



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

---

# Contents

Purpose of this document .....	3
Introduction .....	3
Solution overview .....	6
Prerequisites .....	6
Server operations with Cisco Intersight platform .....	6
Configure Cisco UCS C220, C240, and C240 SD through the Cisco Intersight platform .....	9
Install SLE Micro .....	36
Install K3s .....	49
K3s integration into the workload management tool .....	53
Conclusion .....	60
For more information .....	60

---

## 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)

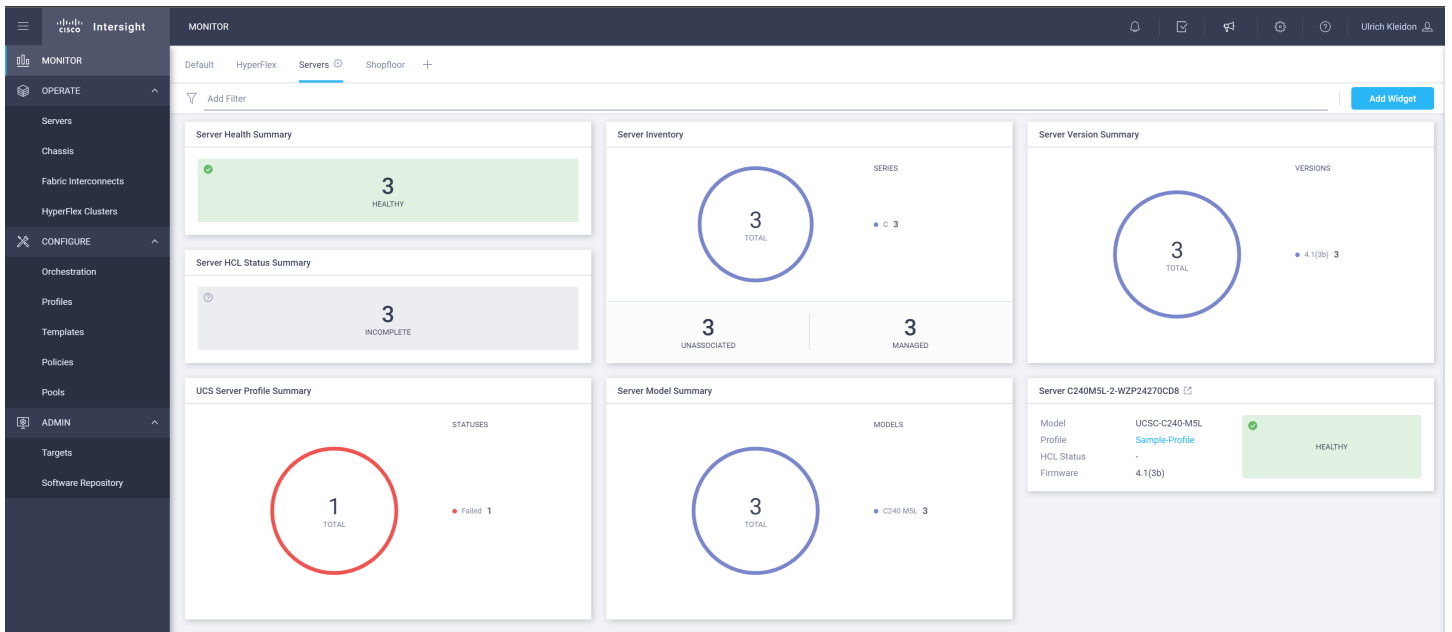
- Cisco® 12-Gbps SAS RAID modular controller and Cisco 12-Gbps SAS host bus adapter (HBA) controller
- 2 Flexible Flash (FlexFlash) Secure Digital (SD) card slots or 2 modular M.2 SATA slots
- 10-Gbps embedded Intel x550 10GBASE-T LAN-on-motherboard (LOM) port
- 1 modular LOM (mLOM) slot
- 6 PCIe Generation 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 the Cisco Intersight platform

The Cisco Intersight™ platform (<https://intersight.com>) is an API-driven, cloud-based system management tool (Figure 1). It is designed to help organizations implement their IT management and operations with a higher level of automation, simplicity, and operational efficiency. It is a new generation of global management tool for Cisco UCS and Cisco HyperFlex™ systems and provides a holistic and unified approach to managing customers' distributed and virtualized environments. The Cisco Intersight platform simplifies the installation, monitoring, troubleshooting, upgrading, and support of your infrastructure through the following benefits:

- **Cloud-based management:** The capability to manage Cisco UCS and Cisco HyperFlex systems from the cloud enables customers to more quickly and simply scale and manage their infrastructure whether in data centers or remote and branch-office locations.
- **Automation:** The unified API in Cisco UCS and Cisco HyperFlex systems enables policy-based configuration and management of the infrastructure and makes the Cisco Intersight platform itself and the devices connected to it fully programmable and DevOps friendly.
- **Analytics and telemetry:** The Cisco Intersight platform monitors the health and relationships of all physical and virtual infrastructure components. It also collects telemetry and configuration information to develop the intelligence of the platform in accordance with Cisco information security requirements.
- **Connected Cisco Technical Assistance Center (TAC):** Solid integration with the Cisco TAC enables more efficient and proactive technical support. The Cisco Intersight platform provides enhanced operations automation by expediting file transmission to accelerate troubleshooting.
- **Recommendation engine:** Driven by analytics and machine learning, the Cisco Intersight recommendation engine provides actionable intelligence for IT operations management through the daily-increasing knowledge base and practical insights learned in the entire system.
- **Management as a service:** The Cisco Intersight platform provides management as a service and is designed to be infinitely scalable and easy to implement. It relieves users of the burden of maintaining systems management software and hardware.



**Figure 1.**  
Cisco Intersight platform

## 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).

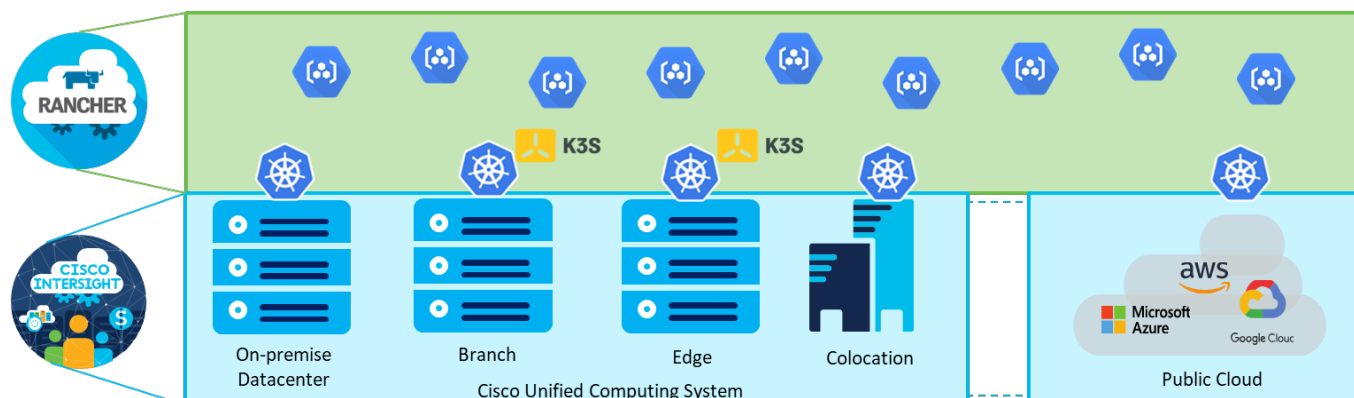
K3s bundles these Kubernetes components:

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

## Solution overview

The Cisco Intersight platform is a cloud-based service for managing Cisco UCS servers located at different locations from a single point. With policy-based architecture and infrastructure management based on profiles, organizations can easily define a server profile for a Cisco UCS server, such as the Cisco UCS C240, running SLE Micro and K3s and deploy it at any location. SUSE Rancher is used to manage the Kubernetes installations in the data center, in branch offices, the edge, or in the public cloud.

Using the sample landscape shown in Figure 2, this document demonstrates the installation of a single server.



**Figure 2.**  
Solution overview with Cisco Intersight platform 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 server racked and cabled
- Domain Host Configuration Protocol (DHCP) server to provide an IP address to the Cisco Integrated Management Controller (IMC)
- Monitor, keyboard, and mouse for initial IMC configuration

## Server operations with Cisco Intersight platform

To monitor and operate a Cisco UCS server from the Cisco Intersight platform, the first step is to claim the device. The following procedure provides the steps for claiming the Cisco UCS C240 server manually in the Cisco Intersight platform.

1. Log on to the Cisco Intersight platform and navigate to Admin > Targets.
2. In the top-right corner of the window, click Claim a New Target.
3. In the next window, choose Compute / Fabric > Cisco UCS Server (Standalone). Then click Start.

## Select Target Type

- In a second tab of the web browser, log on to the Cisco IMC portal of the Cisco UCS C240 SD and navigate to Admin > Device Connector.

- Back in the Cisco Intersight platform, enter the device ID and claim code from the server and click Claim.

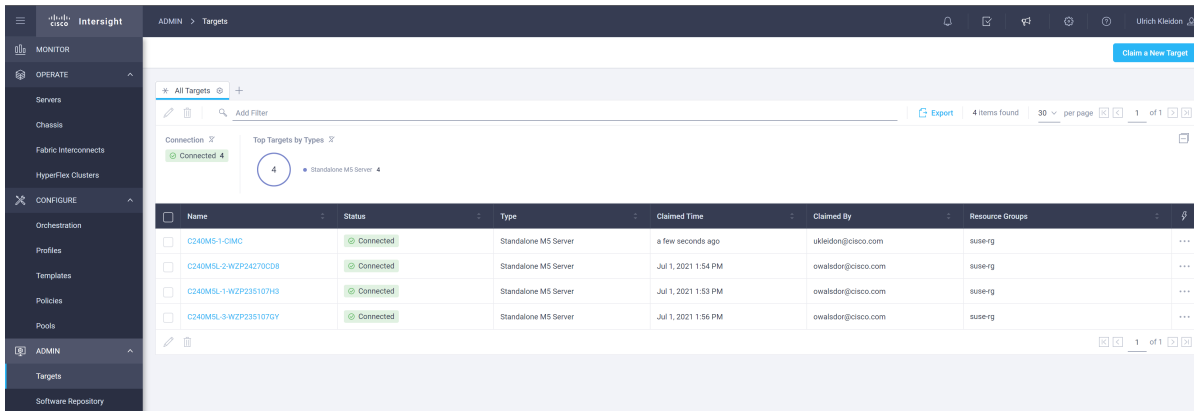


### Claim Cisco UCS Server (Standalone) Target

To claim your target, provide the Device ID, Claim Code and select the appropriate Resource Groups.

Name	Usage	Description
suse-rg	Suse	

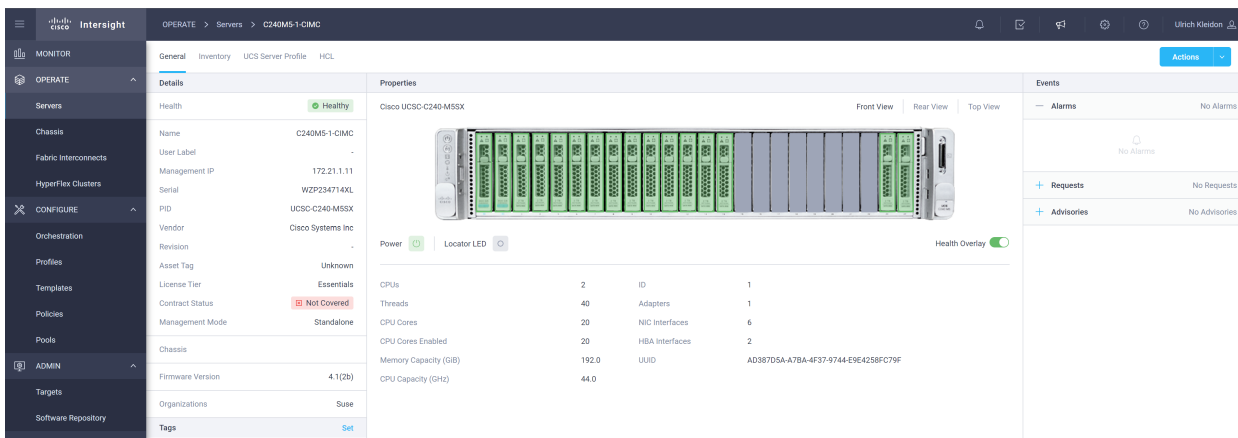
The server is now listed in the Cisco Intersight platform under Targets and under Servers.



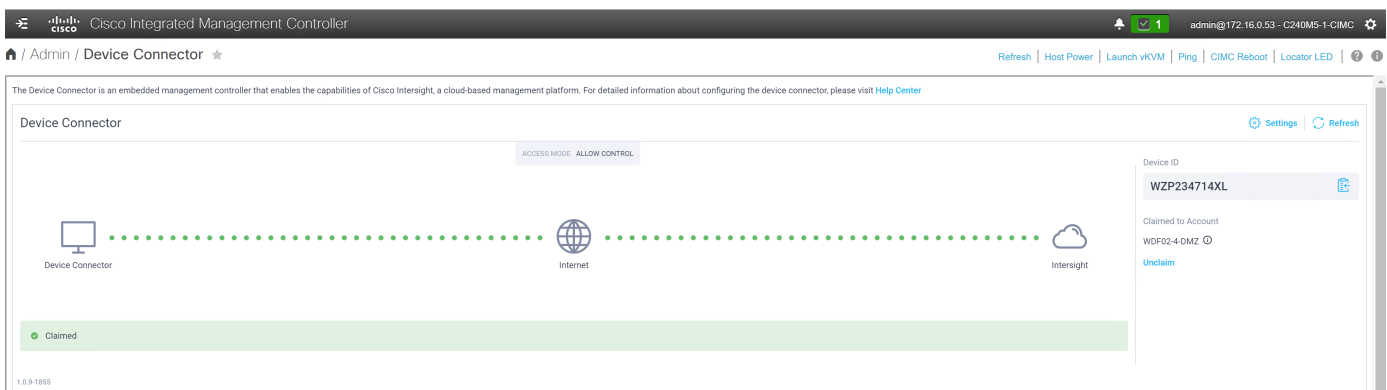
6. Navigate to Operate > Servers and choose the name of the new server to see the details and actions available for this system.

The available actions are based on the Cisco Intersight license level available for this server and the privileges of the user account.

See [https://intersight.com/help/saas/getting\\_started/licensing\\_requirements#intersight\\_licensing](https://intersight.com/help/saas/getting_started/licensing_requirements#intersight_licensing) for an overview of the functions available with the different license tiers.

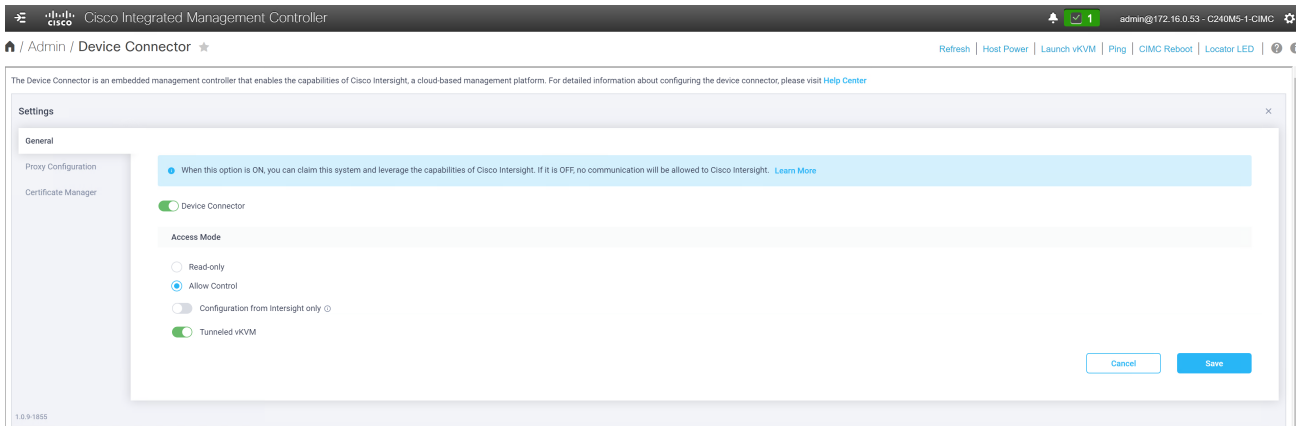


7. In the C240 IMC, click Refresh. The system must be shown as Claimed. Click Settings.



8. Enable Tunnelled vKVM and click Save. Tunnelled vKVM allows the Cisco Intersight platform to open the virtual keyboard, video, and mouse (vKVM) window in the event that the client has no direct network access to the server on the local LAN or virtual private network (VPN).





## Configure Cisco UCS C220, C240, and C240 SD through the Cisco Intersight platform

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

All configuration steps are performed in the Cisco Intersight portal.

### Perform initial setup

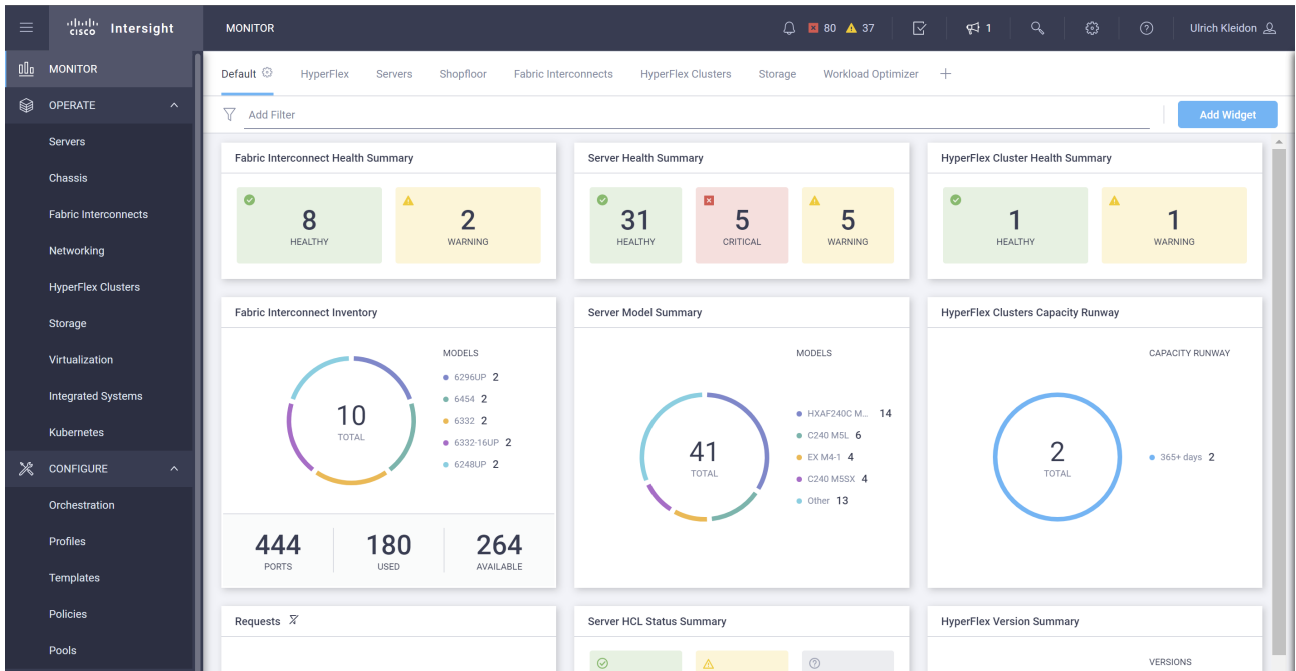
Hardware installation details and the initial server setup process are documented in the server's installation documentation. For the Cisco UCS C240 SD M5, the document can be found here:

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/c/hw/c240sdm5/install/c240sdm5/C240M5\\_chapter\\_01.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/c240sdm5/install/c240sdm5/C240M5_chapter_01.html).

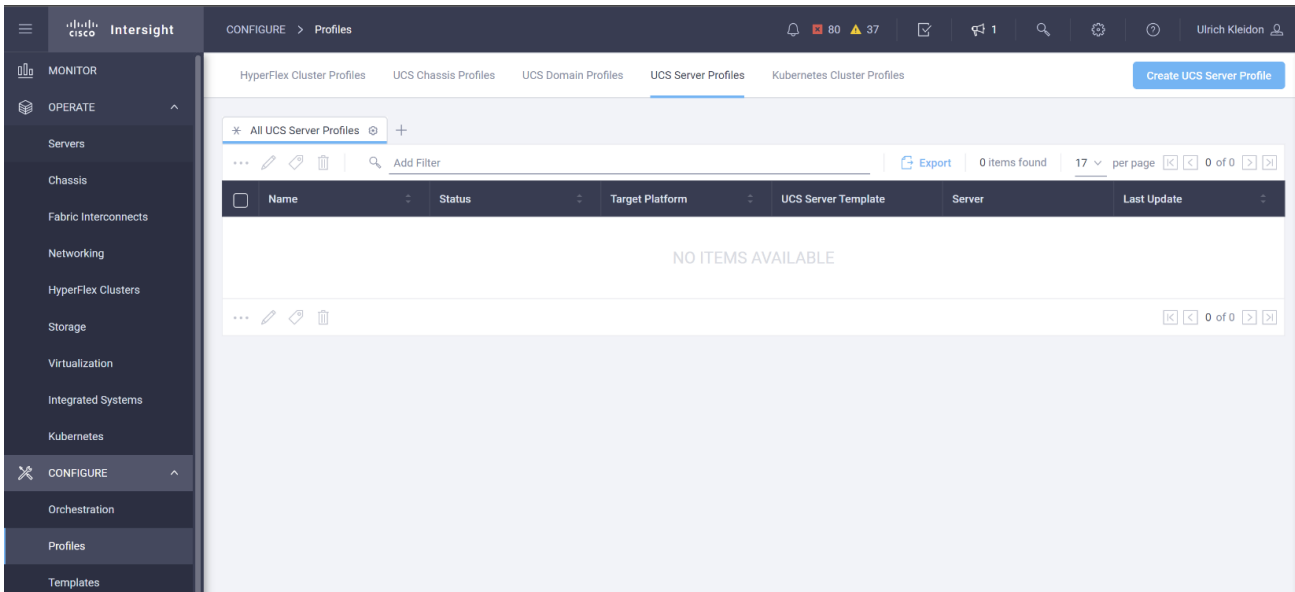
### Configure the storage and network

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

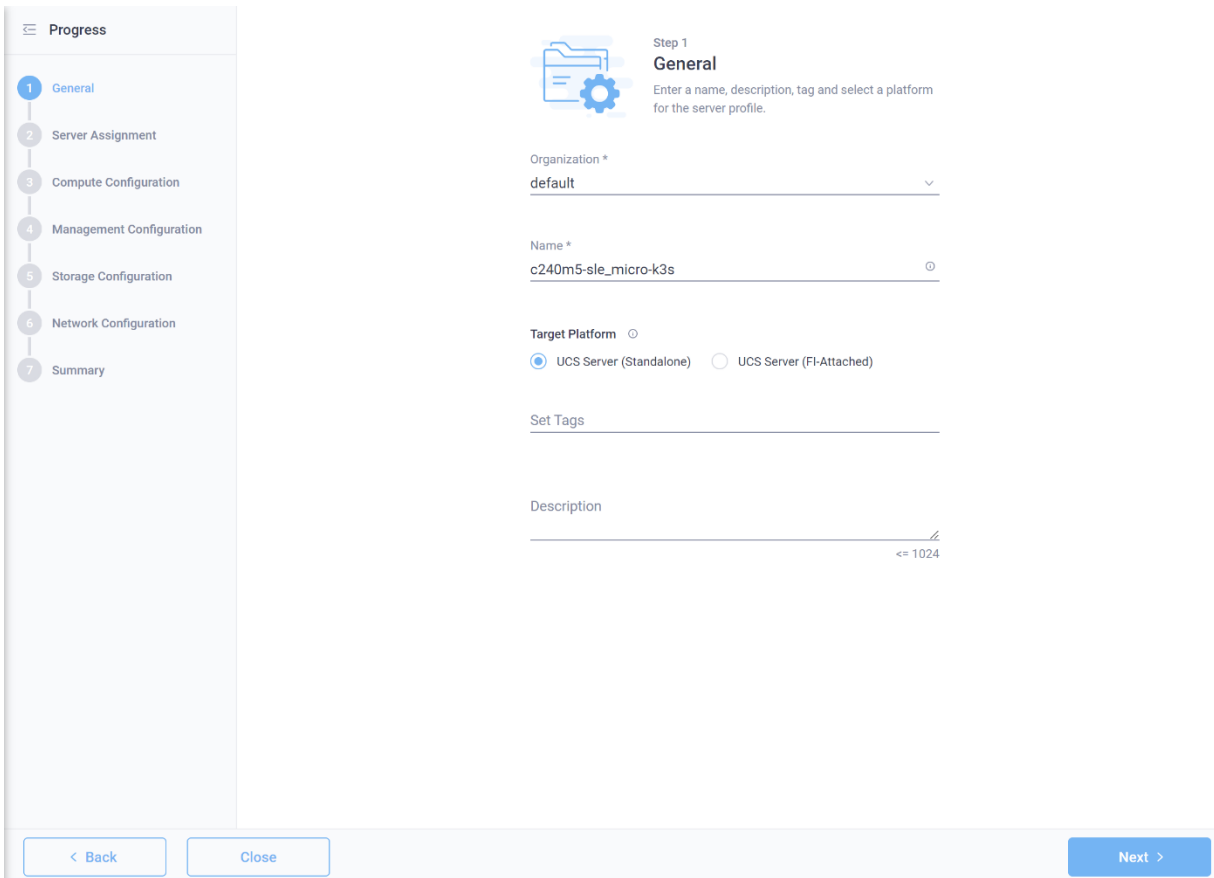
1. Open the Cisco Intersight portal in a web browser and log in as a user with admin permissions.



2. In the left menu, navigate to Configure > Profiles and click the UCS Server Profiles tab. Click Create UCS Server Profile.




3. Enter a name for the new server profile and click Next.



4. Select the server that will be assigned to the server profile and used to install SLE Micro and K3s.

**Progress**

- General
- Server Assignment**
- Compute Configuration
- Management Configuration
- Storage Configuration
- Network Configuration
- Summary



**Step 2**  
**Server Assignment**  
Choose to assign a server now, from a resource pool, or later.

Assign Now
Assign Server from a Resource Pool
Assign Later

Click the appropriate button to assign a server now, from a resource pool, or later. If you choose to assign a server now, select the server, click Next, and select and attach policies to the server profile.

Name	User Label	Health	Model	UCS Domain	Serial Num...	Managemen...
<input type="radio"/> C240M5L-1-WZ...		Healthy	UCSC-C240-M5L		WZP235107H3	172.16.1.80
<input type="radio"/> C240M5L-2-WZ...		Healthy	UCSC-C240-M5L		WZP24270CD8	172.16.1.81
<input type="radio"/> C240M5L-3-WZ...		Healthy	UCSC-C240-M5L		WZP235107GY	172.16.1.82
<input type="radio"/> C240M5-2-CIMC		Healthy	UCSC-C240-M5...		WZP234507TF	172.16.1.84
<input checked="" type="radio"/> C240M5-1-CIMC		Healthy	UCSC-C240-M5...		WZP234714XL	172.16.1.83
<input type="radio"/> C240M5-3-CIMC		Healthy	UCSC-C240-M5...		WZP234507UH	172.16.1.85


Selected 1 of 6    Show Selected    Unselect All    1 of 1

< Back
Close
Next >

5. Click Select Policy to the right of Boot Order.

**Progress**

- General
- Server Assignment
- Compute Configuration**
- Management Configuration
- Storage Configuration
- Network Configuration
- Summary



**Step 3**  
**Compute Configuration**  
Create or select existing Compute policies that you want to associate with this profile.

BIOS	
Boot Order	Select Policy
Persistent Memory	
Virtual Media	

6. Click Create New (or select a profile that fits your needs).

**Progress**

- General
- Server Assignment
- Compute Configuration**
- Management Configuration
- Storage Configuration
- Network Configuration
- Summary



**Step 3**  
**Compute Configuration**  
Create or select existing Compute policies that you want to associate with this profile.

BIOS	
Boot Order	
Persistent Memory	
Virtual Media	

Select Boot Order

Policies 0 Create New

Search

7. Enter the name of the new boot profile and click Next.

CONFIGURE > Create UCS Server Profile > Create Boot Order Policy

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

Name \*  
C240M5-Boot

Set Tags

Description  
≤ 1024

Cancel Next >

8. Select Unified Extensible Firmware Image (UEFI) under Configured Boot Mode and click Add Boot Device.

CONFIGURE > Create UCS Server Profile > Create Boot Order Policy

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)


Configured Boot Mode

Legacy  Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot

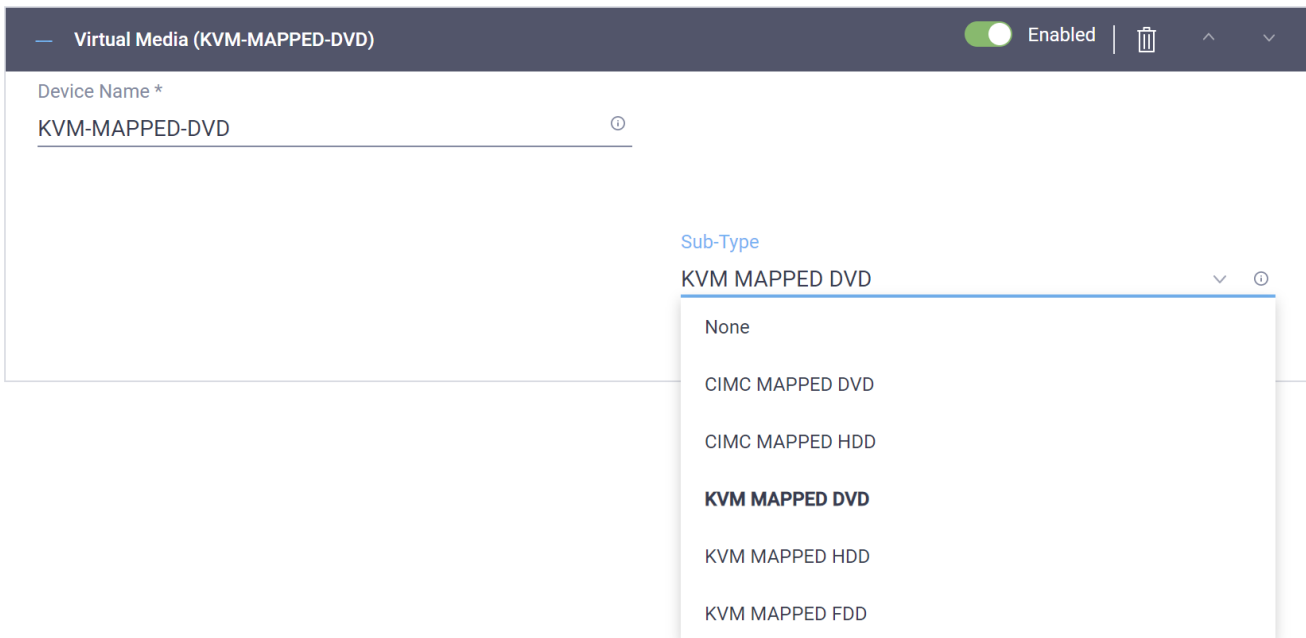
Add Boot Device

9. Select UEFI Shell from the drop-down menu, enter **UEFI-SHELL** as the name, and click Add.



10. Click Add Boot Device again and select Virtual Media.


11. Enter KVM-MAPPED-DVD as the device name and select KVM MAPPED DVD as the subtype from the drop-down list.  
Click Add.



12. The KVM-MAPPED-DVD should be the first option and the UEFI-Shell the second option in the list. Click Create.

**Progress**

1 General  
2 Policy Details



Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

**Configured Boot Mode** ⊙

Legacy  Unified Extensible Firmware Interface (UEFI)

Enable Secure Boot ⊙


▾

+ Virtual Media (KVM-MAPPED-DVD)	<input checked="" type="checkbox"/> Enabled	🗑️ ^ ▾
+ UEFI Shell (UEFI-SHELL)	<input checked="" type="checkbox"/> Enabled	🗑️ ^ ▾

13. Click Next.

**Progress**

1 General  
2 Server Assignment  
3 Compute Configuration  
4 Management Configuration  
5 Storage Configuration  
6 Network Configuration  
7 Summary



Step 3  
**Compute Configuration**  
Create or select existing Compute policies that you want to associate with this profile.

BIOS	📄
Boot Order	📄 C240M5-Boot
Persistent Memory	📄
Virtual Media	📄

14. Click Next.

The screenshot shows a configuration wizard interface. On the left is a 'Progress' sidebar with seven steps: 1. General, 2. Server Assignment, 3. Compute Configuration, 4. Management Configuration (highlighted in blue), 5. Storage Configuration, 6. Network Configuration, and 7. Summary. The main area is titled 'Step 4 Management Configuration' with a sub-instruction: 'Create or select existing Management policies that you want to associate with this profile.' Below this is a list of management policies, each with a document icon to its right: Device Connector, IPMI Over LAN, LDAP, Local User, Network Connectivity, NTP, Serial Over LAN, SMTP, SNMP, SSH, Syslog, and Virtual KVM. At the bottom of the wizard are three buttons: '< Back', 'Close', and 'Next >'.

15. Move your mouse to the Storage entry and click Select Policy and Create New.

The screenshot shows the next step in the configuration wizard, 'Step 5 Storage Configuration', with the sub-instruction: 'Create or select existing Storage policies that you want to associate with this profile.' The 'Progress' sidebar on the left now highlights '5. Storage Configuration'. The main area shows a list of storage policies: 'SD Card' and 'Storage'. The 'Storage' entry is highlighted with a light blue background and has a 'Select Policy' button with a document icon to its right.

16. Enter a name for the new storage policy and click Next.

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

Name \*  
C240M5-Sle\_Micro-Storage

Set Tags

Description  
≤ 1024

Cancel Next >

17. If the drive state in this server is unknown, you can enable “Use JBOD drives for Virtual Drive creation.”  
As the drive state of this server is known, you do not need to enable this option

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

**General Configuration**

Use JBOD drives for Virtual Drive creation

Unused Disks State  
No Change

M.2 RAID Configuration  Enable

MRAID/RAID Controller Configuration  Enable

MRAID/RAID Single Drive RAID0 Configuration  Enable

18. Enable MRAID/RAID Controller Configuration and click Add Drive Group.



The screenshot shows the 'MRAID/RAID Controller Configuration' page. On the left, a sidebar indicates the current step is 'Policy Details' (2) under the 'General' section (1). The main content area has a title bar with an 'Enable' toggle. Below this, there are two sections for adding drive groups and virtual drives. Each section contains a table with columns for drive group name, RAID level, number of spans, dedicated hot spares, and drive array spans. Both tables currently show '0 items found' and 'NO ITEMS AVAILABLE'.

19. Enter a drive group Name, select RAID1 as the RAID level, and enter the drive numbers of the two drives that will form the RAID1 group for boot. Then click Add.

The 'Add Drive Group' dialog box is shown. It has a title bar with a close button. The 'Configuration' section contains two fields: 'Drive Group Name \*' with the value 'C240M5-Boot-DG' and 'RAID Level' with a dropdown menu set to 'RAID1'. The 'Drive Selection' section contains two fields: 'Drive Array Span 0' with the value '9,10' and 'Dedicated Hot Spares' which is empty. At the bottom, there are 'Cancel' and 'Add' buttons.

You can find the drive numbers by navigating to Operate > Server > [name of your server] > Inventory > Storage Controller > Controller MRAID > Physical Drives.

OPERATE > Servers > C240M5-1-CIMC

General Inventory UCS Server Profile HCL Statistics

Expand All

Motherboard  
Boot  
CIMC  
CPUs  
Memory  
Network Adapters  
PCIe Devices  
Storage Controllers  
Controller MRAID  
TPM

Controller MRAID

General Physical Drives Virtual Drives

Name	Disk Firmware...	Size (MiB)	Model	Serial	Vendor	Protocol	Type	Drive State	
<input type="checkbox"/>	Disk 1	CN05	952720 MB	ST1000NX04...	W472XEF100...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 2	CN05	952720 MB	ST1000NX04...	W472XD180...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 3	CN05	952720 MB	ST1000NX04...	W472XE0P0...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 4	CN05	952720 MB	ST1000NX04...	W472XDX20...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 5	CN05	952720 MB	ST1000NX04...	W472XEA30...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 6	CN05	952720 MB	ST1000NX04...	W472XDC40...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 7	CN05	952720 MB	ST1000NX04...	W472XE5W0...	SEAGATE	SAS	HDD	Unconfigured...
<input type="checkbox"/>	Disk 8	CN05	952720 MB	ST1000NX04...	W472XH9B0...	SEAGATE	SAS	HDD	Unconfigured...
<input checked="" type="checkbox"/>	Disk 9	0103	914573 MB	PX05SRB096	7830A063TS...	TOSHIBA	SAS	SSD	Unconfigured...
<input checked="" type="checkbox"/>	Disk 10	0103	914573 MB	PX05SRB096	7830A05WT...	TOSHIBA	SAS	SSD	Unconfigured...
<input type="checkbox"/>	Disk 11	0103	914573 MB	PX05SRB096	7810A02CTS...	TOSHIBA	SAS	SSD	Unconfigured...

20. Click Add Virtual Drive.

Progress

1 General  
2 Policy Details

MRAID/RAID Controller Configuration  Enable

Global Hot Spares

Add Drive Group

Drive Group Name	RAID Level	Number of Spans	Dedicated Hot Spares	Drive Array Spans
<input type="checkbox"/>	C240M5-Boot-DG	RAID1		{9,10}

Add Virtual Drive

Virtual Drive Name	Drive Group	Size (MiB)	Expand to Available	Set as Boot Drive
NO ITEMS AVAILABLE				

21. Enter a name for the new virtual drive, select the created drive group, and enable Expand to Available to use the complete capacity of the RAID1 drive group. Click Add.

### Add Virtual Drive

Drive Groups \*  
 C240M5-Boot-DG    Number of Copies  
 0 - 10

**Virtual Drive Configuration**

Virtual Drive Name \*  
 boot-vd

Expand to Available

Set as Boot Drive

Strip Size: 64KiB    Access Policy: Default    Read Policy: Default

Write Policy: Default    Disk Cache: Default

22. Click Create.

Progress

- 1 General
- 2 Policy Details

M.2 RAID Configuration  Enable

MRAID/RAID Controller Configuration  Enable

Global Hot Spares

Drive Group Name	RAID Level	Number of Spans	Dedicated Hot Spares	Drive Array Spans
<input type="checkbox"/> C240M5-Boot-DG	RAID1			{ 9,10 }


Virtual Drive Name	Drive Group	Size (MiB)	Expand to Available	Set as Boot Drive	
<input type="checkbox"/> boot-vd	C240M5-Boot-DG	-	Yes	Yes	...

MRAID/RAID Single Drive RAID0 Configuration  Enable



23. Click Next.

**Progress**

- 1 General
- 2 Server Assignment
- 3 Compute Configuration
- 4 Management Configuration
- 5 Storage Configuration**
- 6 Network Configuration
- 7 Summary



Step 5  
**Storage Configuration**  
Create or select existing Storage policies that you want to associate with this profile.


SD Card	
Storage	<input checked="" type="radio"/> C240M5-Sle_Micro-Storage 

[< Back](#)      [Close](#)      [Next >](#)





24. Click Select Policy to the right of Adapter Configuration and then click Create New.

**Progress**

- 1 General
- 2 Server Assignment
- 3 Compute Configuration
- 4 Management Configuration
- 5 Storage Configuration
- 6 Network Configuration**
- 7 Summary



Step 6  
**Network Configuration**  
Create or select existing Network Configuration policies that you want to associate with this profile.

Adapter Configuration	<a href="#">Select Policy</a> 
LAN Connectivity	
SAN Connectivity	
Auto Placement Configuration for vNICs & vHBAs 	

25. Enter the name for the new adapter configuration policy and click Next.

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

Name \*  
C240M5-Sle\_Micro-VIC

Set Tags

Description  
≤ 1024

Cancel

Next >

26. Click Add VIC Adapter Configuration.

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

This policy is applicable only for UCS Servers (Standalone)

Adapter Configurations

Add VIC Adapter Configuration

<input type="checkbox"/>	PCI Slot	LLDP	FIP	Port Channel
NO ITEMS AVAILABLE				

27. Enter MLOM as the PCI slot, enable the port channel settings, and click Add.



### Add VIC Adapter Configuration

PCI Slot \*  
MLOM

**Ethernet Settings**

Enable LLDP

**Fibre Channel Settings**

Enable FIP

**Port Channel Settings**

Enable Port Channel

**DCE Interface Settings**

DCE Interface 1	DCE Interface 2
FEC Mode * cl91	FEC Mode * cl91
DCE Interface 3	DCE Interface 4

[Cancel](#) [Add](#)

## 28. Click Create.

**Progress**

- 1 General
- 2 Policy Details

**Step 2 Policy Details**  
Add policy details

This policy is applicable only for UCS Servers (Standalone)

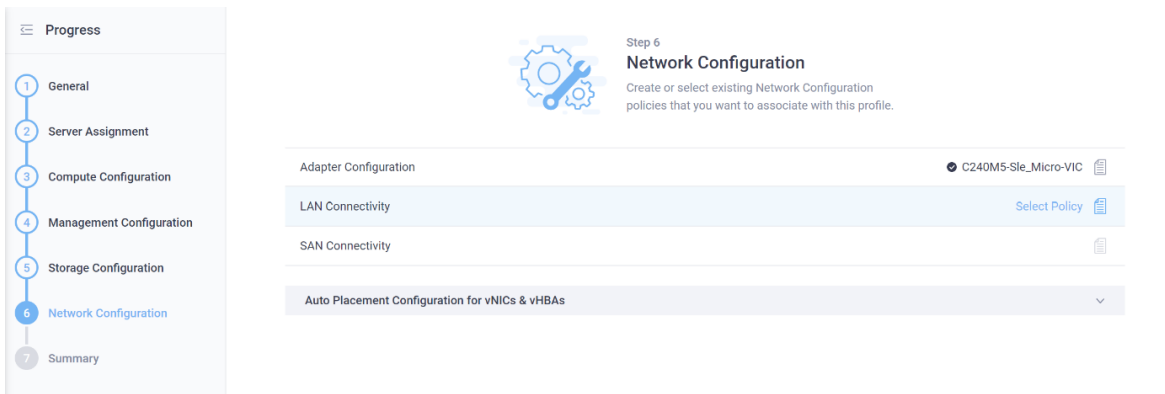
**Adapter Configurations**

[Add VIC Adapter Configuration](#)

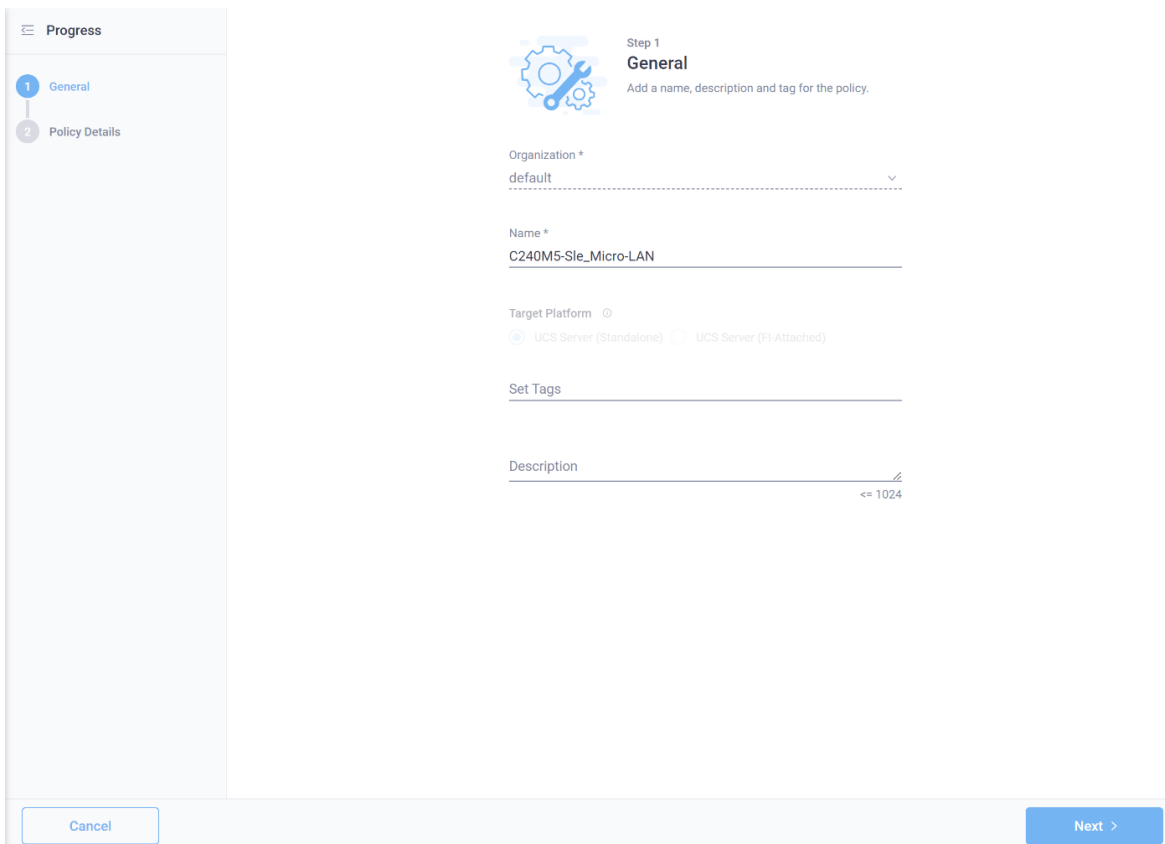
PCI Slot	LLDP	FIP	Port Channel	
<input type="checkbox"/> MLOM	Enabled	Enabled	Enabled	...

[< Back](#) [Cancel](#) [Create](#)

## 29. Click Select Policy to the right of LAN Connectivity and then click Create New.



30. Enter the name for the new policy used for this type of server and click Next.



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 virtual network interface cards (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. Table 1 shows the required information for creating the vNICs.

**Table 1.** vNIC information

vNIC name	Uplink port	PCI order	LAN	VLAN ID
Eth0	0	0	Admin	211
Eht1	1	1	Admin	211
Eth2	0	2	Access	210
Eth3	1	3	Access	210
Eth4	0	4	Storage	212
Eth5	1	5	Storage	212

31. Click Add vNIC.

32. Enter a name for this vNIC and select 0 as the uplink port.



33. Click Select Policy under Ethernet Network and click Create New.

Ethernet Network \*

Select Policy

Ethernet QoS \*

Select Policy

Ethernet Adapter \*

Select Policy

Connection

Disabled usNIC VMQ

Cancel Add

34. Enter a name for the administration LAN policy and click Next.

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

Name \*  
DC-Admin-LAN

Set Tags

Description  
<= 1024

Cancel Next >

35. Keep the VLAN mode at Access and enter the VLAN ID for the administration LAN. Here, 211 is used. Click Create.

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

This policy is applicable only for UCS Servers (Standalone)

**VLAN Settings**

VLAN Mode: Access

Default VLAN: 211

0 - 4094

< Back    Cancel    Create

36. Click Select Policy under QoS Policy and click Create New.

37. In the new view, enter a name for the new quality-of-service (QoS) policy and click Next.

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

Name \*  
Best-Effort-1500

Set Tags

Description  
<= 1024

Cancel    Next >

38. Keep MTU Bytes at 1500 and change the class of service only if advised by your network team, otherwise, use 0. Click Create.

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

QoS Settings

MTU, Bytes  
1500  1500 - 9000

Rate Limit, Mbps  
0  0 - 100000

Class of Service  
0  0 - 6

Enable Trust Host CoS

< Back Cancel Create

39. Click Select Policy under Adapter Policy and click Create New.

40. In the new view, enter a name for the new QoS policy and click Select Default Configuration.

Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

Organization \*  
default

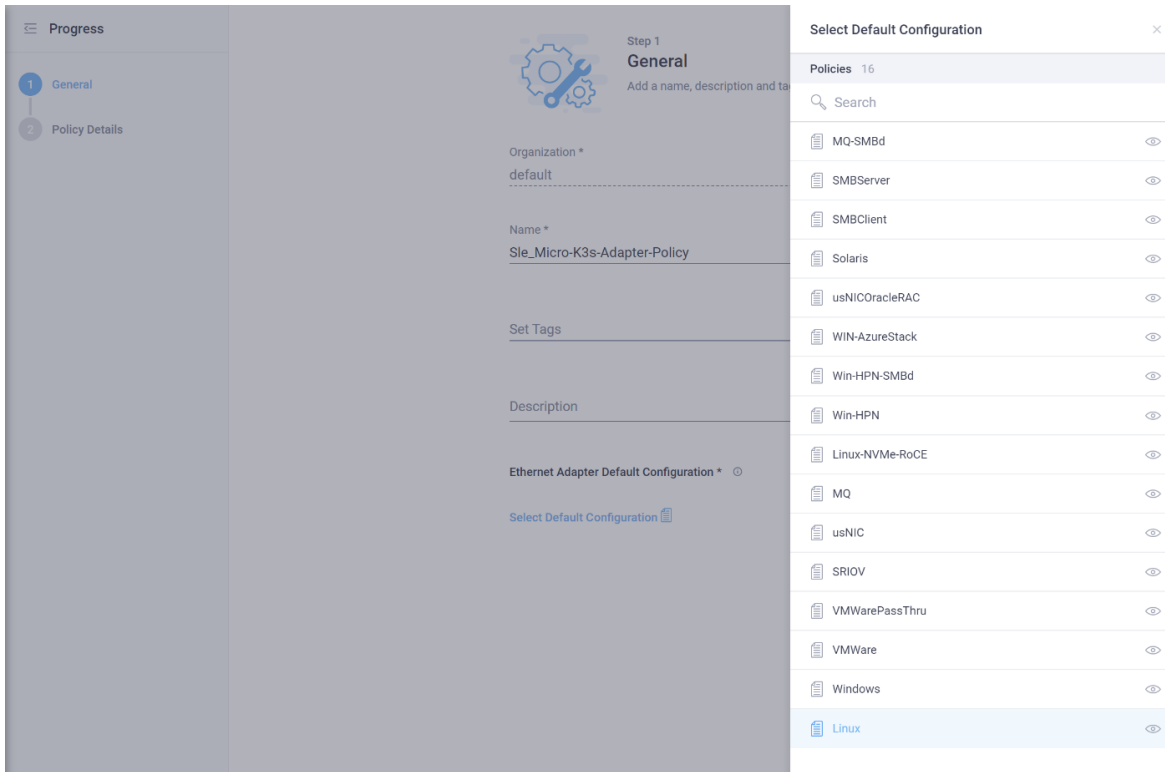
Name \*  
Sle\_Micro-K3s-Adapter-Policy

Set Tags

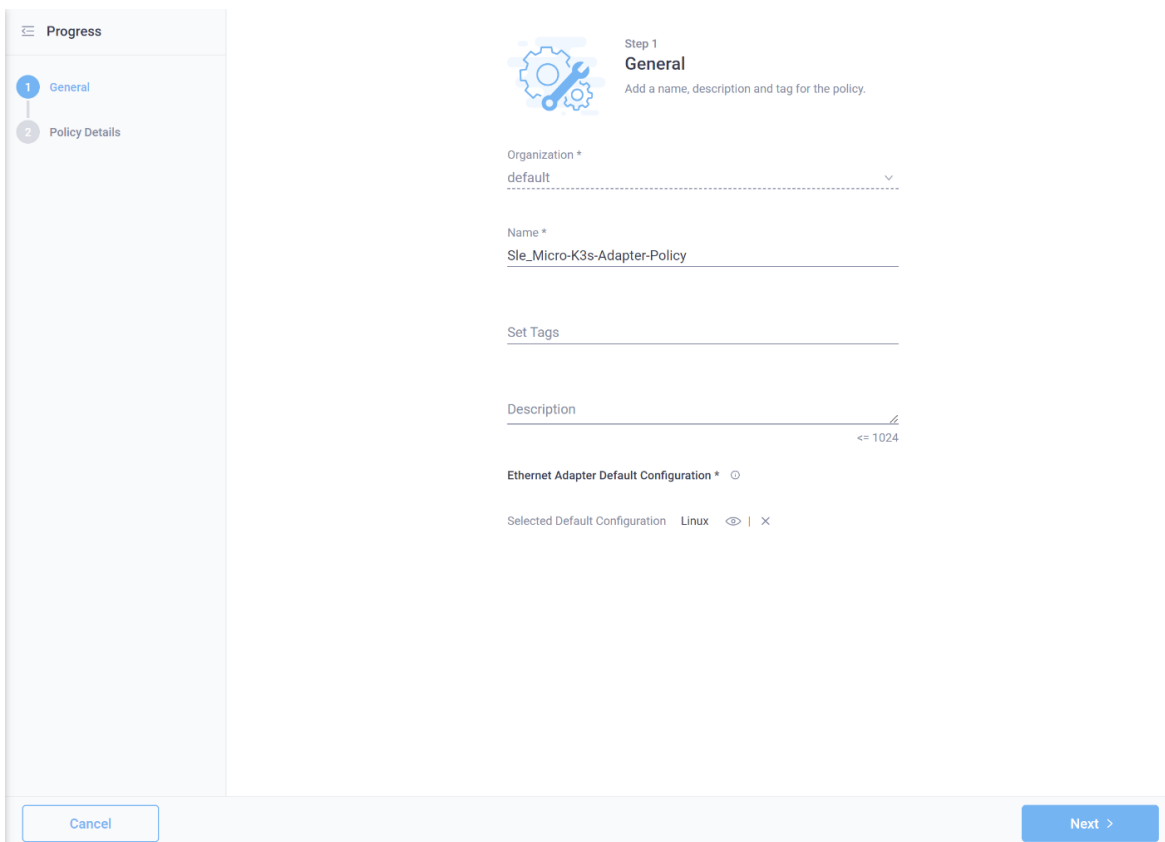
Description  
<= 1024

Ethernet Adapter Default Configuration \*  
Select Default Configuration

41. Click Linux.



42. Click Next.



43. Click Create.

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

All Platforms | UCS Server (Standalone) | UCS Server (FI-Attached)

Enable Virtual Extensible LAN

Enable Network Virtualization using Generic Routing Encapsulation

Enable Accelerated Receive Flow Steering

Enable Precision Time Protocol

Enable Advanced Filter

Enable Interrupt Scaling

**RoCE Settings**

Enable RDMA over Converged Ethernet

**Interrupt Settings**

Interrupts: 4 (1 - 1024)

Interrupt Mode: MSIx

Interrupt Timer, us: 125 (0 - 65535)

Interrupt Coalescing Type: Min

**Receive**

Receive Queue Count: 1 (1 - 1000)

Receive Ring Size: 512 (64 - 16384)

**Transmit**

< Back | Cancel | Create

44. Click Add to create the vNIC.

**Ethernet Network \***

Selected Policy: Access-LAN

**Ethernet QoS \***

Selected Policy: Best-Effort-1500

**Ethernet Adapter \***

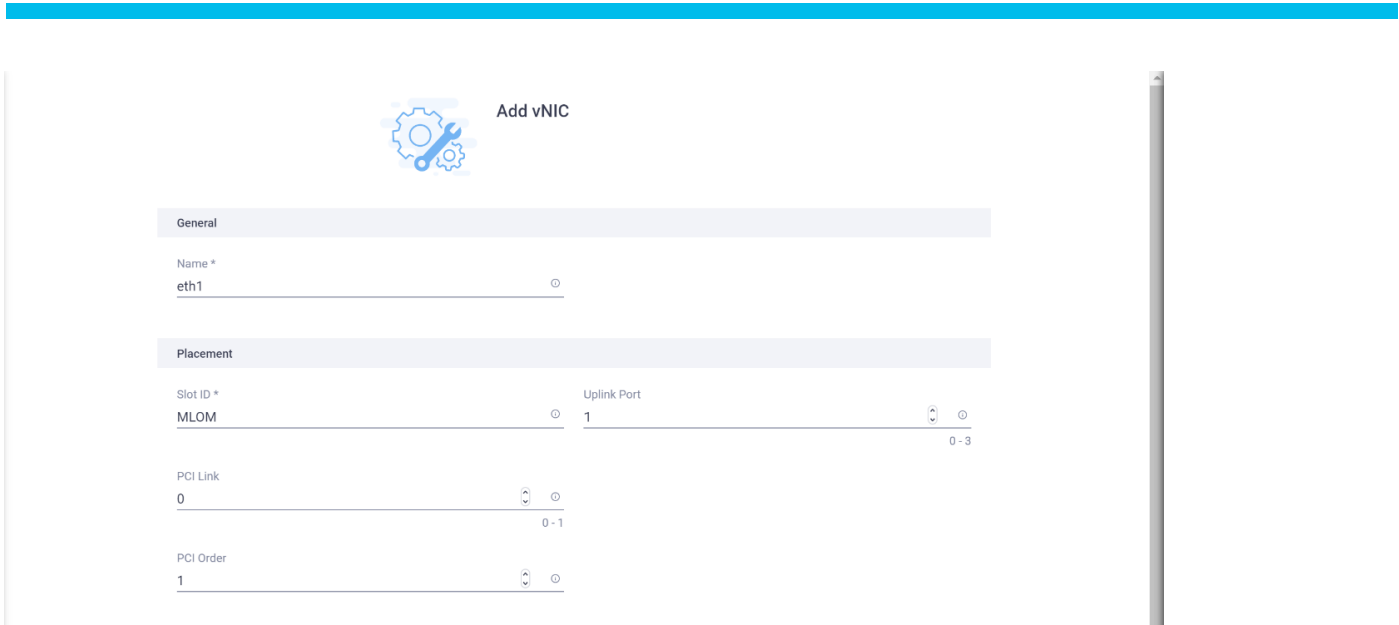
Selected Policy: Sle\_Micro-K3s-Adapter-Policy

**Connection**

Disabled | usNIC | VMQ

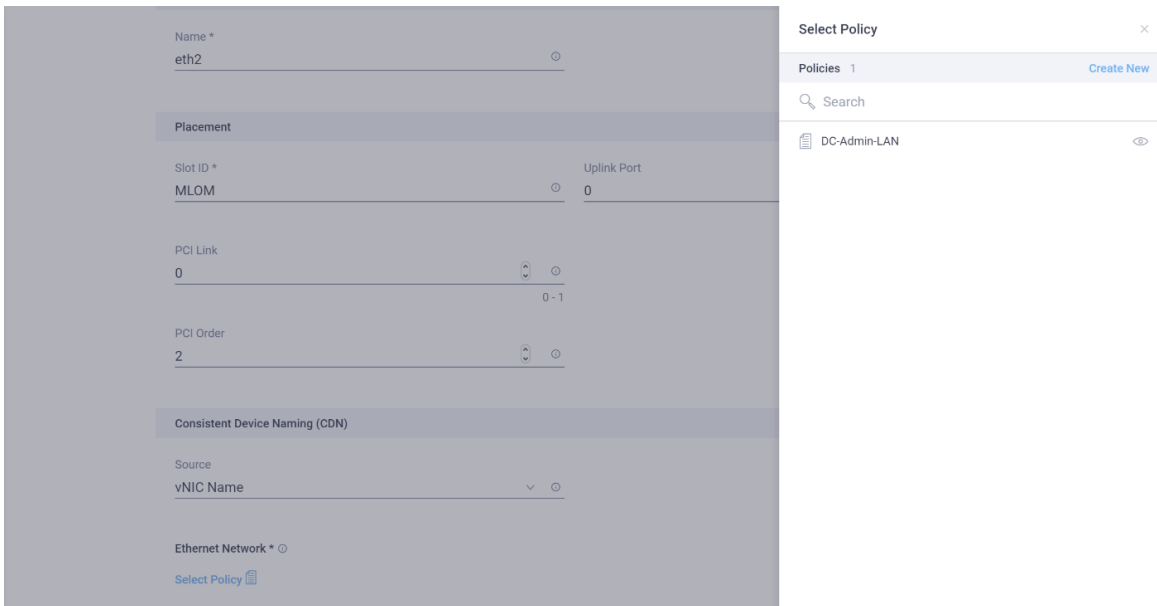
Cancel | Add

45. Create vNIC eth1 with uplink port 1 and PCI order 1 as shown in Table 1. Then select the same Ethernet, QoS, and Adapter Policy as for eth0.

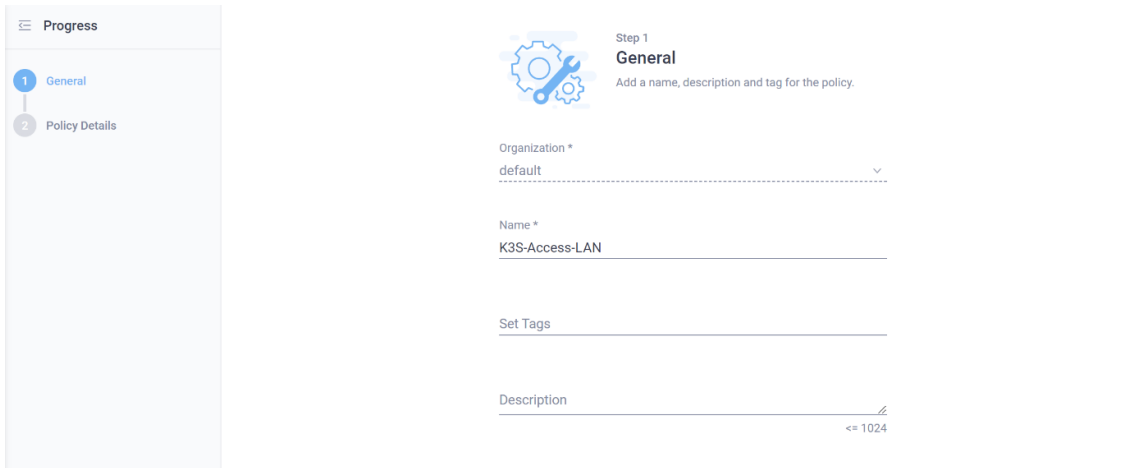


46. Click Add vNIC to create vNIC eth2 with uplink port 0 and PCI order 2.

47. Click Select Policy under Ethernet Network and click Create New.



48. Enter a name for the access LAN policy and click Next.



Progress

1 General

2 Policy Details

Step 1  
**General**  
Add a name, description and tag for the policy.

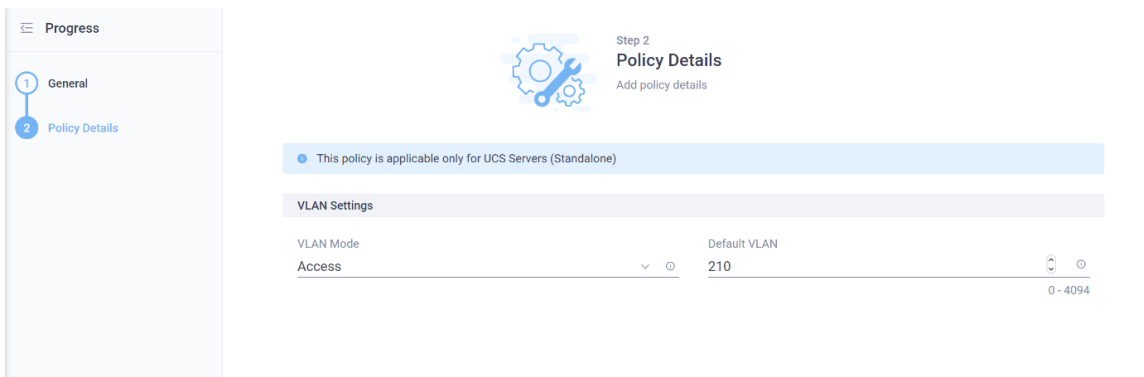
Organization \*  
default

Name \*  
K3S-Access-LAN

Set Tags

Description  
≤ 1024

49. Keep VLAN Mode as Access and enter the VLAN ID for the access network; here, 210 is used. Click Create.



Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

This policy is applicable only for UCS Servers (Standalone)

VLAN Settings

VLAN Mode  
Access

Default VLAN  
210  
0 - 4094

50. Use the same QoS and adapter policies as for eth0 and eth1 and click Add.

51. Create eth3 with the settings from Table 1 and the same Ethernet, QoS, and adapter policies as for eth2.
52. For eth4 and eth5, the Ethernet policy for the storage LAN is required. To create vNIC eth4, click Select Policy and then Create New.

53. Enter a name for the new storage LAN policy and click Next.



54. Keep VLAN Mode at Access and enter the VLAN ID for the storage LAN; here, 212 is used. Click Create.

Progress

1 General

2 Policy Details

Step 2  
**Policy Details**  
Add policy details

This policy is applicable only for UCS Servers (Standalone)

**VLAN Settings**

VLAN Mode: Access

Default VLAN: 212 (0 - 4094)

55. Use the same QoS and adapter policies as for the other vNICs and click Add.

**General**

Name \*  
eth4

**Placement**

Slot ID \*  
MLOM

Uplink Port  
0 (0 - 3)

PCI Link  
0 (0 - 1)

PCI Order  
4

**Consistent Device Naming (CDN)**

Source  
vNIC Name

**Ethernet Network \* 0**

Selected Policy Global-Storage-LAN

**Ethernet QoS \* 0**

Selected Policy Best-Effort-1500

**Ethernet Adapter \* 0**

Selected Policy Sle\_Micro-K3s-Adapter-Policy


Cancel Add

56. Use the settings from Table 1 and the policies from eth4 to create vNIC eth5.

57. The Policy Details screen 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. This list will be used later to validate the bond configuration at the operating system layer. Click Create.

**Progress**

- 1 General
- 2 Policy Details



**Step 2 Policy Details**  
Add policy details

At a minimum two vNICs are required named eth0 and eth1. Learn more at [Help Center](#)

**Add vNIC**


<input type="checkbox"/>	Name	Slot ID	Uplink Port	PCI Link	PCI Order	
<input type="checkbox"/>	eth0	MLOM	0	0	0	...
<input type="checkbox"/>	eth1	MLOM	1	0	1	...
<input type="checkbox"/>	eth2	MLOM	0	0	2	...
<input type="checkbox"/>	eth3	MLOM	1	0	3	...
<input type="checkbox"/>	eth4	MLOM	0	0	4	...
<input type="checkbox"/>	eth5	MLOM	1	0	5	...

< Back    Cancel    Create

58. Back in the Server Profile view, click Next.

**Progress**

- 1 General
- 2 Server Assignment
- 3 Compute Configuration
- 4 Management Configuration
- 5 Storage Configuration
- 6 Network Configuration
- 7 Summary



**Step 6 Network Configuration**  
Create or select existing Network Configuration policies that you want to associate with this profile.

Adapter Configuration ● C240M5-Sle\_Micro-VIC

LAN Connectivity ● C240M5-Sle\_Micro-LAN

SAN Connectivity 📄

Auto Placement Configuration for vNICs & vHBAs ▼

< Back    Close    Next >

59. Review the information on the Summary page and click Deploy.

**Step 7 Summary**  
Verify details of the profile and the policies, resolve errors and deploy.

General			
Organization	default	Status	Not Deployed
Name	c240m5-sle_micro-k3s	Management IP	172.16.1.83
Assigned Server	C240M5-1-CIMC		
Target Platform	UCS Server (Standalone)		

Below the table are tabs for: Compute Configuration (selected), Management Configuration, Storage Configuration, Network Configuration, and Errors/Warnings (0). A 'Boot Order' section shows 'C240M5-Boot'.

Buttons at the bottom: < Back, Close, Deploy

60. Browse to Operate > Servers and select the assigned server. Then click Actions and choose Launch vKVM.

**Operate > Servers > C240M5-1-CIMC**

General | Inventory | UCS Server Profile | HCL | Statistics

**Details**

Health: ● Healthy

Name: C240M5-1-CIMC

User Label: -

Management IP: 172.16.1.83

Serial: W2P234714XL

PID: UCSC-C240-M5SX

Vendor: Cisco Systems Inc

Revision: -

Asset Tag: Unknown

License Tier: Premier

Management Mode: Standalone

**Chassis**

Profile: c240m5-sle\_micro-k3s

Profile Status: ● OK

Firmware Version: 4.1(2b)

**Properties**

Cisco UCSC-C240-M5SX

Power: ● On | Locator LED: ● Off | Health Overlay: ● On

CPU	2	ID	1
Threads	40	Adapters	1
CPU Cores	20	NIC Interfaces	6
CPU Cores Enabled	20	HBA Interfaces	2
Memory Capacity	192.0	UUID	AD387D5A-A7BA-4F37-9744-E9E4258FC79F
CPU Capacity (GHz)	44.0		

**Actions**

- Power
- System
- Import Server Profile
- Install Operating System
- Upgrade Firmware
- Launch IMC
- Launch vKVM
- Launch Tunneled vKVM
- Open TAC Case
- Set License Tier

61. In the new window, take the necessary steps to continue with an untrusted certificate and close it at the end.

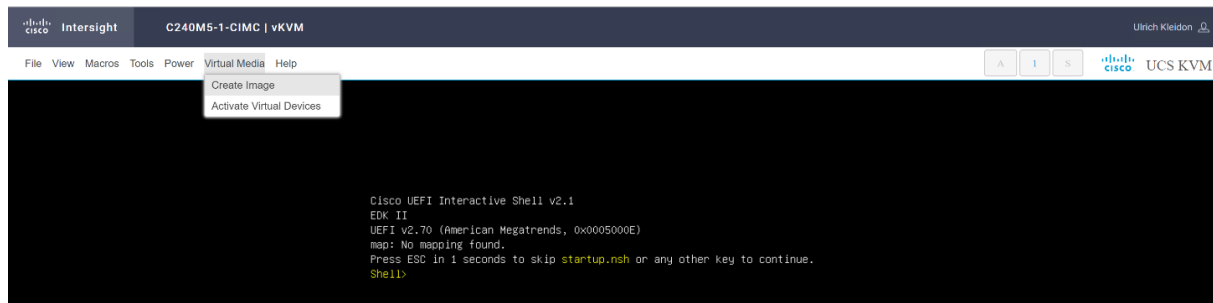
You have a SSL certificate for remote presence port. You should close this window now.

Close

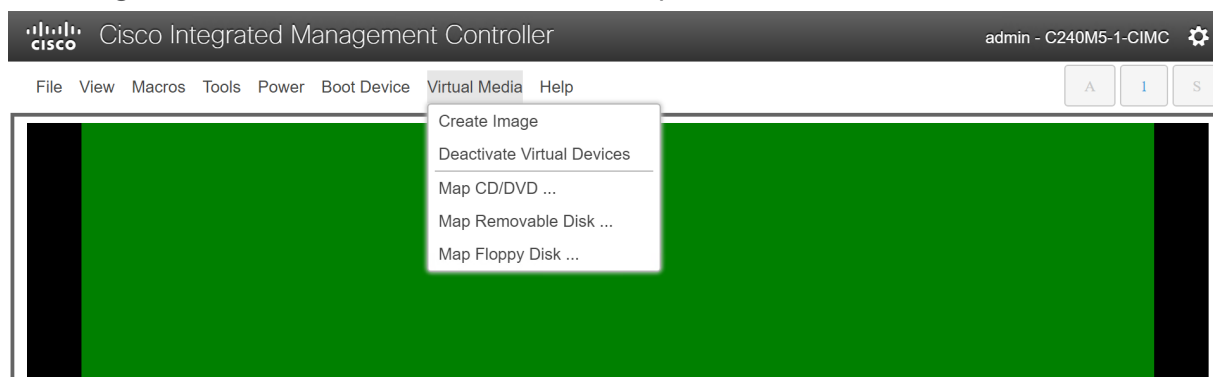
## Install SLE Micro

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


1. In the vKVM window, click Virtual Media and choose Activate Virtual Devices.



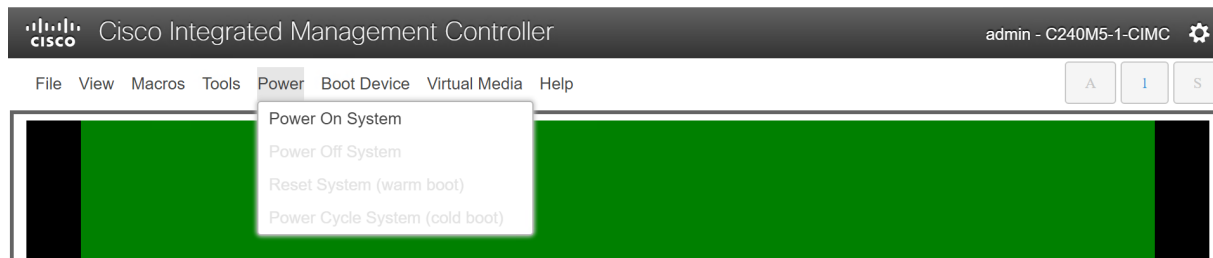
2. Again click Virtual Media and now choose Map CD/DVD.



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.



Copyright (c) 2020 Cisco Systems, Inc.

Press <F2> BIOS Setup : <F6> Boot Menu : <F7> Diagnostics  
Press <F8> CIMC Setup : <F12> Network Boot  
Bios Version : C240M5.4.1.2b.0.0917201934  
Platform ID : C240M5

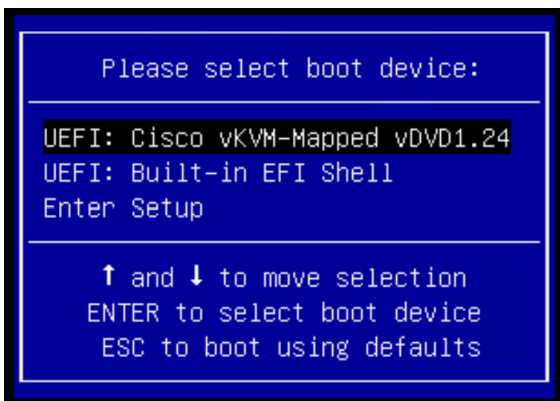
Processor(s) Intel(R) Xeon(R) Silver 4210 CPU @ 2.20GHz  
Total Memory = 192 GB Effective Memory = 192 GB  
Memory Operating Speed 2400 Mhz  
M.2 SWRAID configuration is not detected. Switching to AHCI mode.

Cisco IMC IPv4 Address : 172.16.1.83  
Cisco IMC MAC Address : 5C:71:0D:DA:54:FD

Entering Boot Menu ...

A2

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



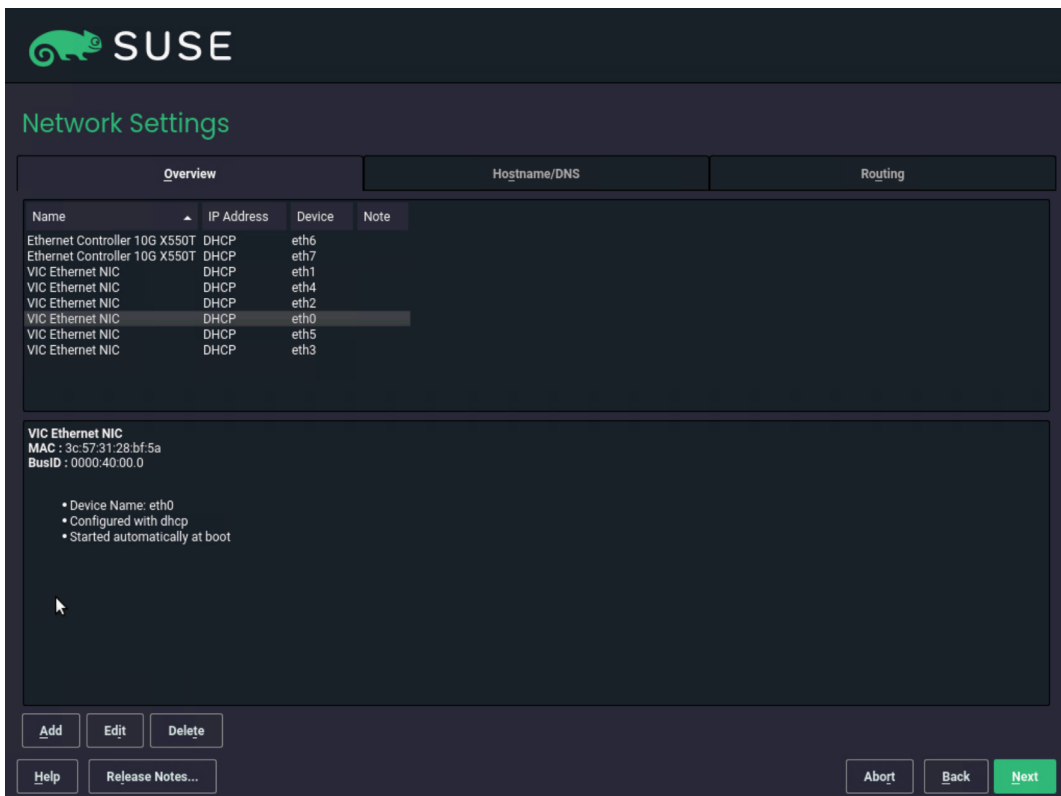
The SUSE installation process will start automatically.



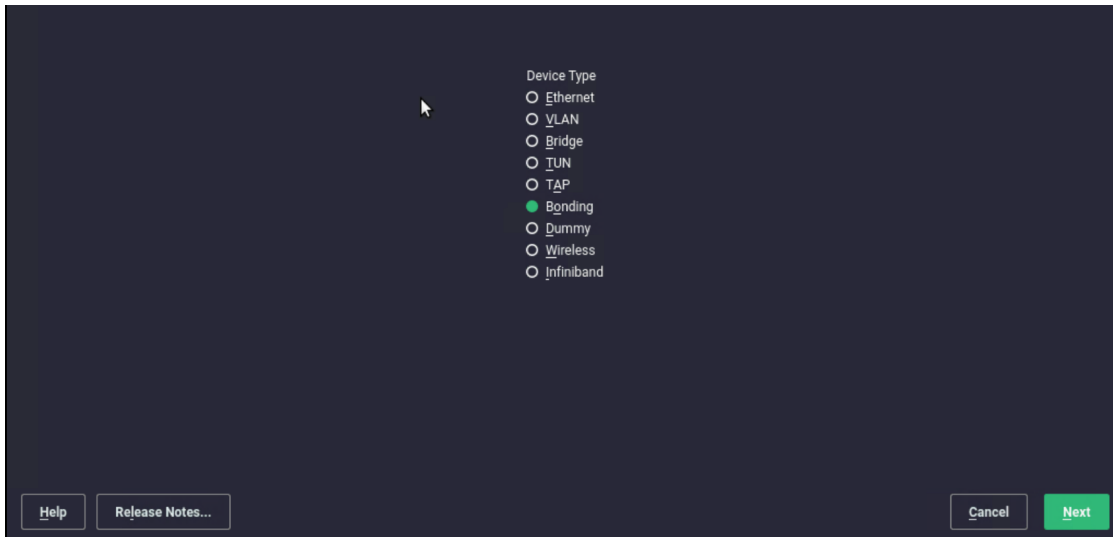
- Proceed with the installation process until you see the Installation Settings screen. Then click Network Configuration.



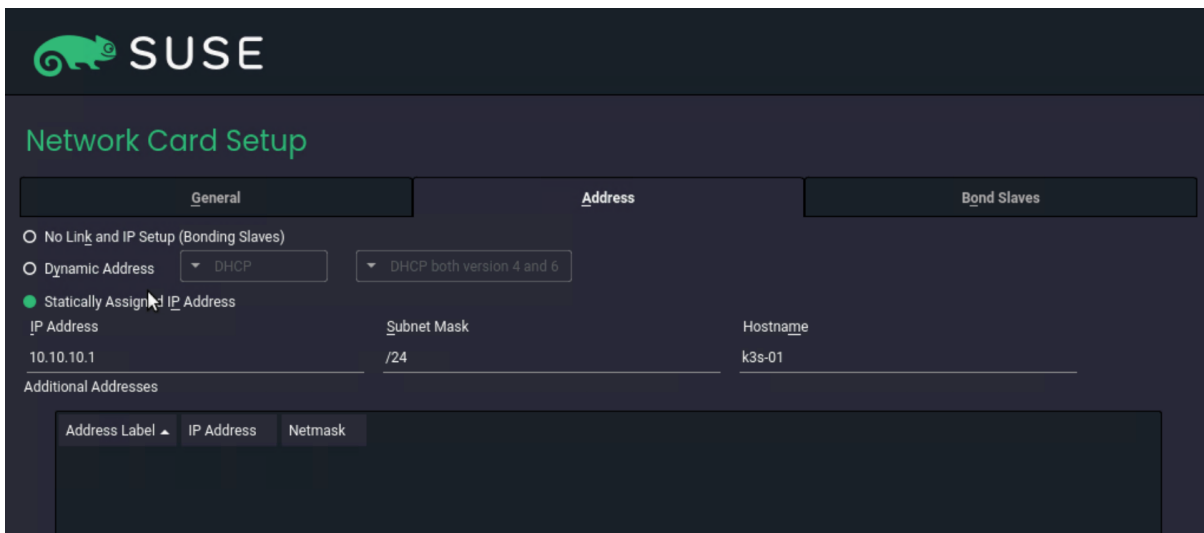
- Click the various devices in the network view and compare the names and MAC addresses with the vNIC list from the IMC. Click Add.



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



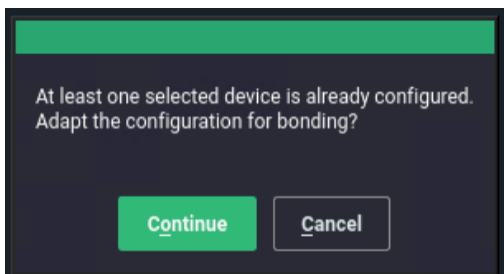
10. Enter the IP address, the netmask for the administration traffic network, and a hostname. Click Bond Slaves.



11. 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.

## Network Card Setup

General	Address	Bond Slaves
<b>Bond Slaves and Order</b>		
<input checked="" type="checkbox"/> eth0 - eth0 configured		
<input checked="" type="checkbox"/> eth1 - eth1 configured		
<input type="checkbox"/> eth2 - eth2 configured		
<input type="checkbox"/> eth3 - eth3 configured		
<input type="checkbox"/> eth4 - eth4 configured		
<input type="checkbox"/> eth5 - eth5 configured		
<input type="checkbox"/> eth6 - eth6 configured		
<input type="checkbox"/> eth7 - eth7 configured		
<div style="text-align: center;"><input type="button" value="Up"/> <input type="button" value="Down"/></div>		
<b>Bond Driver Options</b>		
<input type="text" value="mode=active-backup miimon=100"/>		
<input type="button" value="Help"/>	<input type="button" value="Release Notes..."/>	<input type="button" value="Cancel"/> <input type="button" value="Back"/> <input type="button" value="Next"/>



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



**SUSE**

## Network Settings

Overview      Hostname/DNS      Routing

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	

**bond0**  
(No hardware information)

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

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

**SUSE**

## Network Card Setup

General      Address      Bond Slaves

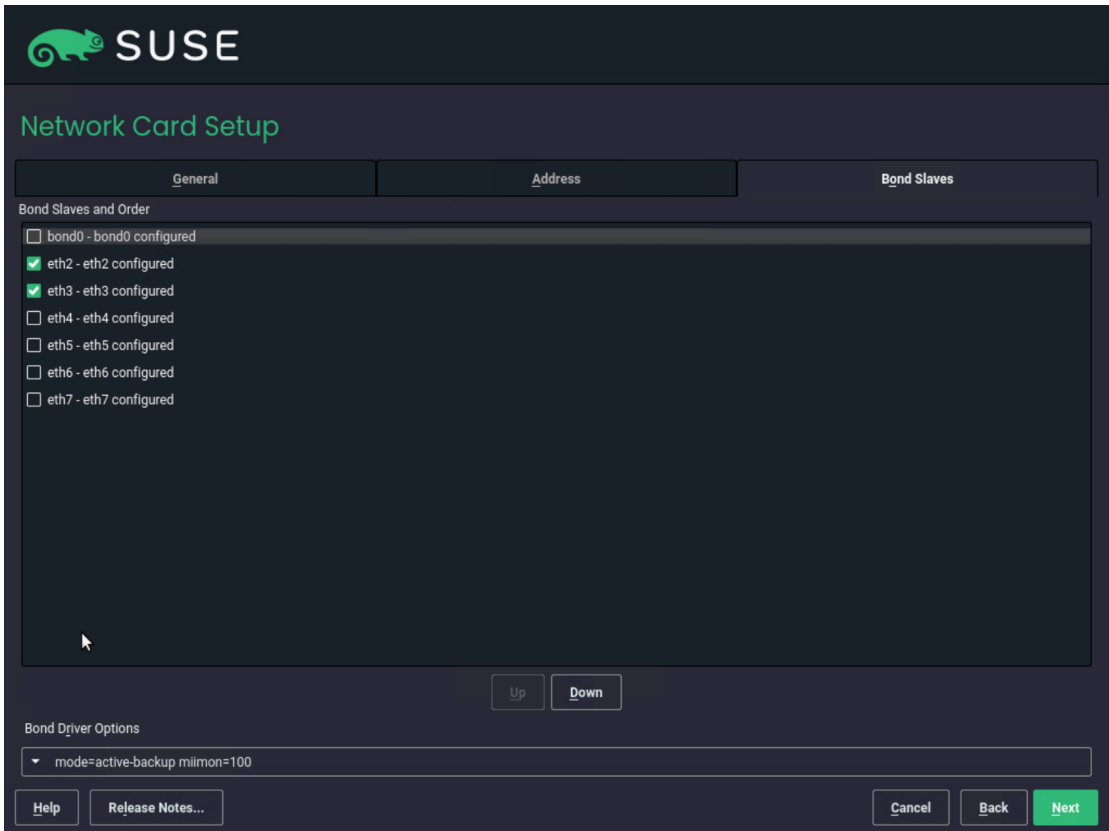
No Link and IP Setup (Bonding Slaves)  
 Dynamic Address    DHCP    DHCP both version 4 and 6  
 Statically Assigned IP Address

IP Address: 10.20.10.1    Subnet Mask: /24    Hostname: k3sa01-access

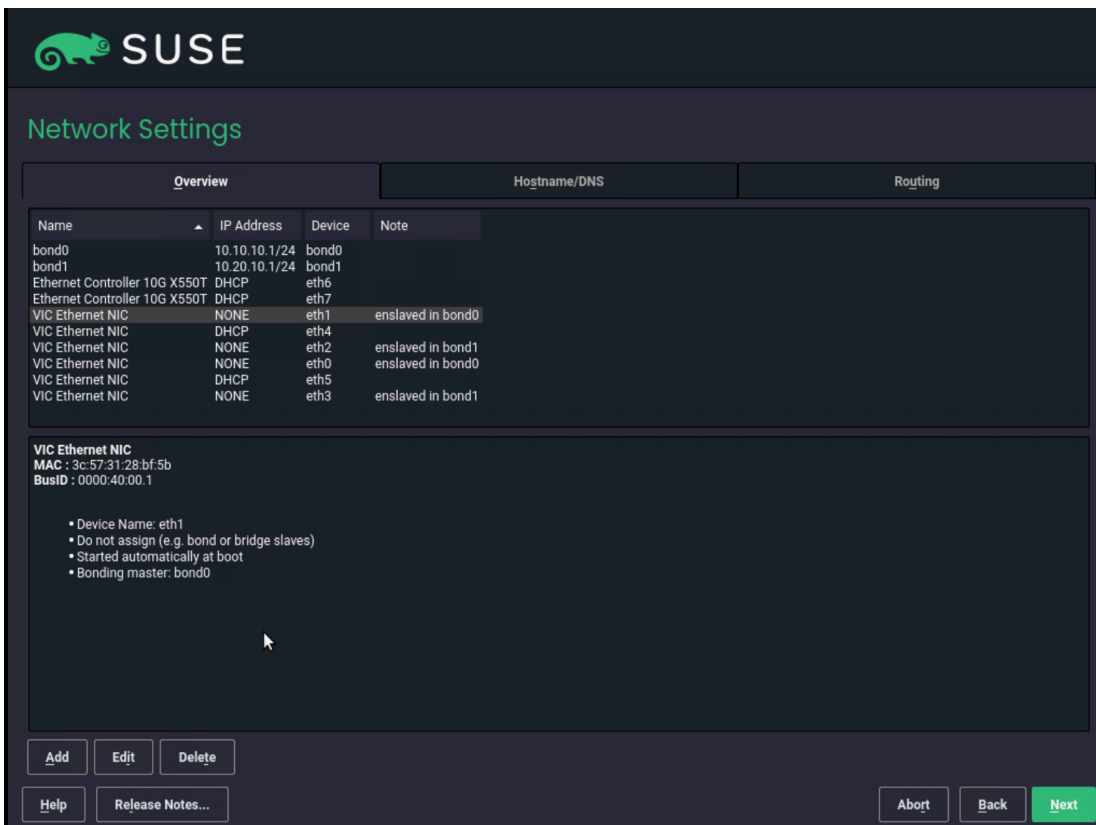
Additional Addresses

Address Label	IP Address	Netmask

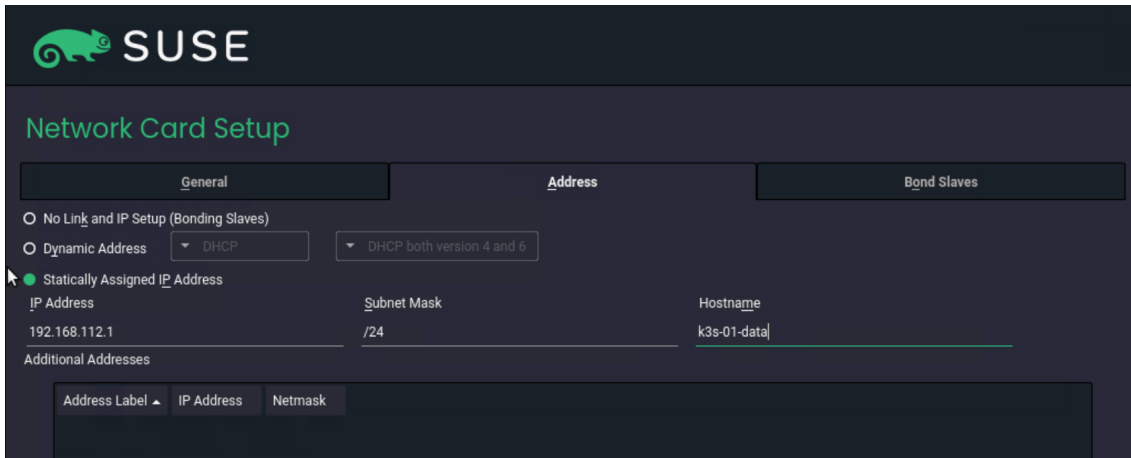
14. Select the two interfaces created for access traffic (eth2 and eth3) and use active-backup as the mode. Click Next. In the pop-up window, click Continue.



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

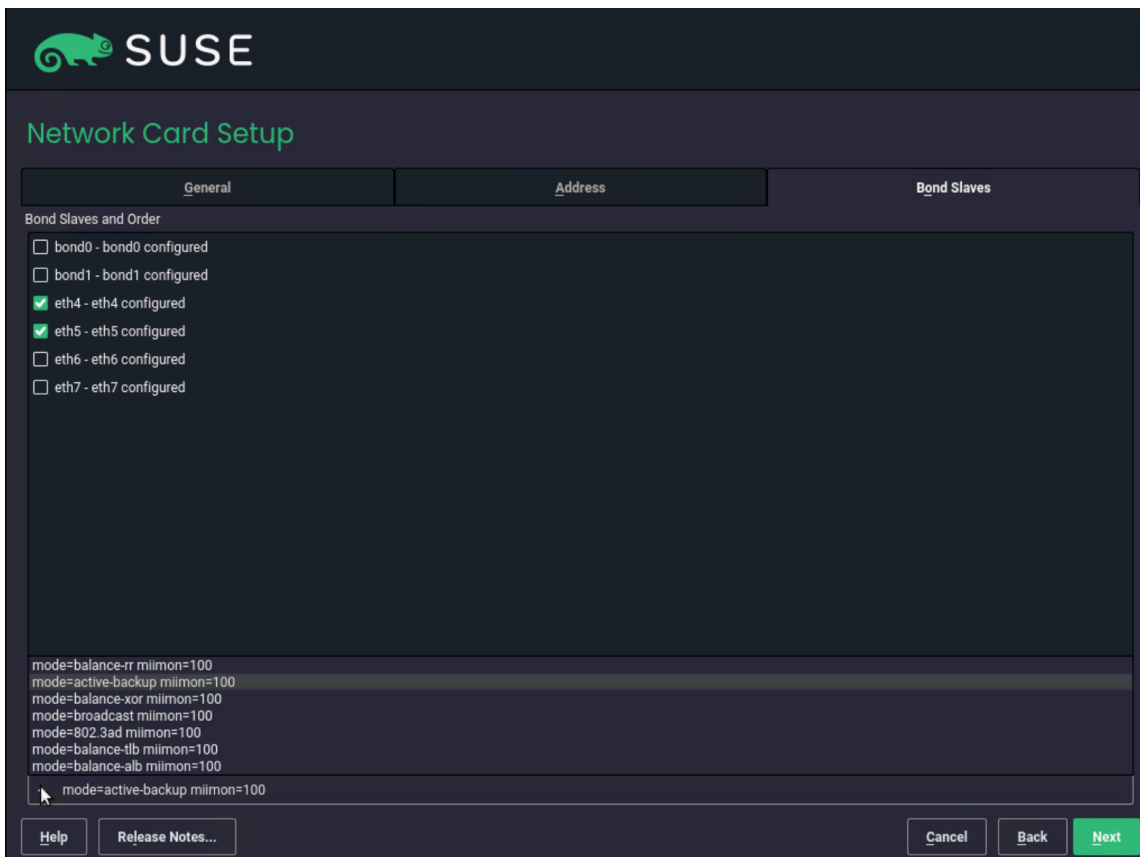


16. Enter the IP address and netmask for the storage traffic connection and a hostname. Click Bond Slaves.

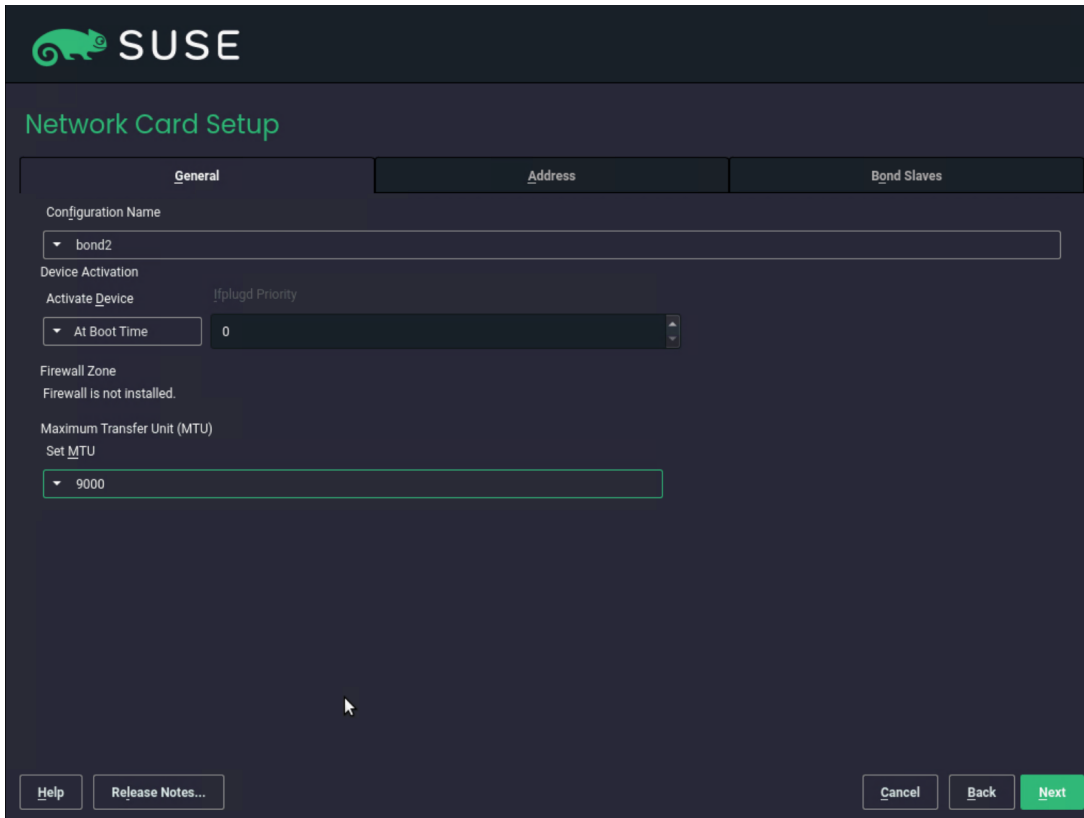


17. Select the two interfaces created for storage traffic (eth4 and eth5). Check with your networking and storage teams to determine whether an active-active bonding option for storage access is possible. An active-active option will increase the maximum throughput between this server and the storage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

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

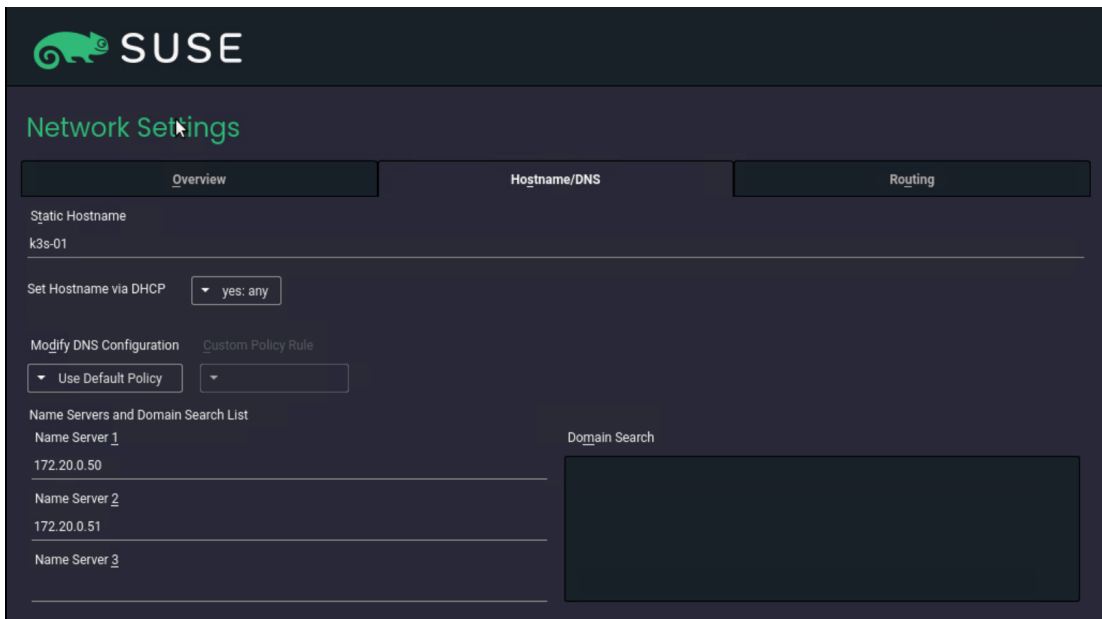


19. Enter **9000** in the field under Set MTU and click Next.

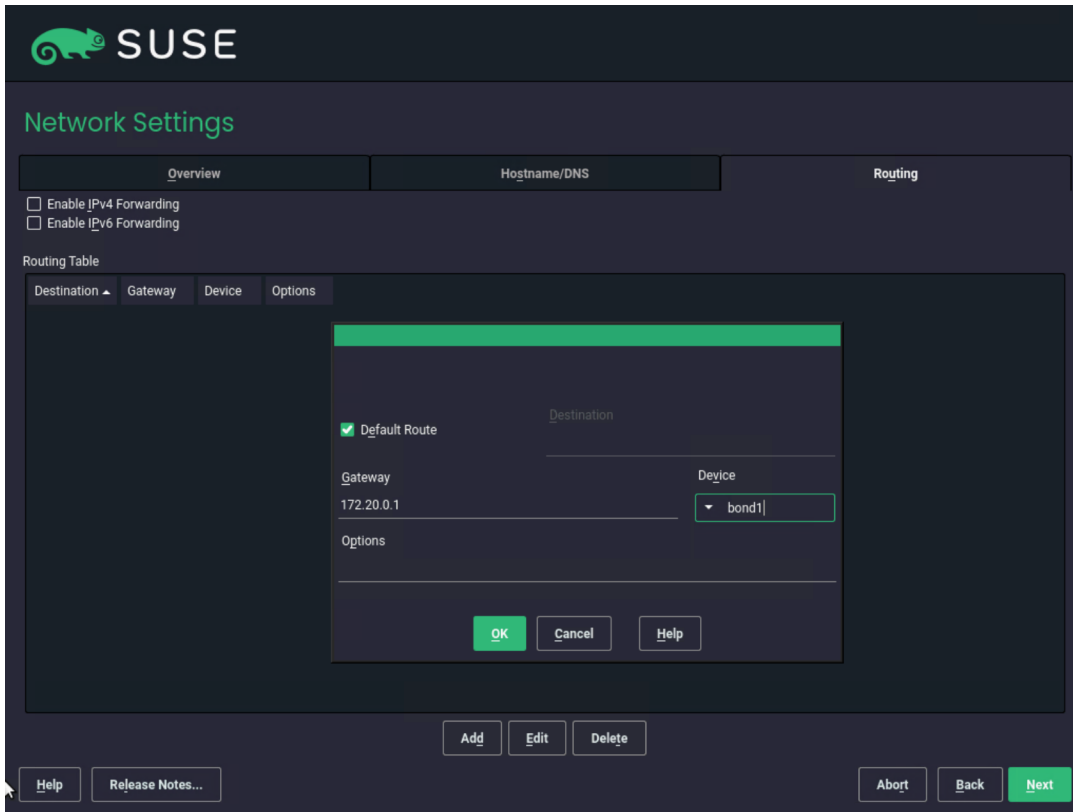


20. Click Hostname / DNS.

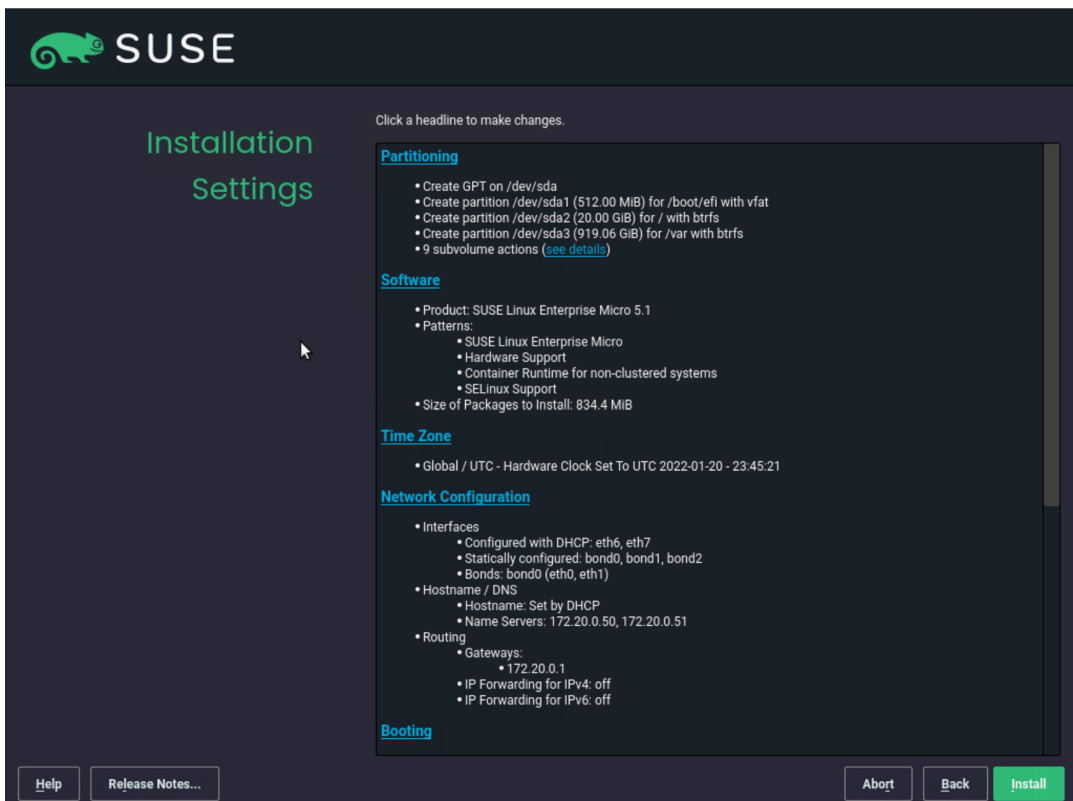
21. Enter the static hostname for this system and the IP address for at least one name server. Click Routing.



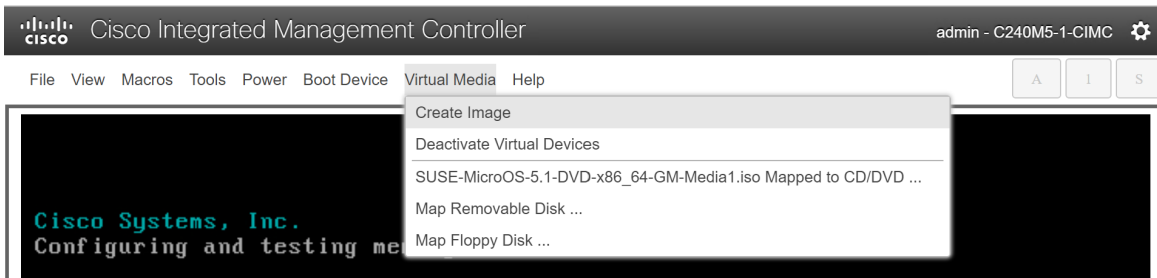
22. Click Add and in the pop-up window enter at least the default route for your network. Click OK. Click Next.



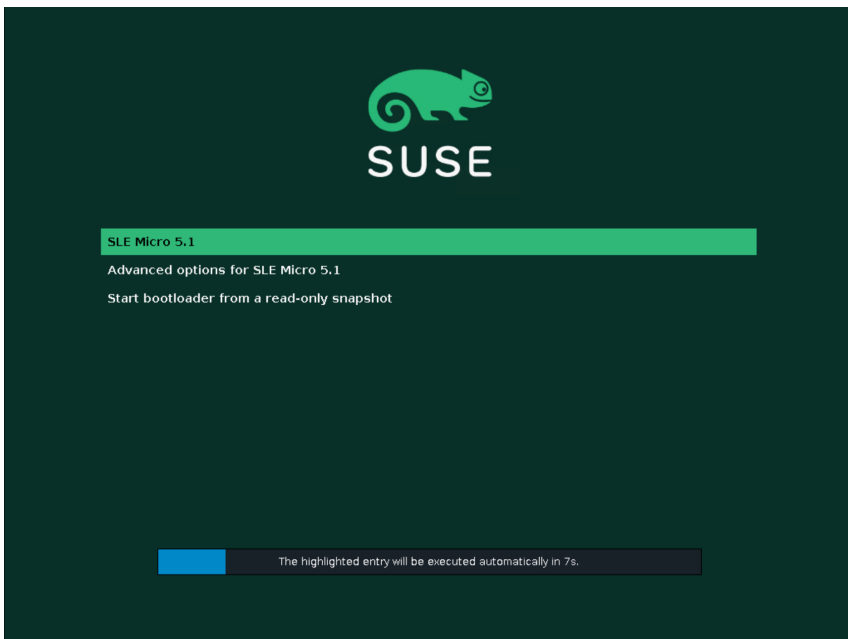
23. Back on the Installation Settings page, check all information and start the installation by clicking Install. Follow the next screens until the installation process is finished.



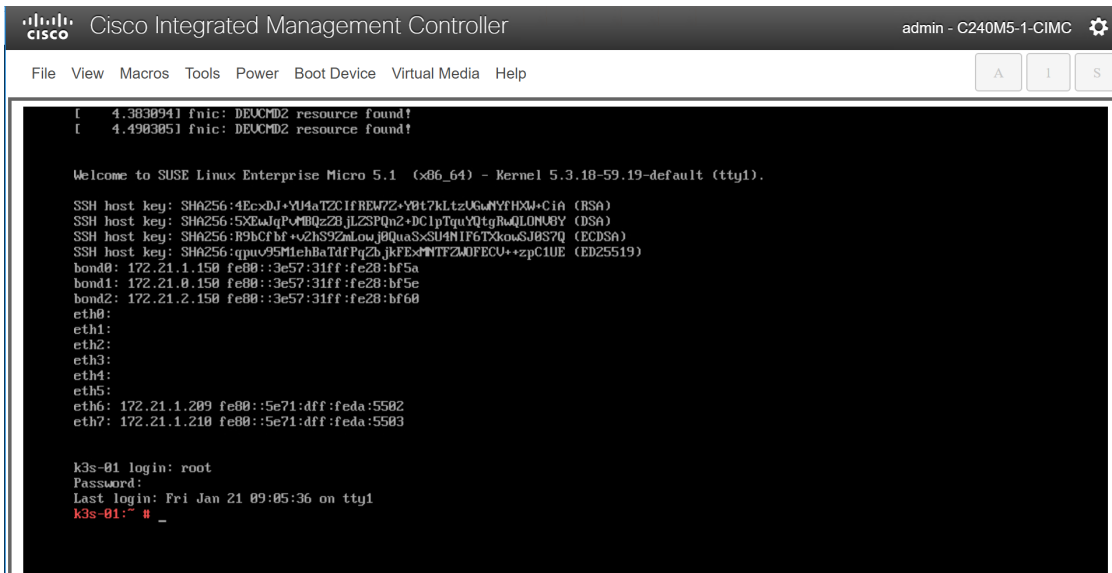
24. You must “eject” the CD/DVD as soon the installation process is finished and the reboot is initiated. Click Virtual Media > \*\*\*\*\* Mapped to CD/DVD and confirm the ejection by clicking OK in the pop-up window.



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



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



## 26. 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 it.

```
k3s-01:~ # curl -sL https://get.k3s.io | sh -s - --write-kubeconfig-mode 644
[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 youare 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/kubectl 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	11
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	NODE SELECTOR	AGE			
kube-system	daemonset.apps/svclb-traefik	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
kube-system	replicaset.apps/local-path-provisioner-64ffb68fd	1	1	1
kube-system	replicaset.apps/coredns-85cb69466	1	1	1
kube-system	replicaset.apps/metrics-server-9cf544f65	1	1	1
kube-system	replicaset.apps/traefik-786ff64748	1	1	1

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.

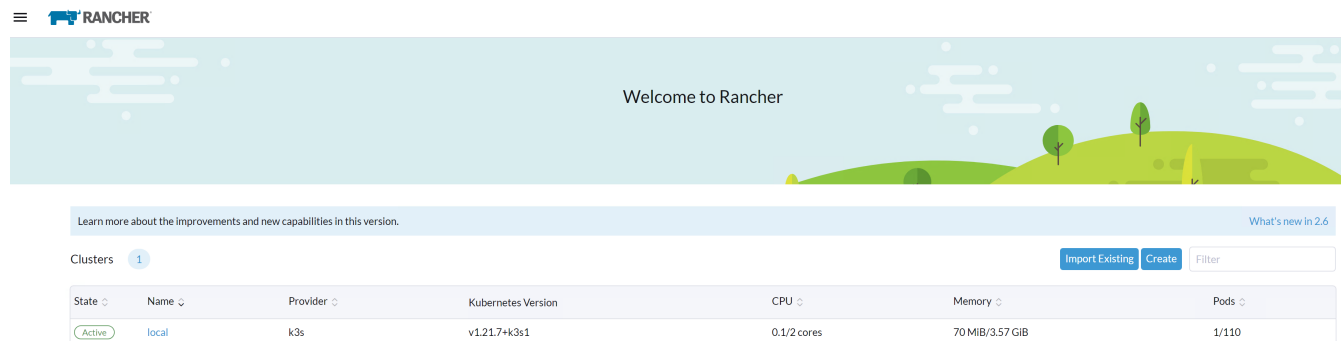
## K3s integration into the workload management tool

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 the SUSE Rancher Kubernetes Operations Platform and integration into the Rafay Kubernetes Operations Platform.

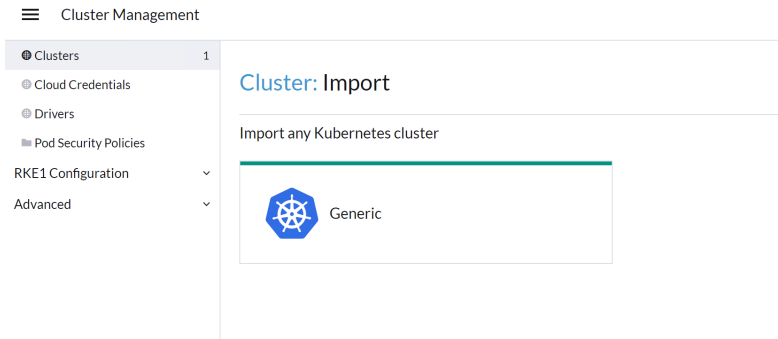
### Integrate into SUSE Rancher Kubernetes Operations Platform

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 SUSE Rancher Kubernetes Operations Platform.

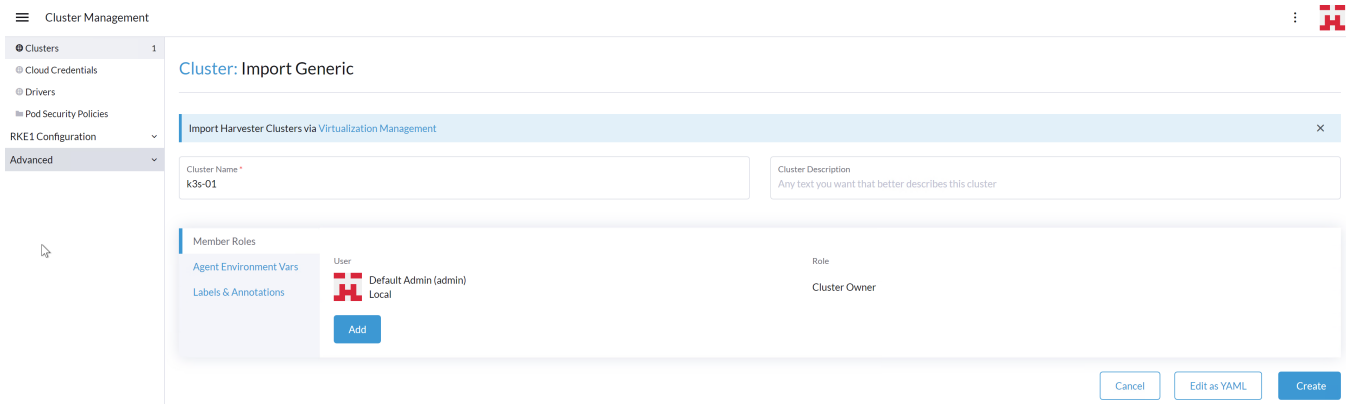
1. In the SUSE Rancher console, navigate to the list of clusters and click Import Existing.



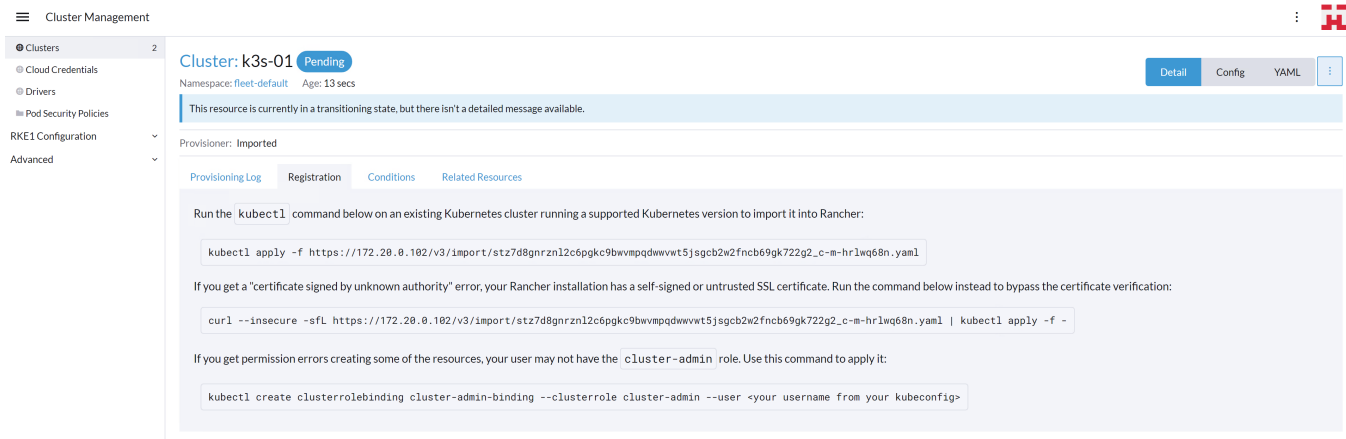
2. Click Generic.



### 3. Enter a cluster name and click Create.



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



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

```
k3s-01:~ # kubectl 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-
hrlwq68n.yaml | kubectl 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.nodeSelectorTerms[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

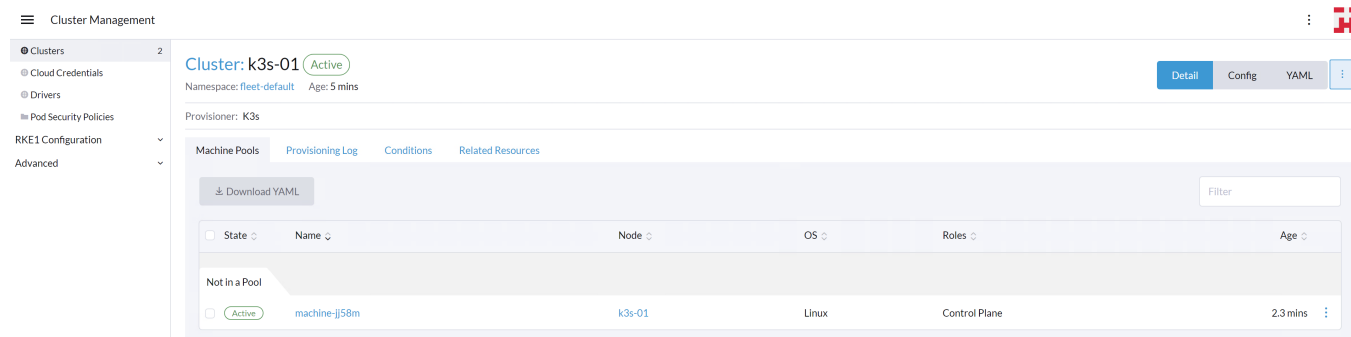
NAME                                TYPE          CLUSTER-IP    EXTERNAL-IP  PORT(S)          AGE
service/cattle-cluster-agent           ClusterIP     10.43.118.86  <none>       80/TCP,443/TCP  3m20s

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 SUSE Rancher user interface. The cluster is now shown as Active.



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

Welcome to Rancher

Learn more about the improvements and new capabilities in this version.

What's new in 2.6

Clusters 2

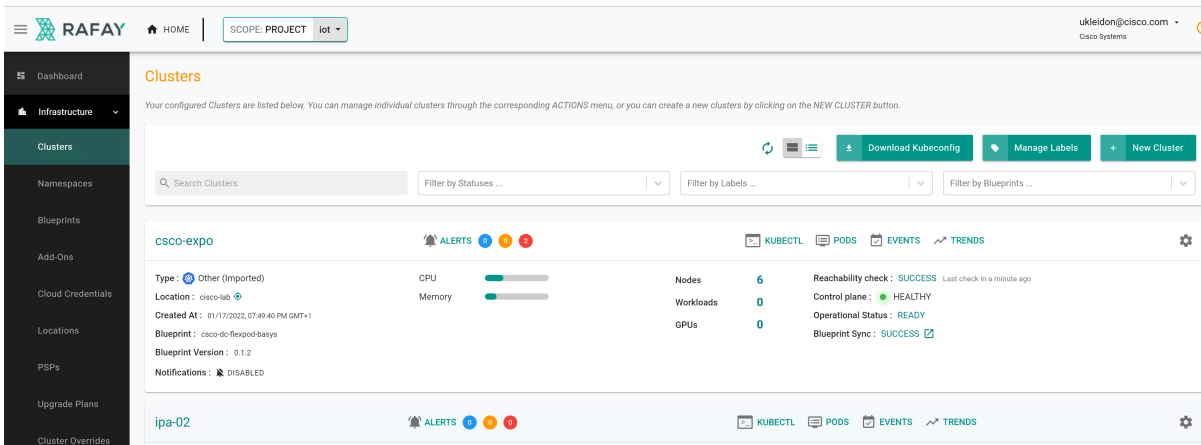
Import Existing Create Filter

State	Name	Provider	Kubernetes Version	CPU	Memory	Pods
Active	k3s-01	k3s	v1.22.5+k3s1	0/2 cores	0 B/3.57 GiB	0/110
Active	local	k3s	v1.21.7+k3s1	0.1/2 cores	70 MiB/3.57 GiB	1/110

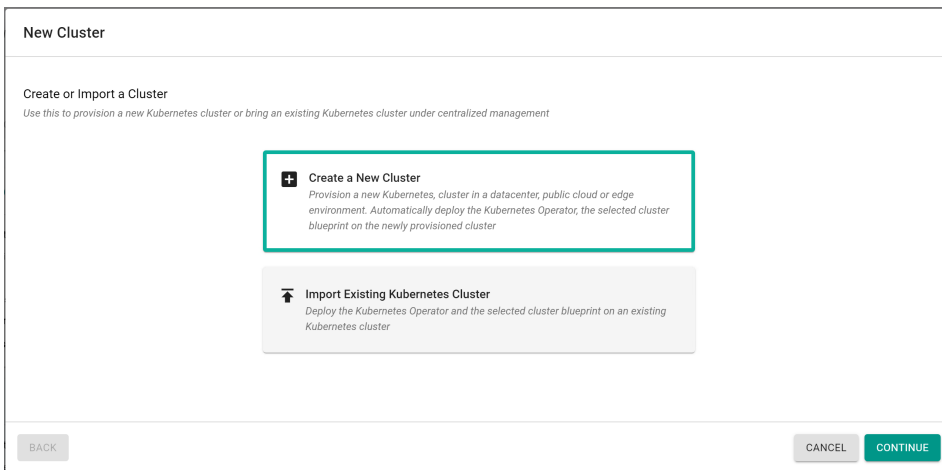
## Integrate into Rafay Kubernetes Operations Platform

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.



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



4. 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, it is best practice to start with the blueprint minimal or default. Those are the basic blueprints from Rafay to make the system work (minimal) or add components such as monitoring and reporting (default). Then click Continue.

5. Download the Bootstrap YAML file to the K3s system.

6. 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.

Clusters > k3s-01

Cluster Status **PROVISIONING**

- ✓ Cluster Register Complete
- ✓ Cluster Checkin Complete
- ✓ Cluster Namespace Sync Complete
- ⏳ Cluster Blueprint Sync Pending

Cluster Registration Instructions

- Click the button below to download the bootstrap YAML file to register the cluster.

[Download Bootstrap YAML](#)

- Run `"kubectl apply -f [path to file]/k3s-01-bootstrap.yaml"` on your kubernetes cluster.
- Once the bootstrap YAML file is installed, the status of the cluster will be reflected in the status panel. Generally it might take 3-5 mins for the registration to complete.

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

Clusters

Your configured Clusters are listed below. You can manage individual clusters through the corresponding ACTIONS menu, or you can create a new clusters by clicking on the NEW CLUSTER button.

Download kubeconfig Manage Labels + New Cluster

Search Clusters Filter by Statuses ... Filter by Labels ... Filter by Blueprints ...

**k3s-01** ALERTS 0 0 0 KUBECTL PODS EVENTS TRENDS

Type: Other (Imported) CPU Memory Nodes: 1 Workloads: 0 GPUs: 0 Reachability check: SUCCESS Last check in a few seconds ago Control plane: HEALTHY Operational Status: READY Blueprint Sync: SUCCESS

Location: cisco-lab Created At: 01/21/2022, 12:10:00 PM GMT+1 Blueprint: minimal Blueprint Version: snapshot-2022-01-08T05:04:47Z Notifications: DISABLED

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)	AGE	
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
NODE SELECTOR	AGE					
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 C240 SD server can run modern cloud-native applications developed for Kubernetes in a single server deployed in a short-depth network cabinet. With the Cisco Intersight platform, all servers can be monitored and operated from a single place, regardless of where they are deployed.

## 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://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/hw/c240sdm5/install/c240sdm5.html)
- <https://www.intersight.com/help>
- <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)