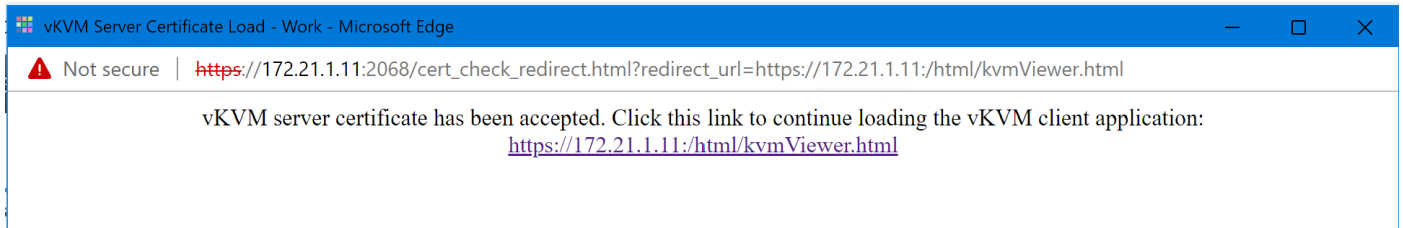
24. The new settings become active with the next power cycle of the server. At the top right side of the window, click Host Power > Power Off. In the pop-up window, click OK.



25. Click Launch vKVM.



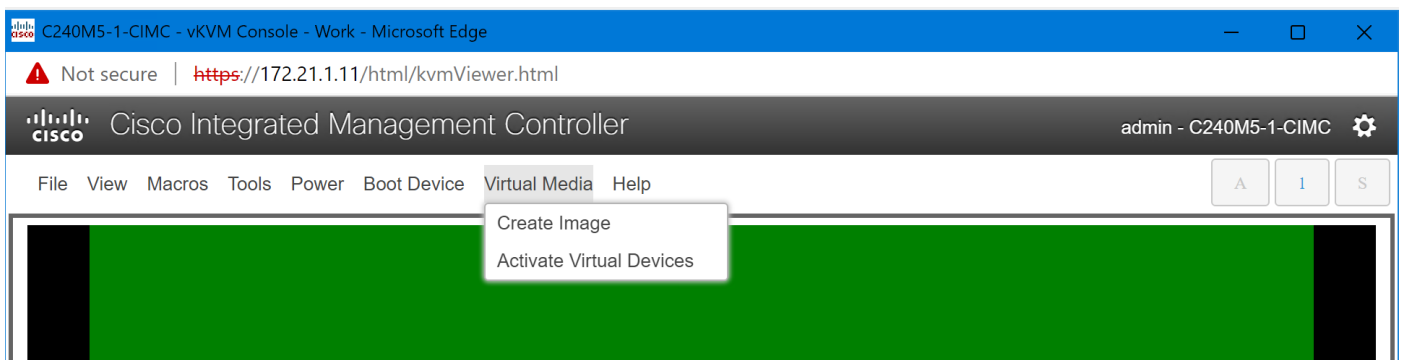
26. In the new window, take the necessary steps to continue with an untrusted certificate.



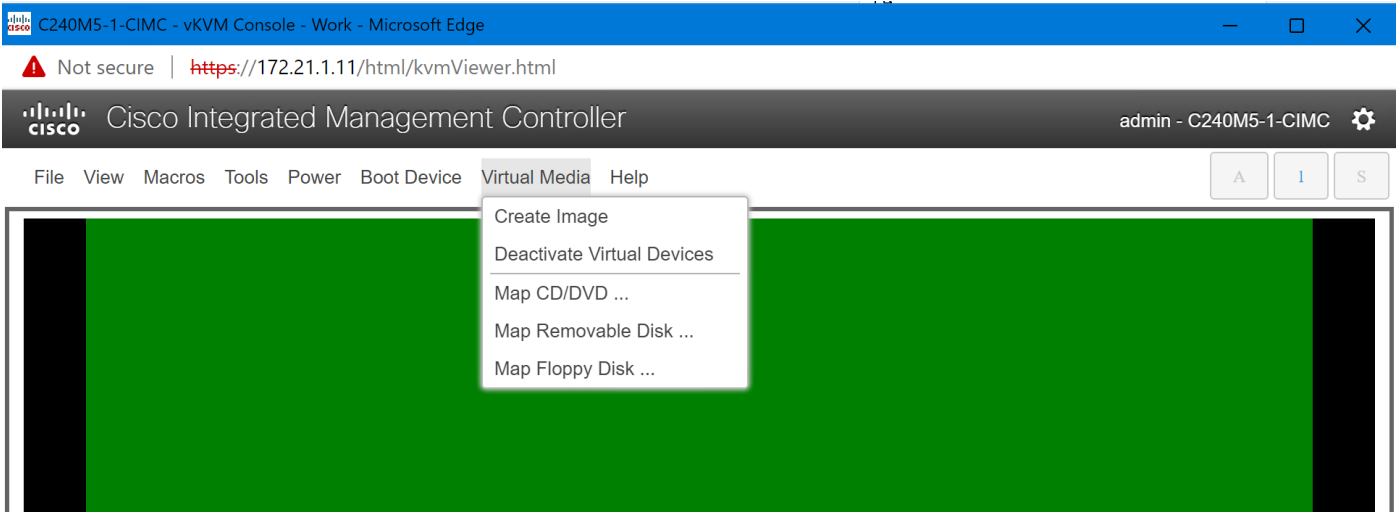
Install SLE Micro

Follow the steps here to install the SLE Micro operating system on the prepared server.


1. In the virtual keyboard, video, and mouse (vKVM) window, click Virtual Media > Activate Virtual Devices.



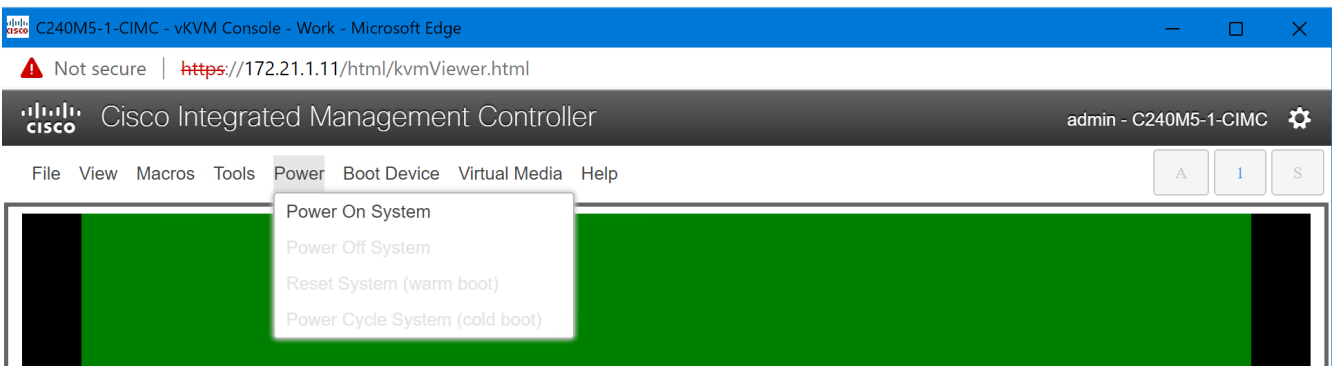
2. Again click Virtual Media and now click Map CD/DVD.



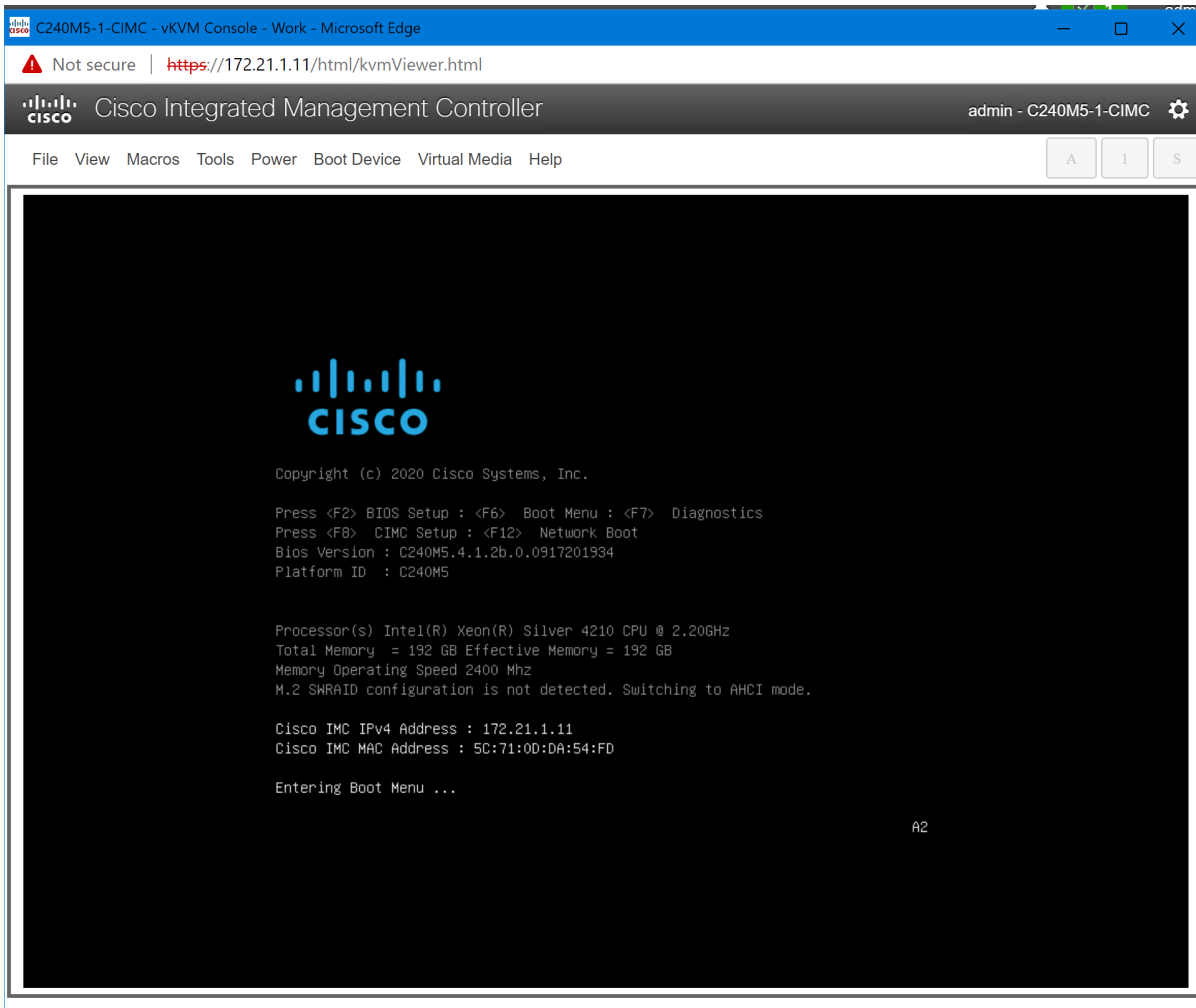
3. Click Browse, select the SLE Micro Media ISO image, and click Map Drive.



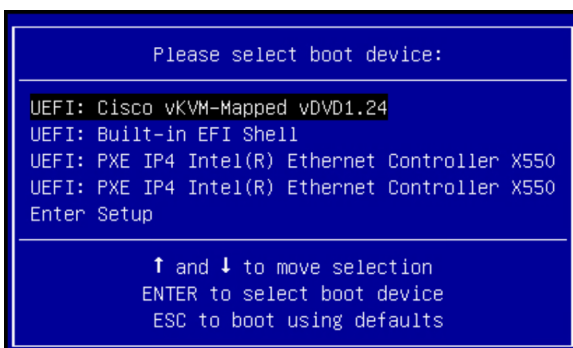
4. Click Power and choose Power On System. In the pop-up window, click OK.



5. As soon the selection menu appears, press F6 to enter the Boot Menu.



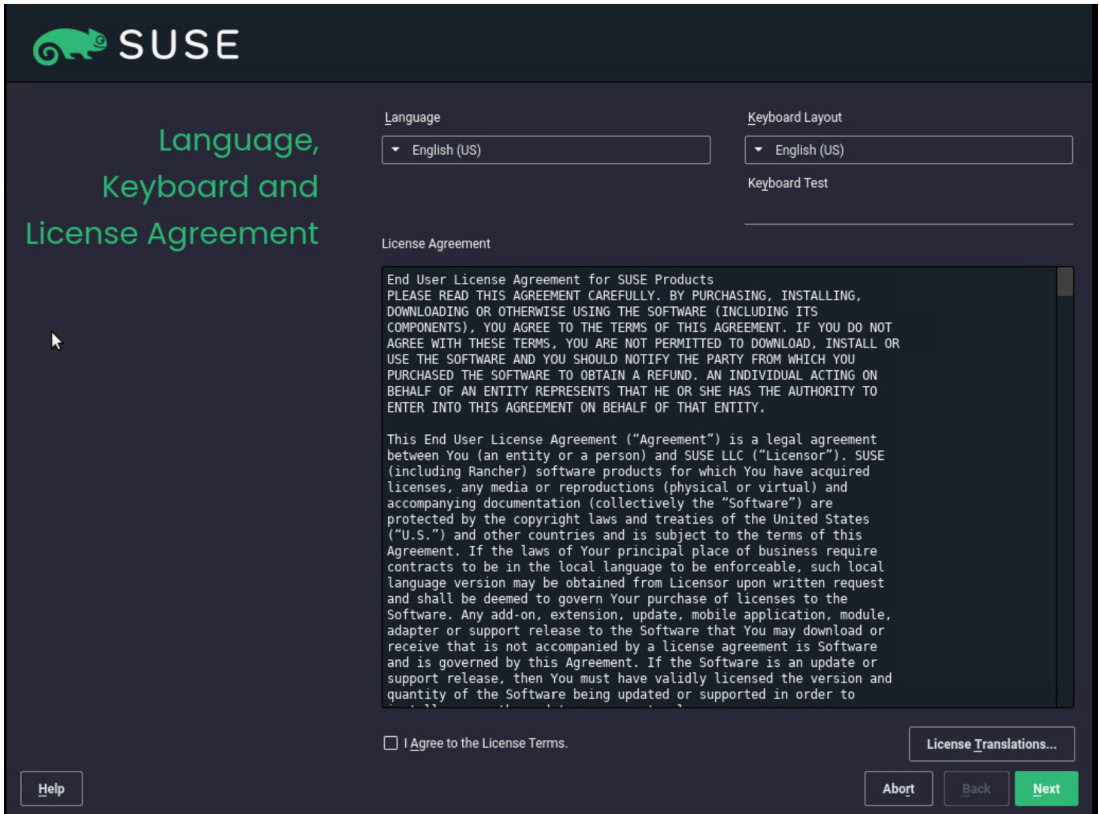
6. Select UEFI: Cisco vKVM-Mapped vDVD and press Enter.



The SUSE installation process will start automatically.



7. Select your language and keyboard layout, agree to the license terms, and click Next.



8. Enter your information to register this installation or select Skip Registration and click Next.

SUSE

Registration

Network Configuration...

SUSE Linux Enterprise Micro 5.1

Please select your preferred method of registration.

Register System via scc.suse.com

E-mail Address

Registration Code

Register System via local RMT Server

Local Registration Server URL

https://rmt.example.com

Skip Registration

Help Abort Back Next

9. Enter the list of Network Time Protocol (NTP) servers and click Next.

SUSE

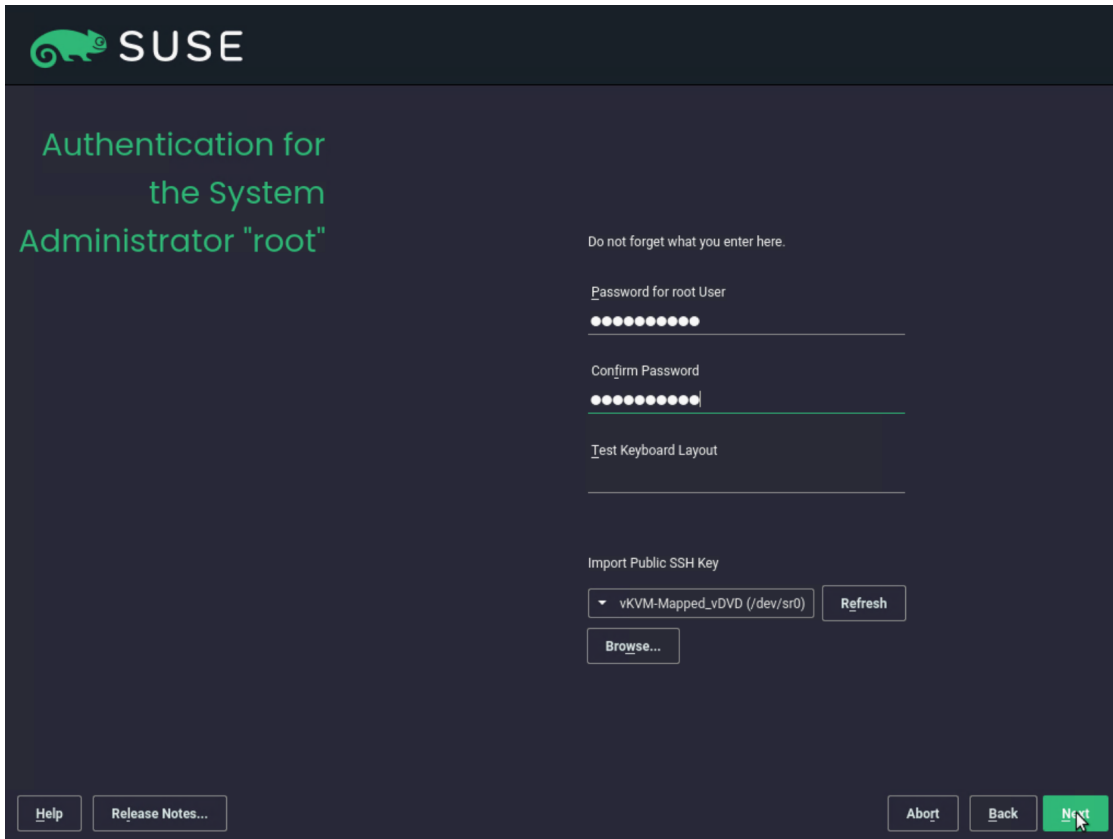
NTP Configuration

NTP Servers (comma or space separated)

1.suse.pool.ntp.org

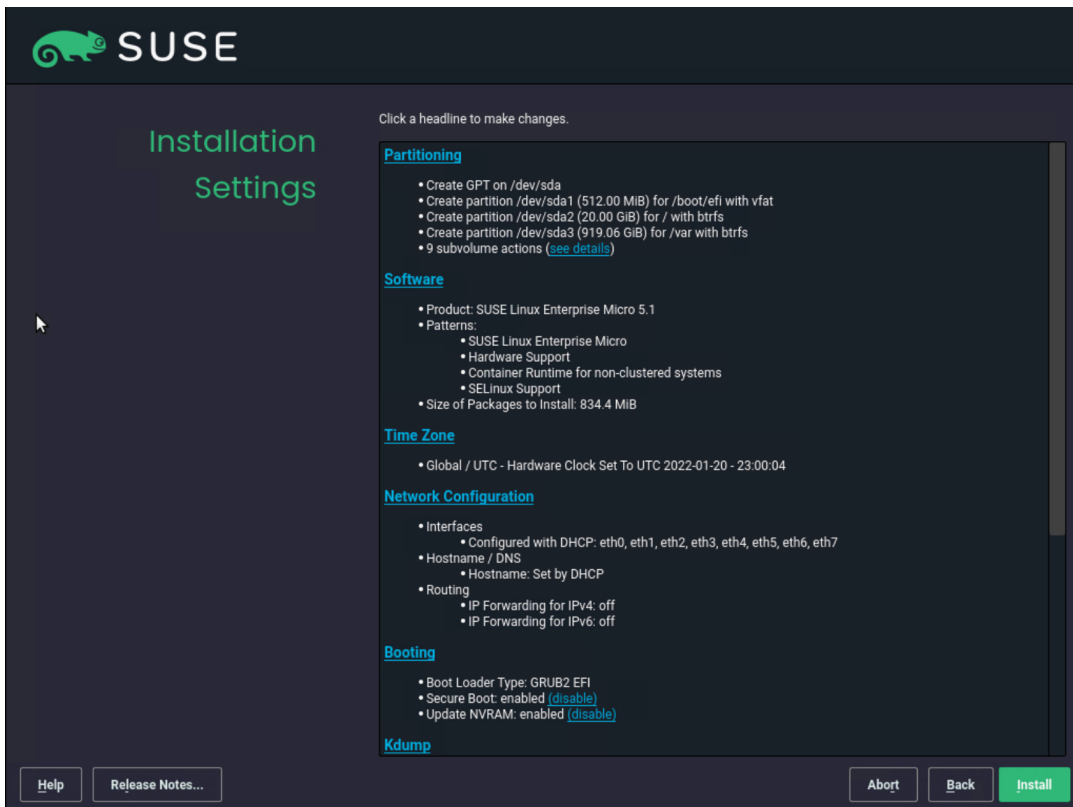
Help Release Notes... Abort Back Next

10. Enter and confirm the password for the user root. Then click Next.



The screenshot shows the 'Authentication for the System Administrator "root"' screen in the SUSE installer. The SUSE logo is at the top left. The title is 'Authentication for the System Administrator "root"'. Below the title, there is a warning: 'Do not forget what you enter here.' There are three input fields: 'Password for root User' with 10 dots, 'Confirm Password' with 10 dots, and 'Test Keyboard Layout' with a blank space. Below these is the 'Import Public SSH Key' section, which has a dropdown menu showing 'vKVM-Mapped_vDVD (/dev/sr0)', a 'Refresh' button, and a 'Browse...' button. At the bottom, there are buttons for 'Help', 'Release Notes...', 'Abort', 'Back', and 'Next'.

11. The partitioning suggested by SUSE is already optimized for container deployments, and no changes are required. Configure the time zone as required for this setup. Click Network Configuration.



The screenshot shows the 'Installation Settings' screen in the SUSE installer. The SUSE logo is at the top left. The title is 'Installation Settings'. Below the title, there is a note: 'Click a headline to make changes.' The screen is divided into several sections: 'Partitioning' (with a list of disk actions), 'Software' (with product and pattern information), 'Time Zone' (with a global/UTC setting), 'Network Configuration' (with interface, hostname, and routing settings), 'Booting' (with boot loader and secure boot settings), and 'Kdump'. At the bottom, there are buttons for 'Help', 'Release Notes...', 'Abort', 'Back', and 'Install'.

12. Click the various devices in the network view and compare the names and MAC addresses with the vNIC list from the IMC. Click Add.

The screenshot shows the SUSE Network Settings application. At the top, there is a header with the SUSE logo and the title "Network Settings". Below the header, there are three tabs: "Overview", "Hostname/DNS", and "Routing". The "Overview" tab is active, displaying a table of network devices. The table has columns for Name, IP Address, Device, and Note. The devices listed are:

Name	IP Address	Device	Note
Ethernet Controller 10G X550T	DHCP	eth6	
Ethernet Controller 10G X550T	DHCP	eth7	
VIC Ethernet NIC	DHCP	eth1	
VIC Ethernet NIC	DHCP	eth4	
VIC Ethernet NIC	DHCP	eth2	
VIC Ethernet NIC	DHCP	eth0	
VIC Ethernet NIC	DHCP	eth5	
VIC Ethernet NIC	DHCP	eth3	

Below the table, there is a detailed view of the selected device, "VIC Ethernet NIC". It shows the MAC address as 3c:57:31:28:bf:5a and the BusID as 0000:40:00:0. Below this, there are three bullet points: "Device Name: eth0", "Configured with dhcp", and "Started automatically at boot". At the bottom of the interface, there are buttons for "Add", "Edit", "Delete", "Help", "Release Notes...", "Abort", "Back", and "Next".

13. We want to create bonding devices for high availability. Select Bonding and click Next.

The screenshot shows the SUSE Add Interface Configuration application. At the top, there is a header with the SUSE logo and the title "Add Interface Configuration". Below the header, there is a list of "Device Type" options with radio buttons next to each:

- Ethernet
- VLAN
- Bridge
- TUN
- TAP
- Bonding
- Dummy
- Wireless
- Infiniband

At the bottom of the interface, there are buttons for "Help", "Release Notes...", "Cancel", and "Next".

14. Enter the IP address, the netmask for the administration traffic network, and a hostname. Click Bond Slaves.

The screenshot shows the 'Address' tab of the 'Network Card Setup' window. The 'General' tab is selected. Under 'Dynamic Address', there are two dropdown menus: 'DHCP' and 'DHCP: both version 4 and 6'. The 'Statically Assign IP Address' option is selected. The 'IP Address' field contains '10.10.10.1', the 'Subnet Mask' field contains '/24', and the 'Hostname' field contains 'k3s-01'. Below these fields is an 'Additional Addresses' section with a table header: 'Address Label', 'IP Address', and 'Netmask'. The table is currently empty.

15. Select the two interfaces created for administration traffic (eth0 and eth1) and use active-backup as the mode. Click Next. In the pop-up window, click Continue.

The screenshot shows the 'Bond Slaves' tab of the 'Network Card Setup' window. The 'Bond Slaves and Order' section contains a list of network interfaces with checkboxes: 'eth0 - eth0 configured' (checked), 'eth1 - eth1 configured' (checked), 'eth2 - eth2 configured' (unchecked), 'eth3 - eth3 configured' (unchecked), 'eth4 - eth4 configured' (unchecked), 'eth5 - eth5 configured' (unchecked), 'eth6 - eth6 configured' (unchecked), and 'eth7 - eth7 configured' (unchecked). Below the list are 'Up' and 'Down' buttons. The 'Bond Driver Options' section has a dropdown menu set to 'mode=active-backup miimon=100'. At the bottom, there are 'Help', 'Release Notes...', 'Cancel', 'Back', and 'Next' buttons.

The dialog box has a green header bar. The text inside reads: 'At least one selected device is already configured. Adapt the configuration for bonding?'. At the bottom, there are two buttons: 'Continue' and 'Cancel'.

16. Click Add to create the bonding device for the access traffic. Select Bonding on the next screen and click Next.

The screenshot shows the SUSE Network Settings application. The 'Overview' tab is active, displaying a table of network interfaces. The 'bond0' interface is highlighted, showing its IP address (10.10.10.1/24) and the devices it is bonded to (eth0 and eth1). Below the table, there is a section for 'bond0 (No hardware information)' with a list of configuration details: Device Name: bond0, Configured with address 10.10.10.1/24, Started automatically at boot, and Bonding Slaves: eth0 eth1. At the bottom, there are buttons for 'Add', 'Edit', 'Delete', 'Help', 'Release Notes...', 'Abort', 'Back', and 'Next'.

Name	IP Address	Device	Note
bond0	10.10.10.1/24	bond0	
Ethernet Controller 10G X550T	DHCP	eth6	
Ethernet Controller 10G X550T	DHCP	eth7	
VIC Ethernet NIC	NONE	eth1	enslaved in bond0
VIC Ethernet NIC	DHCP	eth4	
VIC Ethernet NIC	DHCP	eth2	
VIC Ethernet NIC	NONE	eth0	enslaved in bond0
VIC Ethernet NIC	DHCP	eth5	
VIC Ethernet NIC	DHCP	eth3	

- Device Name: bond0
- Configured with address 10.10.10.1/24
- Started automatically at boot
- Bonding Slaves: eth0 eth1

17. Enter the IP address and netmask for the access traffic connection and a hostname. Click Bond Slaves.

The screenshot shows the SUSE Network Card Setup application. The 'Address' tab is active, displaying configuration options for the network card. The 'Statically Assigned IP Address' option is selected. The IP address is set to 10.20.10.1, the Subnet Mask is /24, and the Hostname is k3s801-access. There is also a section for 'Additional Addresses' with a table for adding more IP addresses.

No Link and IP Setup (Bonding Slaves)

Dynamic Address

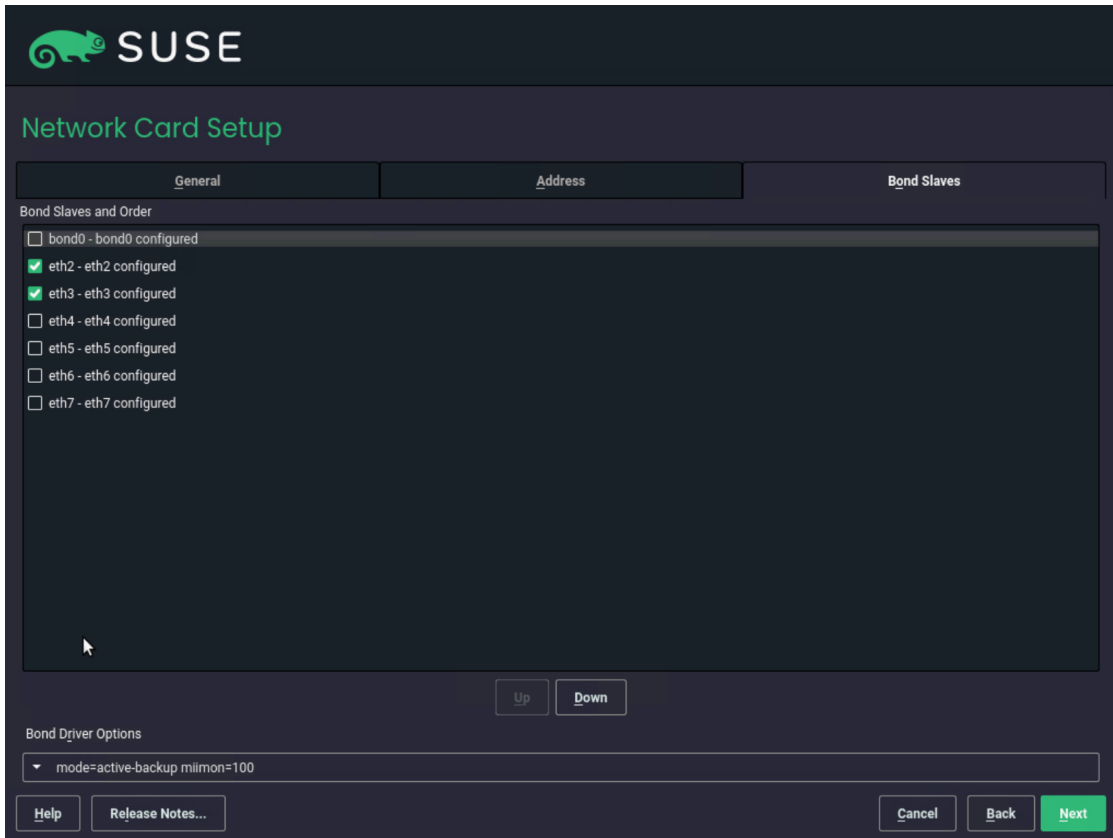
Statically Assigned IP Address

IP Address: 10.20.10.1 Subnet Mask: /24 Hostname: k3s801-access

Additional Addresses

Address Label	IP Address	Netmask
---------------	------------	---------

18. Select the two interfaces created for access traffic (eth2 and eth3) and use active-backup as the mode. Click Next. In the pop-up window, click Continue.



19. Back on the Network Settings screen, click Add to create the bonding for storage traffic. Select Bonding on the next screen and click Next.

The screenshot shows the SUSE Network Settings interface. At the top, there is a header with the SUSE logo and the title "Network Settings". Below the header, there are three tabs: "Overview", "Hostname/DNS", and "Routing". The "Overview" tab is active, displaying a table of network interfaces.

Name	IP Address	Device	Note
bond0	10.10.10.1/24	bond0	
bond1	10.20.10.1/24	bond1	
Ethernet Controller 10G X550T	DHCP	eth6	
Ethernet Controller 10G X550T	DHCP	eth7	
VIC Ethernet NIC	NONE	eth1	enslaved in bond0
VIC Ethernet NIC	DHCP	eth4	
VIC Ethernet NIC	NONE	eth2	enslaved in bond1
VIC Ethernet NIC	NONE	eth0	enslaved in bond0
VIC Ethernet NIC	DHCP	eth5	
VIC Ethernet NIC	NONE	eth3	enslaved in bond1

Below the table, there is a section for "VIC Ethernet NIC" with the following details:

- MAC : 3c:57:31:28:bf:5b
- BusID : 0000:40:00.1
- Device Name: eth1
- Do not assign (e.g. bond or bridge slaves)
- Started automatically at boot
- Bonding master: bond0

At the bottom of the interface, there are buttons for "Add", "Edit", "Delete", "Help", "Release Notes...", "Abort", "Back", and "Next".

20. Enter the IP address and netmask for the storage traffic connection and a hostname. Click Bond Slaves.

The screenshot shows the SUSE Network Card Setup interface. At the top, there is a header with the SUSE logo and the title "Network Card Setup". Below the header, there are three tabs: "General", "Address", and "Bond Slaves". The "Address" tab is active, displaying the configuration for a network card.

The configuration options are:

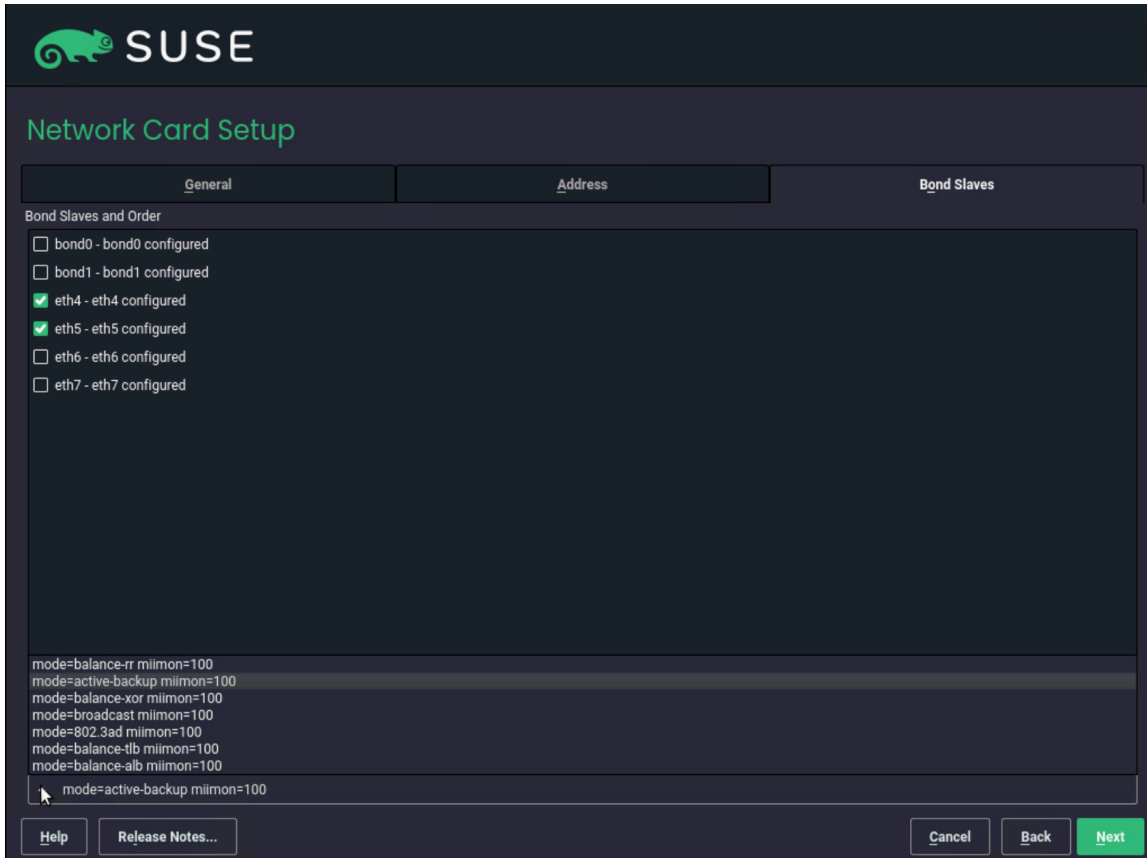
- No Link and IP Setup (Bonding Slaves)
- Dynamic Address
 - DHCP
 - DHCP both version 4 and 6
- Statically Assigned IP Address
 - IP Address: 192.168.112.1
 - Subnet Mask: /24
 - Hostname: k3s-01-data

Below the configuration options, there is a section for "Additional Addresses" with a table for adding more addresses:

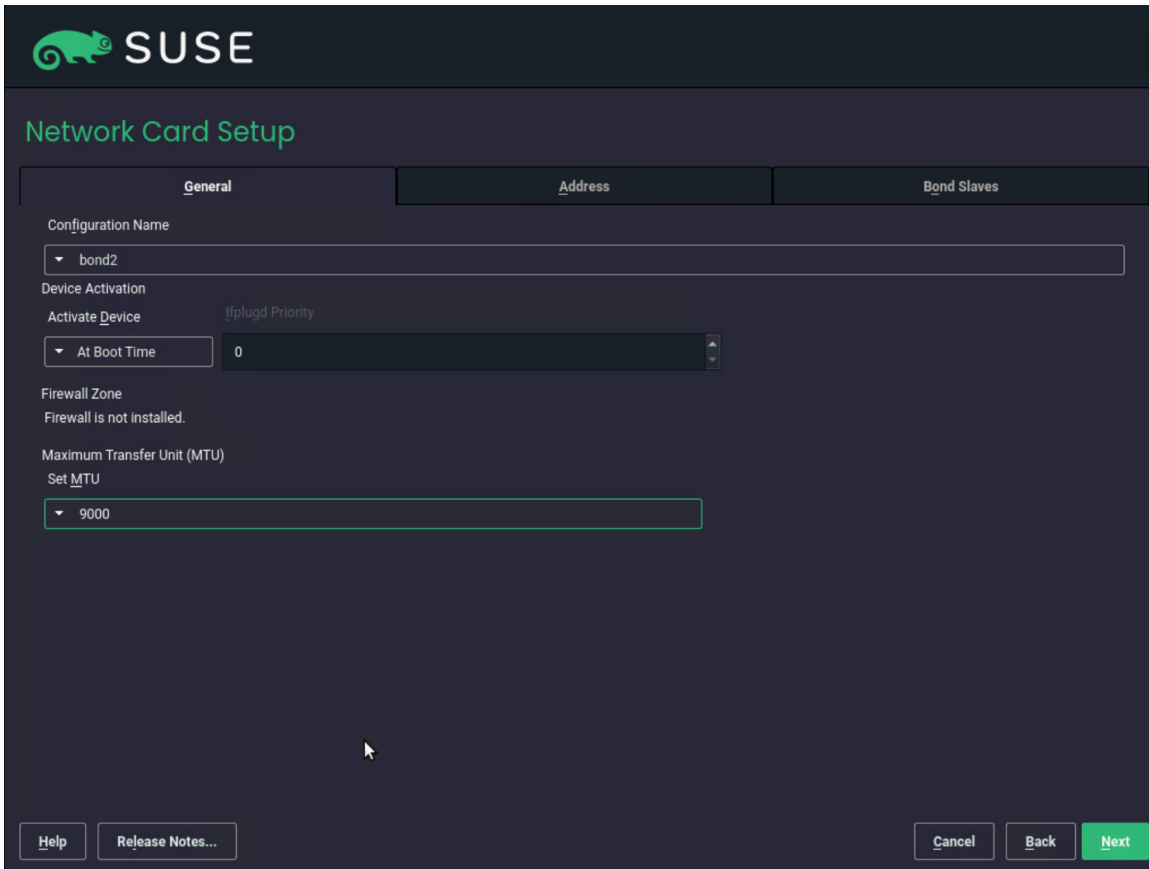
Address Label	IP Address	Netmask

21. Select the two interfaces created for storage traffic (eth4 and eth5). Check with your networking and storage teams to determine whether an active-active bonding option for storage access is possible. An active-active option will increase the maximum throughput between this server and the storage system. In the absence of a clear answer from the network team, use active-backup as the mode.

22. Click General and in the pop-up window click Continue.

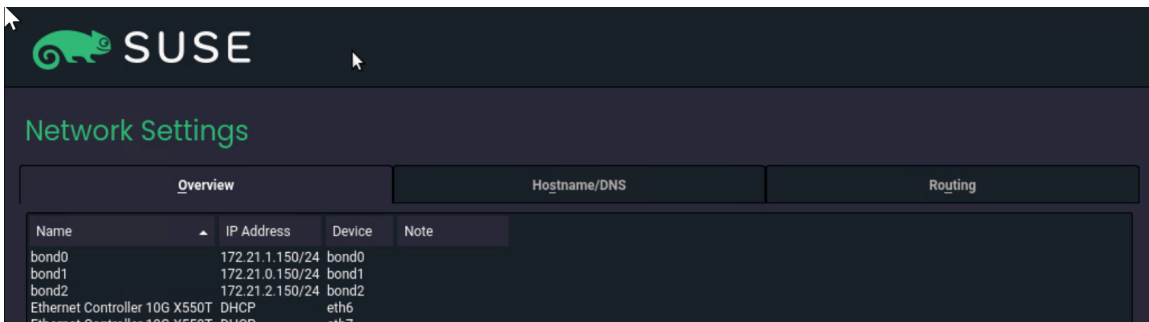


23. Enter **9000** in the field under Set MTU and click Next.



The screenshot shows the 'Network Card Setup' window in SUSE, with the 'General' tab selected. The 'Configuration Name' is set to 'bond2'. Under 'Device Activation', 'Activate Device' is set to 'At Boot Time' and 'Ifplugd Priority' is set to '0'. The 'Firewall Zone' is 'Firewall is not installed'. Under 'Maximum Transfer Unit (MTU)', 'Set MTU' is set to '9000'. At the bottom, there are buttons for 'Help', 'Release Notes...', 'Cancel', 'Back', and 'Next'.

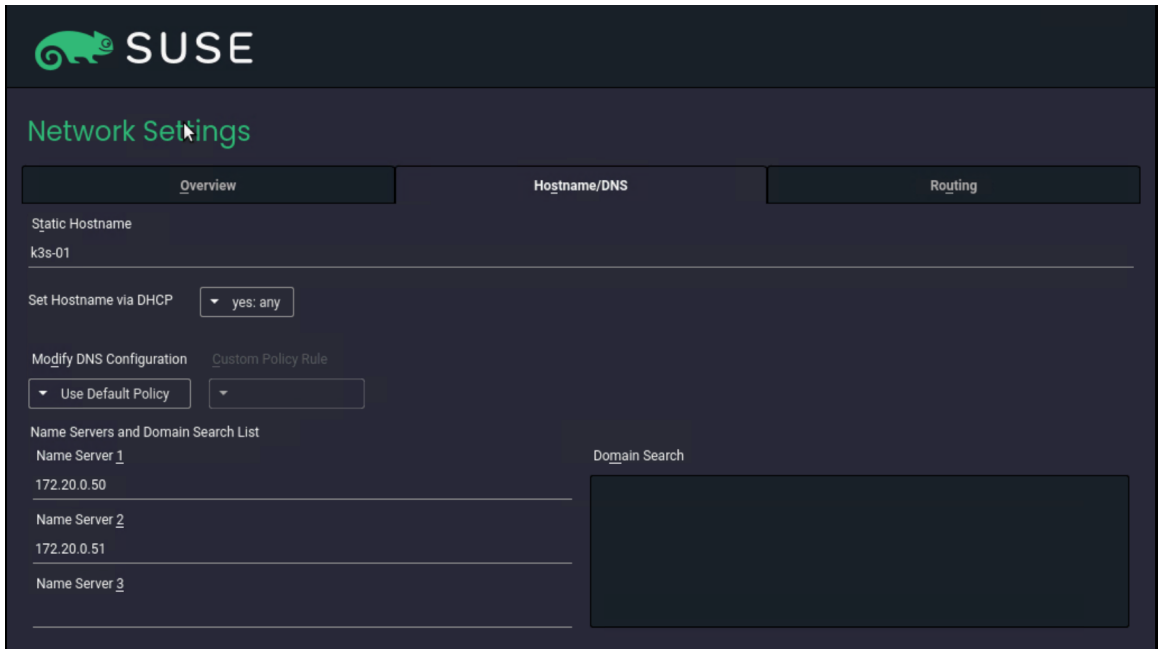
24. Click Hostname / DNS.



The screenshot shows the 'Network Settings' window in SUSE, with the 'Hostname/DNS' tab selected. Below the tabs, there is a table with the following data:

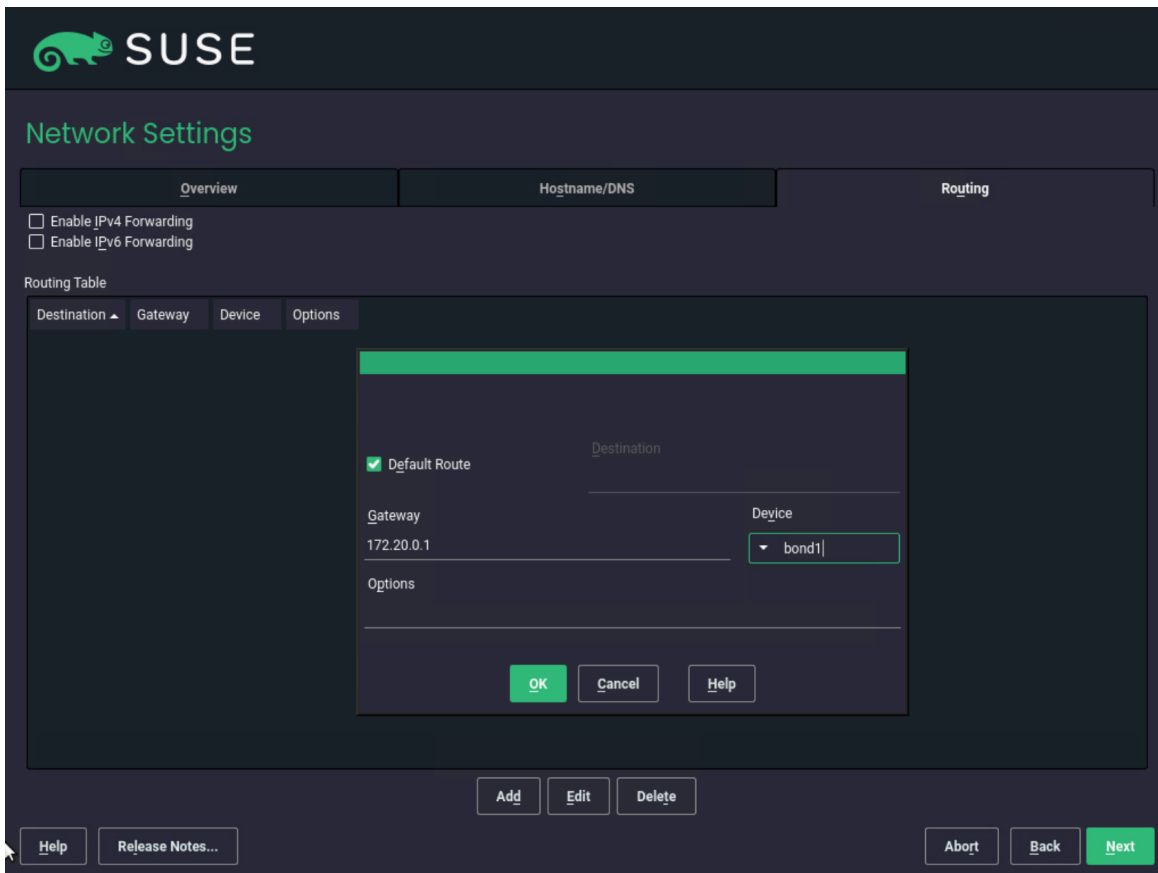
Name	IP Address	Device	Note
bond0	172.21.1.150/24	bond0	
bond1	172.21.0.150/24	bond1	
bond2	172.21.2.150/24	bond2	
Ethernet Controller 10G X550T	DHCP	eth6	
Ethernet Controller 10G X550T	DHCP	eth7	

25. Enter the static hostname for this system and the IP address for at least one name server. Click Routing.



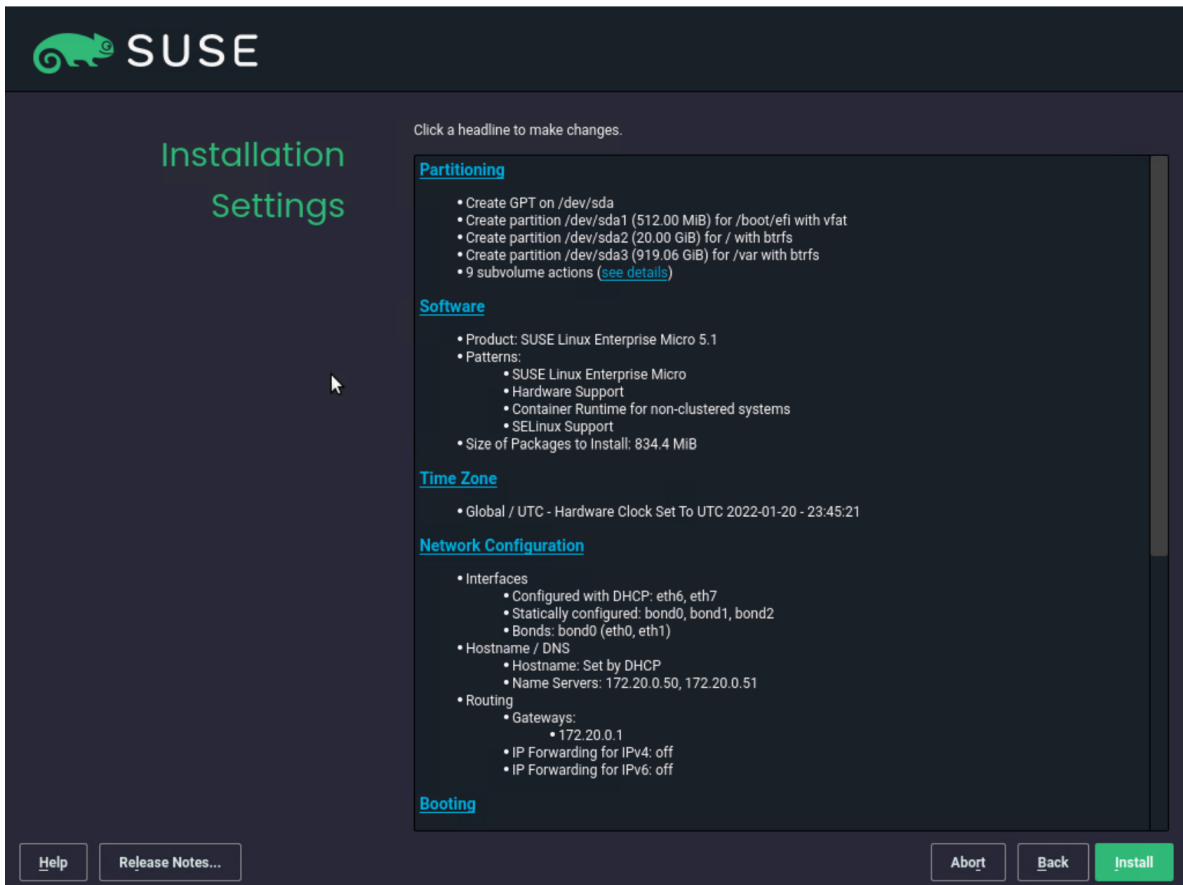
The screenshot shows the SUSE Network Settings interface with the 'Hostname/DNS' tab selected. The 'Static Hostname' field contains 'k3s-01'. The 'Set Hostname via DHCP' dropdown is set to 'yes: any'. Under 'Modify DNS Configuration', 'Use Default Policy' is selected. The 'Name Servers and Domain Search List' section shows three name servers: '172.20.0.50', '172.20.0.51', and an empty field for 'Name Server 3'. A large empty text area is provided for 'Domain Search'.

26. Click Add and in the pop-up window enter at least the default route for your network. Click OK. Click Next.

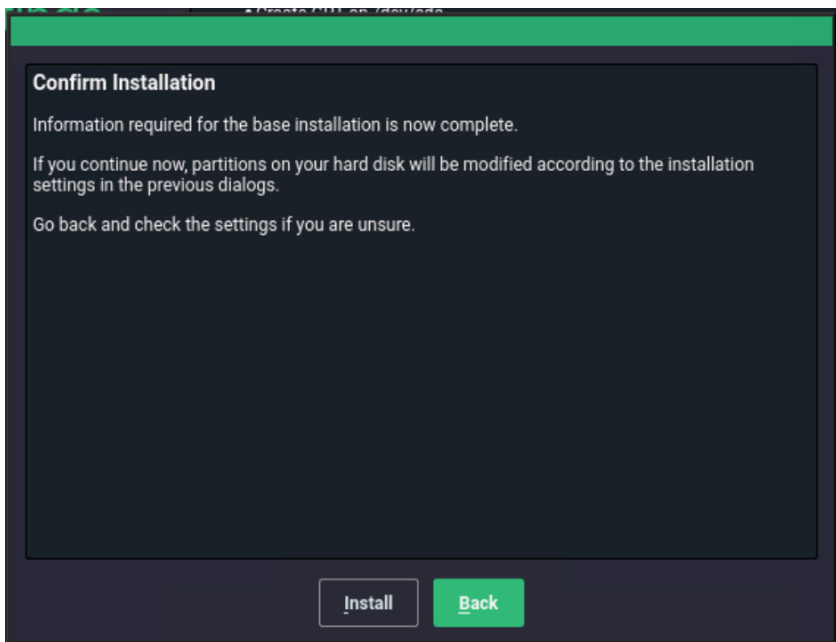


The screenshot shows the SUSE Network Settings interface with the 'Routing' tab selected. The 'Routing Table' section is visible, showing columns for Destination, Gateway, Device, and Options. A pop-up window is open for adding a new route. The 'Default Route' checkbox is checked. The 'Gateway' field contains '172.20.0.1' and the 'Device' dropdown is set to 'bond1'. The 'Options' field is empty. The pop-up window has 'OK', 'Cancel', and 'Help' buttons. At the bottom of the main interface, there are 'Add', 'Edit', and 'Delete' buttons, and a 'Next' button highlighted in green.

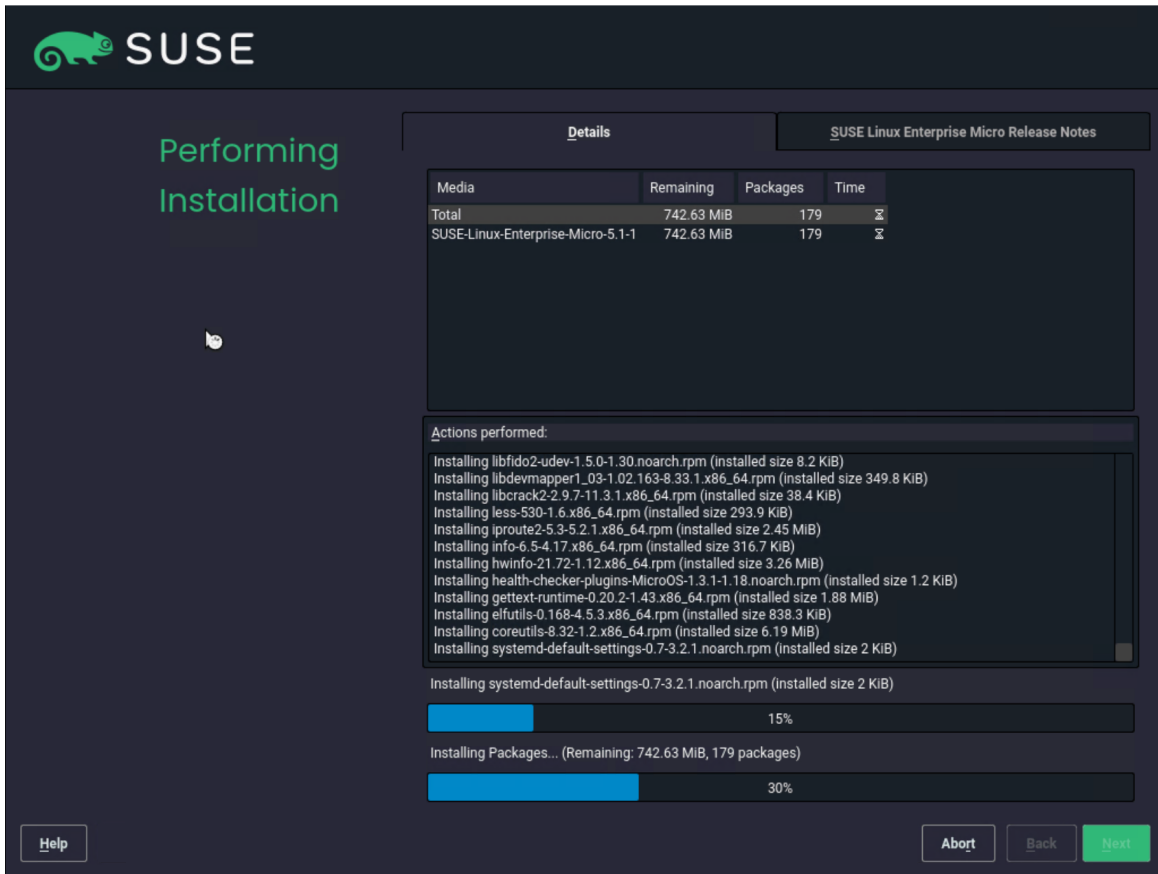
27. Check all the information and then click Install.



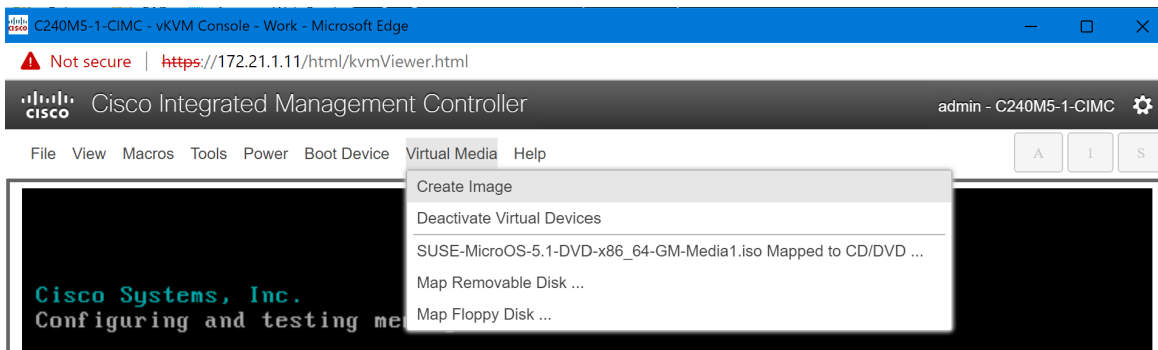
28. Click Install.



The installation process proceeds.



29. You must “eject” the CD/DVD as soon the installation process has finished and the reboot is initiated. Click Virtual Media > “SUSE-MicorOS-5.1-DVD-x86_64-GM-Media1.iso Mapped to CD/DVD” and confirm the ejection by clicking OK in the pop-up window.



After the installation is complete, the system will reboot automatically.



SLE Micro 5.1

Advanced options for SLE Micro 5.1

Start bootloader from a read-only snapshot

The highlighted entry will be executed automatically in 7s.

30. Log on to the system as the user root and using the password provided during the installation process.

```
C240M5-1-CIMC - vKVM Console - Work - Microsoft Edge
Not secure | https://172.21.1.11/html/kvmViewer.html
Cisco Integrated Management Controller admin - C240M5-1-CIMC
File View Macros Tools Power Boot Device Virtual Media Help
[ 4.383094] fnic: DEUCMD2 resource found?
[ 4.490305] fnic: DEUCMD2 resource found?

Welcome to SUSE Linux Enterprise Micro 5.1 (x86_64) - Kernel 5.3.18-59.19-default (tty1).

SSH host key: SHA256:4EcxDJ+YU4aTZCIfREW7Z+Y0t7kLtzUGwNYfHXW+CiA (RSA)
SSH host key: SHA256:5XEwJqPvMBQzZBjLZSPQn2+DC1pTquYQtgBwQLONUBY (DSA)
SSH host key: SHA256:R9bCfBf+v2hS9ZmLowj0QuaSxSU4NIF6TXkouSJ0S7Q (ECDSA)
SSH host key: SHA256:qpuv95M1ehBaTdfPqZb_jkFExMNTFZMOFECU+*zpc1UE (ED25519)
bond0: 172.21.1.150 fe80::3e57:31ff:fe28:bf5a
bond1: 172.21.0.150 fe80::3e57:31ff:fe28:bf5e
bond2: 172.21.2.150 fe80::3e57:31ff:fe28:bf60
eth0:
eth1:
eth2:
eth3:
eth4:
eth5:
eth6: 172.21.1.209 fe80::5e71:dff:feda:5502
eth7: 172.21.1.210 fe80::5e71:dff:feda:5503

k3s-01 login: root
Password:
Last login: Fri Jan 21 09:05:36 on tty1
k3s-01:~# _
```

31. Run the following commands to check the network configuration:

```
k3s-01:~ # ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN group default qlen 1000
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,SLAVE,UP,LOWER_UP> mtu 1500 qdisc mq master bond0 state UP
group default qlen 1000
    link/ether 3c:57:31:28:bf:5a brd ff:ff:ff:ff:ff:ff
    altname enp64s0f0
.
.
.
10: bond2: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 9000 qdisc noqueue state UP group
default qlen 1000
    link/ether 3c:57:31:28:bf:60 brd ff:ff:ff:ff:ff:ff
    inet 172.21.2.150/24 brd 172.21.2.255 scope global bond2
        valid_lft forever preferred_lft forever
    inet6 fe80::3e57:31ff:fe28:bf60/64 scope link
        valid_lft forever preferred_lft forever
11: bond1: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group
default qlen 1000
    link/ether 3c:57:31:28:bf:5e brd ff:ff:ff:ff:ff:ff
    inet 172.21.0.150/24 brd 172.21.0.255 scope global bond1
        valid_lft forever preferred_lft forever
    inet6 fe80::3e57:31ff:fe28:bf5e/64 scope link
        valid_lft forever preferred_lft forever
12: bond0: <BROADCAST,MULTICAST,MASTER,UP,LOWER_UP> mtu 1500 qdisc noqueue state UP group
default qlen 1000
    link/ether 3c:57:31:28:bf:5a brd ff:ff:ff:ff:ff:ff
    inet 172.21.1.150/24 brd 172.21.1.255 scope global bond0
        valid_lft forever preferred_lft forever
    inet6 fe80::3e57:31ff:fe28:bf5a/64 scope link
        valid_lft forever preferred_lft forever
k3s-01:~ #

k3s-01:~ # ip route
default via 172.21.1.1 dev eth6 proto dhcp
172.21.0.0/24 dev bond1 proto kernel scope link src 172.21.0.150
172.21.1.0/24 dev eth6 proto kernel scope link src 172.21.1.209
```



```
172.21.1.0/24 dev eth7 proto kernel scope link src 172.21.1.210
172.21.1.0/24 dev bond0 proto kernel scope link src 172.21.1.150
172.21.2.0/24 dev bond2 proto kernel scope link src 172.21.2.150
k3s-01:~ #
```

```
k3s-01:~ # cat /proc/net/bonding/bond0
```

```
Ethernet Channel Bonding Driver: v3.7.1 (April 27, 2011)
```

```
Bonding Mode: fault-tolerance (active-backup)
```

```
Primary Slave: None
```

```
Currently Active Slave: eth0
```

```
MII Status: up
```

```
MII Polling Interval (ms): 100
```

```
Up Delay (ms): 0
```

```
Down Delay (ms): 0
```

```
Peer Notification Delay (ms): 0
```

```
Slave Interface: eth0
```

```
MII Status: up
```

```
Speed: 25000 Mbps
```

```
Duplex: full
```

```
Link Failure Count: 0
```

```
Permanent HW addr: 3c:57:31:28:bf:5a
```

```
Slave queue ID: 0
```

```
Slave Interface: eth1
```

```
MII Status: up
```

```
Speed: 25000 Mbps
```

```
Duplex: full
```

```
Link Failure Count: 0
```

```
Permanent HW addr: 3c:57:31:28:bf:5b
```

```
Slave queue ID: 0
```

```
k3s-01:~ #
```

```
k3s-01:~ # ping wdf02-4-pdc.wdf02-4-dmz.local. -c 3
```

```
PING wdf02-4-pdc.wdf02-4-dmz.local (172.20.0.50) 56(84) bytes of data.
```

```
64 bytes from wdf02-4-pdc.wdf02-4-dmz.local (172.20.0.50): icmp_seq=1 ttl=126 time=0.254 ms
```

```
64 bytes from wdf02-4-pdc.wdf02-4-dmz.local (172.20.0.50): icmp_seq=2 ttl=126 time=0.259 ms
```

```
64 bytes from wdf02-4-pdc.wdf02-4-dmz.local (172.20.0.50): icmp_seq=3 ttl=126 time=0.379 ms
```

```
--- wdf02-4-pdc.wdf02-4-dmz.local ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2040ms
rtt min/avg/max/mdev = 0.254/0.297/0.379/0.059 ms
k3s-01:~ #
```

```
k3s-01:~ # ping www.google.de. -c 3
PING www.google.de (142.250.179.131) 56(84) bytes of data.
64 bytes from ams17s10-in-f3.1e100.net (142.250.179.131): icmp_seq=1 ttl=115 time=16.9 ms
64 bytes from ams17s10-in-f3.1e100.net (142.250.179.131): icmp_seq=2 ttl=115 time=16.9 ms
64 bytes from ams17s10-in-f3.1e100.net (142.250.179.131): icmp_seq=3 ttl=115 time=16.9 ms
```

```
--- www.google.de ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2003ms
rtt min/avg/max/mdev = 16.915/16.937/16.963/0.019 ms
k3s-01:~ #
```

Install K3s

This section presents the installation procedure for the K3s software as described in [Rancher Docs: K3s - Lightweight Kubernetes](#).

1. Use the curl command to download the K3s software package and install

```
k3s-01:~ # curl -sfL https://get.k3s.io | sh -
[INFO] Finding release for channel stable
[INFO] Using v1.22.5+k3s1 as release
[INFO] Downloading hash https://github.com/k3s-
io/k3s/releases/download/v1.22.5+k3s1/sha256sum-amd64.txt
[INFO] Downloading binary https://github.com/k3s-io/k3s/releases/download/v1.22.5+k3s1/k3s
[INFO] Verifying binary download
[INFO] Installing k3s to /usr/local/bin/k3s
transactional-update 3.5.6 started
Options: --no-selfupdate -d run zypper --gpg-auto-import-keys install -y k3s-selinux
Separate /var detected.
2022-01-21 09:27:10 tukit 3.5.6 started
2022-01-21 09:27:10 Options: --discard -c1 open
2022-01-21 09:27:10 Using snapshot 1 as base for new snapshot 3.
2022-01-21 09:27:10 No previous snapshot to sync with - skipping
ID: 3
2022-01-21 09:27:10 Transaction completed.
2022-01-21 09:27:10 tukit 3.5.6 started
2022-01-21 09:27:10 Options: --discard call 3 zypper --gpg-auto-import-keys install -y k3s-selinux
2022-01-21 09:27:11 Executing `zypper --gpg-auto-import-keys install -y k3s-selinux`:
```

Building repository 'Rancher K3s Common (stable)' cache
.....[done]

Loading repository data...

Reading installed packages...

Resolving package dependencies...

The following NEW package is going to be installed:

k3s-selinux

1 new package to install.

Overall download size: 20.0 KiB. Already cached: 0 B. After the operation, additional 85.1 KiB will be used.

Continue? [y/n/v/...? shows all options] (y): y

Retrieving package k3s-selinux-0.5-1.SLE.noarch (1/1), 20.0 KiB (85.1 KiB unpacked)

Retrieving: k3s-selinux-0.5-1.sle.noarch.rpm

.....[done (713 B/s)]

k3s-selinux-0.5-1.sle.noarch.rpm:

Header V4 RSA/SHA1 Signature, key ID e257814a: NOKEY

V4 RSA/SHA1 Signature, key ID e257814a: NOKEY

Looking for gpg key ID E257814A in cache /var/cache/zypp/pubkeys.

Looking for gpg key ID E257814A in repository Rancher K3s Common (stable).

gpgkey=https://rpm.rancher.io/public.key

Retrieving:

public.key.....[done]

Automatically importing the following key:

Repository: Rancher K3s Common (stable)

Key Fingerprint: C8CF F216 4551 26E9 B9C9 18BE 925E A29A E257 814A

Key Name: Rancher (CI) <ci@rancher.com>

Key Algorithm: RSA 3072

Key Created: Tue Mar 10 22:43:06 2020

Key Expires: (does not expire)

Subkey: AA7E9EC8FE21FDCF 2020-03-10 [does not expire]

Rpm Name: gpg-pubkey-e257814a-5e6817fa

Note: A GPG pubkey is clearly identified by it's fingerprint. Do not rely the keys name. If you are not sure whether the presented key is authentic, ask the repository provider or check his web site. Many provider maintain a web page showing the fingerprints of the GPG keys they are using.

Checking for file

conflicts:.....[done]

```
(1/1) Installing: k3s-selinux-0.5-1.sle.noarch
.....[done]
Executing %posttrans
scripts.....[done]
2022-01-21 09:27:21 Application returned with exit status 0.
2022-01-21 09:27:22 Transaction completed.
2022-01-21 09:27:22 tukit 3.5.6 started
2022-01-21 09:27:22 Options: --discard close 3
2022-01-21 09:27:22 New default snapshot is #3 (/.snapshots/3/snapshot).
2022-01-21 09:27:22 Transaction completed.
```

Please reboot your machine to activate the changes and avoid data loss.
New default snapshot is #3 (/.snapshots/3/snapshot).

```
transactional-update finished
[INFO] Creating /usr/local/bin/kubect1 symlink to k3s
[INFO] Creating /usr/local/bin/crictl symlink to k3s
[INFO] Creating /usr/local/bin/ctr symlink to k3s
[INFO] Creating killall script /usr/local/bin/k3s-killall.sh
[INFO] Creating uninstall script /usr/local/bin/k3s-uninstall.sh
[INFO] env: Creating environment file /etc/systemd/system/k3s.service.env
[INFO] systemd: Creating service file /etc/systemd/system/k3s.service
[INFO] systemd: Enabling k3s unit
Created symlink /etc/systemd/system/multi-user.target.wants/k3s.service ->
/etc/systemd/system/k3s.service.
k3s-01:~ #
```

2. Use the **systemctl** command to start the K3s server and check the status.

```
k3s-01:~ # systemctl start k3s
k3s-01:~ # systemctl status k3s
● k3s.service - Lightweight Kubernetes
   Loaded: loaded (/etc/systemd/system/k3s.service; enabled; vendor preset: disabled)
   Active: active (running) since Fri 2022-01-21 09:37:50 UTC; 7min ago
     Docs: https://k3s.io
   Process: 2583 ExecStartPre=/bin/sh -xc ! /usr/bin/systemctl is-enabled --quiet nm-cloud-
setup.service (code=exited, s>
   Process: 2596 ExecStartPre=/sbin/modprobe br_netfilter (code=exited, status=0/SUCCESS)
   Process: 2610 ExecStartPre=/sbin/modprobe overlay (code=exited, status=0/SUCCESS)
  Main PID: 2611 (k3s-server)
     Tasks: 225
...
k3s-01:~ #
```

3. Get basic information from the installed K3s cluster.

```
k3s-01:~ # kubectl cluster-info
```

```
Kubernetes control plane is running at https://127.0.0.1:6443
```

```
CoreDNS is running at https://127.0.0.1:6443/api/v1/namespaces/kube-system/services/kube-dns:dns/proxy
```

```
Metrics-server is running at https://127.0.0.1:6443/api/v1/namespaces/kube-system/services/https:metrics-server:/proxy
```

To further debug and diagnose cluster problems, use 'kubectl cluster-info dump'.

```
k3s-01:~ #
```

```
k3s-01:~ # kubectl get nodes -o wide
```

NAME	STATUS	ROLES	AGE	VERSION	INTERNAL-IP	EXTERNAL-IP
OS-IMAGE			KERNEL-VERSION		CONTAINER-RUNTIME	
k3s-01	Ready	control-plane,master	15m	v1.22.5+k3s1	172.21.1.209	<none>
SUSE Linux Enterprise Micro 5.1			5.3.18-59.19-default		containerd://1.5.8-k3s1	

```
k3s-01:~ #
```

```
k3s-01:~ # kubectl get all -A
```

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	pod/local-path-provisioner-64ffb68fd-7qs4m	1/1	Running	1 (11m ago)	16m
kube-system	pod/metrics-server-9cf544f65-nxrd2	1/1	Running	0	16m
kube-system	pod/helm-install-traefik-crd--1-gln5p	0/1	Completed	0	16m
kube-system	pod/helm-install-traefik--1-sf5dz	0/1	Completed	1	16m
kube-system	pod/svclb-traefik-24sf4	2/2	Running	0	11m
kube-system	pod/coredns-85cb69466-vwbkc	1/1	Running	1 (11m ago)	16m
kube-system	pod/traefik-786ff64748-x4cz5	1/1	Running	0	11m

NAMESPACE	NAME	TYPE	CLUSTER-IP	EXTERNAL-IP	PORT(S)
default	service/kubernetes	ClusterIP	10.43.0.1	<none>	443/TCP
16m					
kube-system	service/kube-dns	ClusterIP	10.43.0.10	<none>	
53/UDP, 53/TCP, 9153/TCP					16m
kube-system	service/metrics-server	ClusterIP	10.43.136.93	<none>	443/TCP
16m					
kube-system	service/traefik	LoadBalancer	10.43.32.86	172.21.1.209	
80:32380/TCP, 443:32713/TCP					11m

NAMESPACE	NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE
kube-system	daemonset.apps/svclb-traefik	1	1	1	1	1
<none>	11m					

NAMESPACE	NAME	READY	UP-TO-DATE	AVAILABLE	AGE
kube-system	deployment.apps/local-path-provisioner	1/1	1	1	16m
kube-system	deployment.apps/coredns	1/1	1	1	16m

```

kube-system deployment.apps/metrics-server 1/1 1 1 16m
kube-system deployment.apps/traefik 1/1 1 1 11m

NAMESPACE NAME DESIRED CURRENT READY AGE
kube-system replicaset.apps/local-path-provisioner-64ffb68fd 1 1 1 16m
kube-system replicaset.apps/coredns-85cb69466 1 1 1 16m
kube-system replicaset.apps/metrics-server-9cf544f65 1 1 1 16m
kube-system replicaset.apps/traefik-786ff64748 1 1 1 11m

NAMESPACE NAME COMPLETIONS DURATION AGE
kube-system job.batch/helm-install-traefik-crd 1/1 5m16s 16m
kube-system job.batch/helm-install-traefik 1/1 5m17s 16m
k3s-01:~ #

```

The system is now installed and is ready for more specific configurations dependent on local requirements.

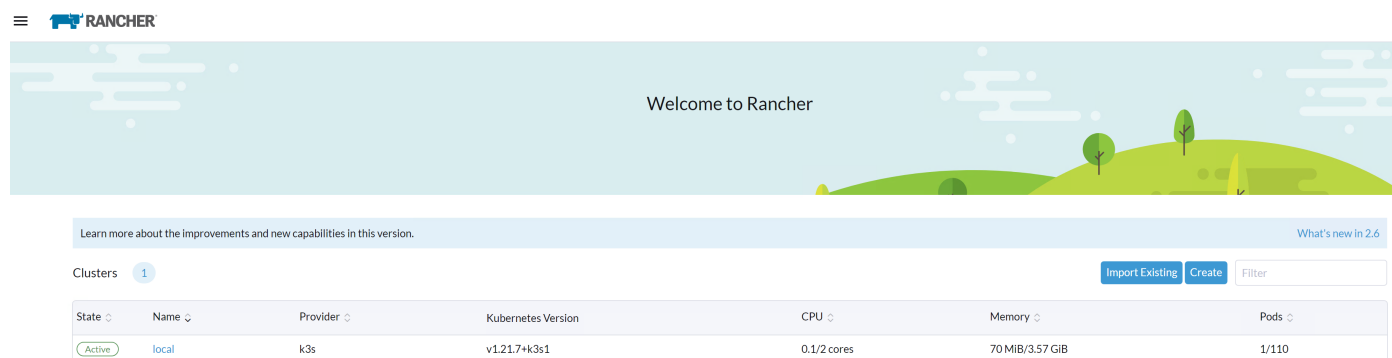
K3s integration into the workload management tool

Many options are available to manage a Kubernetes landscape with multiple clusters, with different workloads, and at different locations. We tested two options: integration into SUSE Rancher Kubernetes Operations Platform and integration into the Rafay Kubernetes Operations Platform.

Integrate into Rancher Kubernetes Operations Platform

The obvious option for managing landscapes with SLE Micro and K3s components is SUSE Rancher. This section shows how to integrate a K3s system into the Rancher Kubernetes Operations Platform.

1. In the Rancher console, navigate to the list of clusters and click Import Existing.



2. Click Generic.

The screenshot shows the Rancher Cluster Management interface. On the left, there is a navigation menu with options: Clusters (1), Cloud Credentials, Drivers, Pod Security Policies, RKE1 Configuration, and Advanced. The main content area is titled 'Cluster: Import' and contains the text 'Import any Kubernetes cluster'. Below this text, there is a box with a Kubernetes logo and the word 'Generic' next to it, which is highlighted with a red border.

3. Enter a cluster name and click Create.

The screenshot shows the Rancher Cluster Management interface with the 'Cluster: Import Generic' form. The 'Cluster Name' field is filled with 'k3s-01'. The 'Cluster Description' field is empty. Below the form, there is a 'Member Roles' section with a table showing the 'Default Admin (admin)' role with 'Local' agent environment vars. A blue 'Add' button is visible. At the bottom right, there are three buttons: 'Cancel', 'Edit as YAML', and 'Create'.

4. Follow the steps shown on the next screen and click Done.

The screenshot shows the Rancher Cluster Management interface with the 'Cluster: k3s-01' page in a 'Pending' state. The page displays instructions for importing a Kubernetes cluster into Rancher. The instructions are as follows:

```
Run the kubect1 command below on an existing Kubernetes cluster running a supported Kubernetes version to import it into Rancher:
```

```
kubect1 apply -f https://172.20.0.102/v3/import/stz7d8gnrzn12c6pgkc9bwvmpqdwwvt5jsgcb2w2fncb69gk722g2_c-m-hr1wq68n.yaml
```

If you get a "certificate signed by unknown authority" error, your Rancher installation has a self-signed or untrusted SSL certificate. Run the command below instead to bypass the certificate verification:

```
curl --insecure -sL https://172.20.0.102/v3/import/stz7d8gnrzn12c6pgkc9bwvmpqdwwvt5jsgcb2w2fncb69gk722g2_c-m-hr1wq68n.yaml | kubect1 apply -f -
```

If you get permission errors creating some of the resources, your user may not have the `cluster-admin` role. Use this command to apply it:

```
kubect1 create clusterrolebinding cluster-admin-binding --clusterrole cluster-admin --user <your username from your kubeconfig>
```

5. Log on to the installed k3s system and run the listed commands from the preceding screen.

```
k3s-01:~ # kubect1 create clusterrolebinding cluster-admin-binding \  
> --clusterrole cluster-admin \  
> --user root  
clusterrolebinding.rbac.authorization.k8s.io/cluster-admin-binding created  
k3s-01:~ #  
k3s-01:~ # curl --insecure -sL  
https://172.20.0.102/v3/import/stz7d8gnrzn12c6pgkc9bwvmpqdwwvt5jsgcb2w2fncb69gk722g2_c-m-  
hr1wq68n.yaml | kubect1 apply -f -
```

```

clusterrole.rbac.authorization.k8s.io/proxy-clusterrole-kubeapiserver created
clusterrolebinding.rbac.authorization.k8s.io/proxy-role-binding-kubernetes-master created
namespace/cattle-system created
serviceaccount/cattle created
clusterrolebinding.rbac.authorization.k8s.io/cattle-admin-binding created
secret/cattle-credentials-fad2056 created
clusterrole.rbac.authorization.k8s.io/cattle-admin created
Warning:
spec.template.spec.affinity.nodeAffinity.requiredDuringSchedulingIgnoredDuringExecution.node
SelectorTerms[0].matchExpressions[0].key: beta.kubernetes.io/os is deprecated since v1.14;
use "kubernetes.io/os" instead
deployment.apps/cattle-cluster-agent created
service/cattle-cluster-agent created
k3s-01:~ #
k3s-01:~ #
k3s-01:~ # kubectl get all -n cattle-system
NAME                                READY   STATUS    RESTARTS   AGE
pod/cattle-cluster-agent-56d66975fc-t56mz  1/1     Running   0           60s

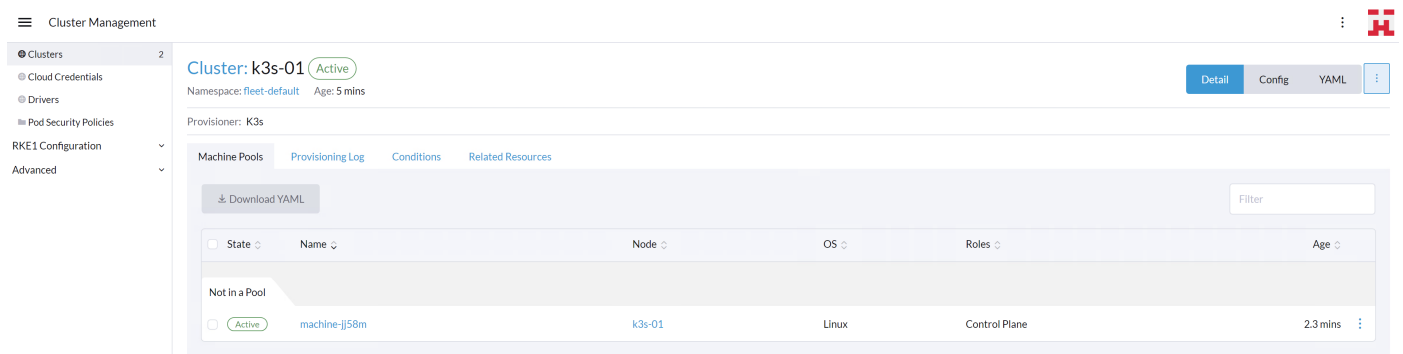
NAME                                TYPE          CLUSTER-IP    EXTERNAL-IP  PORT(S)          AGE
service/cattle-cluster-agent         ClusterIP     10.43.118.86  <none>       80/TCP,443/TCP  3m20s

NAME                                READY   UP-TO-DATE   AVAILABLE   AGE
deployment.apps/cattle-cluster-agent  1/1     1             1           3m20s

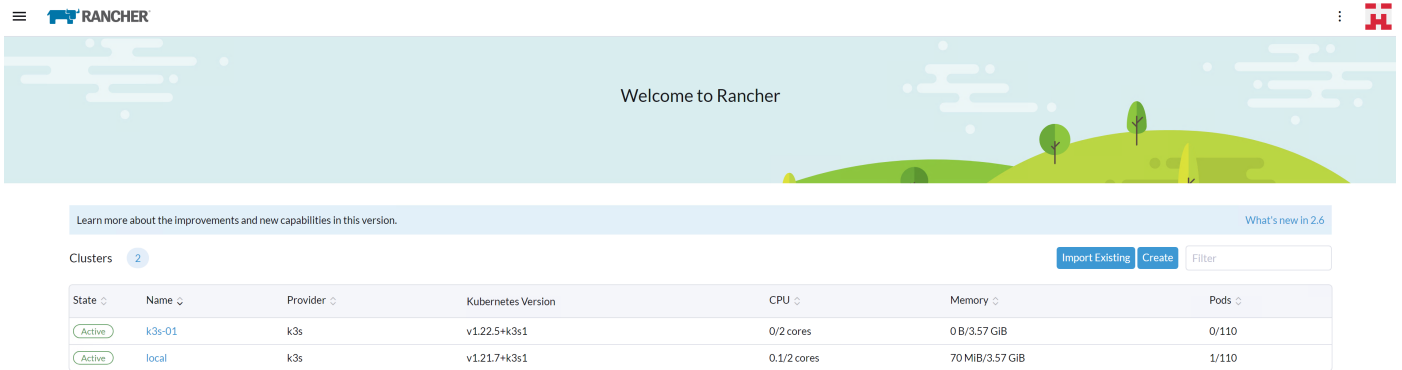
NAME                                DESIRED   CURRENT   READY   AGE
replicaset.apps/cattle-cluster-agent-56d66975fc  1         1         1       60s
replicaset.apps/cattle-cluster-agent-857c647888  0         0         0       3m20s
k3s-01:~ #

```

6. Return to the Rancher user interface. The cluster is now shown as Active.



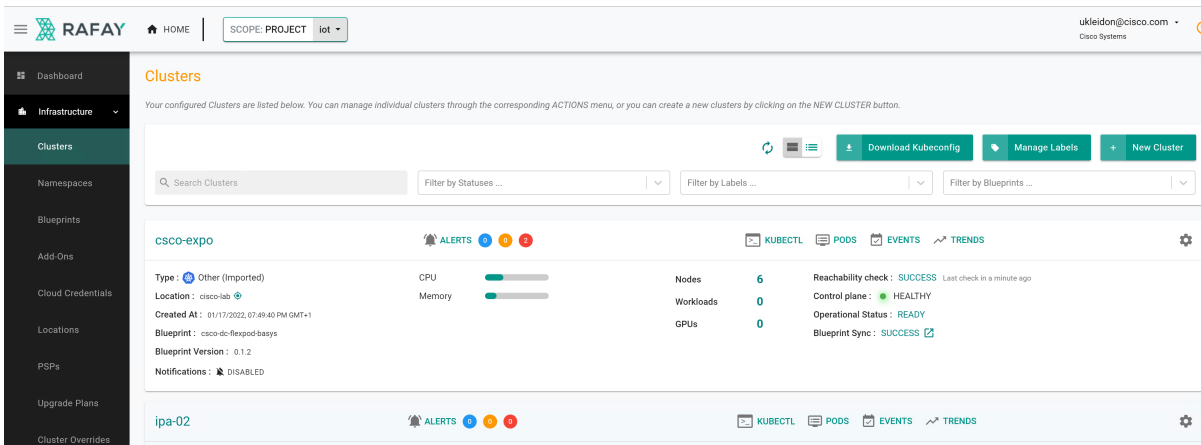
7. View the home screen. The high-level information of the cluster is shown in the home screen of Rancher.



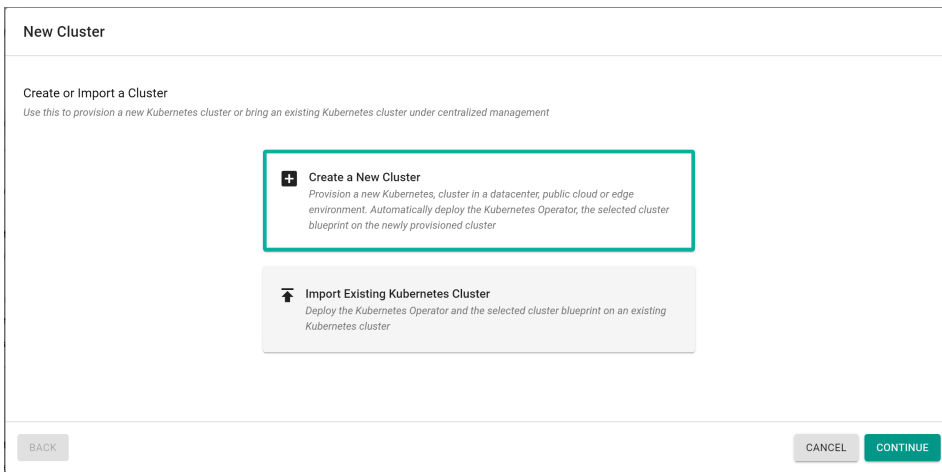
Integrate into Rafay Kubernetes Operations Platform

To demonstrate the manageability of an SLE Micro and K3s system with another tool, this section shows integration into the Rafay Kubernetes management console.

1. In the Rafay console, navigate to the list of clusters in the project of choice and click New Cluster.



2. Click Import Existing Kubernetes Cluster and click Continue.



- Click Data center / Edge and then click Other. Enter a name for the new cluster and a description if wanted. Click Continue.

- Select the location and deployment blueprint for this setup. If this is the first time a K3s cluster will be integrated into the Rafay system, the best practice is to start with the blueprint minimal or default. Those are the basic blueprints from Rafay to make the system work (minimal) or to add components such as monitoring and reporting (default). Then click Continue.

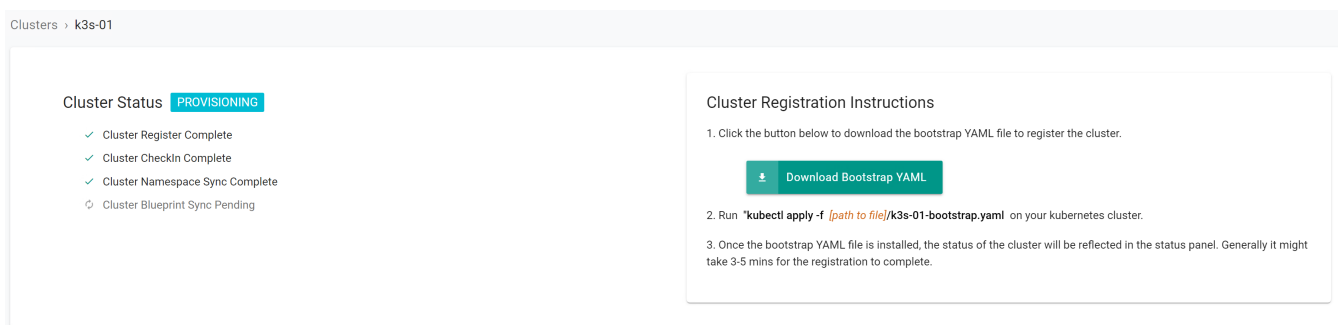
- Download the Bootstrap YAML file to the K3s system.

- Log on to the K3s system and apply the bootstrap file.

```
k3s-01:~ # ls -l /tmp/k3s-01-bootstrap.yaml
-rwxr-xr-x 1 root root 13801 Jan 21 11:16 /tmp/k3s-01-bootstrap.yaml
k3s-01:~ #
```

```
k3s-01:~ # kubectl apply -f /tmp/k3s-01-bootstrap.yaml
namespace/rafay-system created
serviceaccount/system-sa created
Warning: policy/v1beta1 PodSecurityPolicy is deprecated in v1.21+, unavailable in v1.25+
podsecuritypolicy.policy/rafay-privileged-psp created
clusterrole.rbac.authorization.k8s.io/rafay:manager created
clusterrolebinding.rbac.authorization.k8s.io/rafay:rafay-system:manager-rolebinding created
clusterrole.rbac.authorization.k8s.io/rafay:proxy-role created
clusterrolebinding.rbac.authorization.k8s.io/rafay:rafay-system:proxy-rolebinding created
priorityclass.scheduling.k8s.io/rafay-cluster-critical created
role.rbac.authorization.k8s.io/rafay:leader-election-role created
rolebinding.rbac.authorization.k8s.io/rafay:leader-election-rolebinding created
customresourcedefinition.apiextensions.k8s.io/namespaces.cluster.rafay.dev created
customresourcedefinition.apiextensions.k8s.io/tasklets.cluster.rafay.dev created
customresourcedefinition.apiextensions.k8s.io/tasks.cluster.rafay.dev created
service/controller-manager-metrics-service-v3 created
deployment.apps/controller-manager-v3 created
configmap/connector-config-v3 created
configmap/proxy-config-v3 created
deployment.apps/rafay-connector-v3 created
service/rafay-drift-v3 created
validatingwebhookconfiguration.admissionregistration.k8s.io/rafay-drift-validate-v3 created
k3s-01:~ #
```

The process is shown in the Rafay console.



The screenshot displays the Rafay console interface for a cluster named 'k3s-01'. The 'Cluster Status' section is labeled 'PROVISIONING' and lists four items: 'Cluster Register Complete', 'Cluster Checkin Complete', 'Cluster Namespace Sync Complete', and 'Cluster Blueprint Sync Pending'. The 'Cluster Registration Instructions' section provides a three-step guide: 1. Click the 'Download Bootstrap YAML' button. 2. Run the command 'kubectl apply -f [path to file]/k3s-01-bootstrap.yaml' on the Kubernetes cluster. 3. Once the bootstrap YAML file is installed, the cluster status will be reflected in the status panel, typically taking 3-5 minutes to complete.

7. After the deployment is finished, the cluster is shown in the list with basic information about the status.

The screenshot shows a web interface for managing clusters. At the top, there's a search bar and filters for statuses, labels, and blueprints. Below that, a cluster named 'k3s-01' is highlighted. It has 1 node, 0 workloads, and 0 GPUs. The status is 'READY' and 'HEALTHY'. There are also alerts and trends icons visible.

8. On the K3s system, a new namespace rafay-system is created to enable communication between the Rafay Kubernetes Operations Platform and the local K3s system.

```
k3s-01:~ # kubectl get all -A
```

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
kube-system	pod/local-path-provisioner-64ffb68fd-7qs4m	1/1	Running	1 (110m ago)	115m
kube-system	pod/metrics-server-9cf544f65-nxrd2	1/1	Running	0	115m
kube-system	pod/helm-install-traefik-crd--1-gln5p	0/1	Completed	0	115m
kube-system	pod/helm-install-traefik--1-sf5dz	0/1	Completed	1	115m
kube-system	pod/svclb-traefik-24sf4	2/2	Running	0	110m
kube-system	pod/coredns-85cb69466-vwbkc	1/1	Running	1 (110m ago)	115m
kube-system	pod/traefik-786ff64748-x4cz5	1/1	Running	0	110m
rafay-system	pod/edge-client-8c7748dfb-sk416	1/1	Running	0	8m21s
rafay-system	pod/relay-agent-78d645bc89-9w6qw	1/1	Running	0	8m20s
rafay-system	pod/controller-manager-v3-6bb696cc8b-5bsch	1/1	Running	0	6m13s
rafay-system	pod/rafay-connector-v3-6c8dcf8cf9-9m84r	1/1	Running	1 (5m33s ago)	6m14s

NAMESPACE	NAME	TYPE	CLUSTER-IP
EXTERNAL-IP	PORT(S)		
default	service/kubernetes	ClusterIP	10.43.0.1
<none>	443/TCP		115m
kube-system	service/kube-dns	ClusterIP	10.43.0.10
<none>	53/UDP, 53/TCP, 9153/TCP		115m
kube-system	service/metrics-server	ClusterIP	10.43.136.93
<none>	443/TCP		115m
kube-system	service/traefik	LoadBalancer	10.43.32.86
172.21.1.209	80:32380/TCP, 443:32713/TCP		110m
rafay-system	service/controller-manager-metrics-service-v3	ClusterIP	10.43.9.227
<none>	8443/TCP		10m
rafay-system	service/rafay-drift-v3	ClusterIP	10.43.2.198
<none>	8081/TCP		10m

NAMESPACE	NAME	DESIRED	CURRENT	READY	UP-TO-DATE	AVAILABLE
kube-system	daemonset.apps/svclb-traefik	1	1	1	1	1
<none>	110m					

NAMESPACE	NAME	READY	UP-TO-DATE	AVAILABLE	AGE
kube-system	deployment.apps/local-path-provisioner	1/1	1	1	115m
kube-system	deployment.apps/coredns	1/1	1	1	115m
kube-system	deployment.apps/metrics-server	1/1	1	1	115m
kube-system	deployment.apps/traefik	1/1	1	1	110m
rafay-system	deployment.apps/edge-client	1/1	1	1	8m22s
rafay-system	deployment.apps/relay-agent	1/1	1	1	8m21s
rafay-system	deployment.apps/controller-manager-v3	1/1	1	1	10m
rafay-system	deployment.apps/rafay-connector-v3	1/1	1	1	10m

NAMESPACE	NAME	DESIRED	CURRENT	READY	AGE
kube-system	replicaset.apps/local-path-provisioner-64ffb68fd	1	1	1	115m
kube-system	replicaset.apps/coredns-85cb69466	1	1	1	115m
kube-system	replicaset.apps/metrics-server-9cf544f65	1	1	1	115m
kube-system	replicaset.apps/traefik-786ff64748	1	1	1	110m
rafay-system	replicaset.apps/edge-client-8c7748dfb	1	1	1	8m22s
rafay-system	replicaset.apps/relay-agent-78d645bc89	1	1	1	8m21s
rafay-system	replicaset.apps/rafay-connector-v3-88ff764c5	0	0	0	10m
rafay-system	replicaset.apps/controller-manager-v3-6bb696cc8b	1	1	1	6m14s
rafay-system	replicaset.apps/controller-manager-v3-7785d7b9d4	0	0	0	10m
rafay-system	replicaset.apps/rafay-connector-v3-6c8dcf8cf9	1	1	1	6m15s

NAMESPACE	NAME	COMPLETIONS	DURATION	AGE
kube-system	job.batch/helm-install-traefik-crd	1/1	5m16s	115m
kube-system	job.batch/helm-install-traefik	1/1	5m17s	115m

k3s-01:~ #

Conclusion

The combination of SUSE Linux Enterprise Micro, the lightweight Kubernetes system K3s, and the Cisco UCS C220, C240, and C240 SD servers can run modern cloud-native applications developed for Kubernetes in a single server deployed in a short-depth network cabinet.

For more information

For additional information, see the following resources:

- <https://suse.com/products/micro>
- <https://k3s.io/>
- https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/c240sdm5/install/c240sdm5.html
- <https://rafay.co/>

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)