



Event Driven Automation

Release 24.12.2

Installation Guide

3HE 20938 AAAC TQZZA

Edition: 1

February 2025

© 2025 Nokia.

Use subject to Terms available at: www.nokia.com/terms.

Nokia is committed to diversity and inclusion. We are continuously reviewing our customer documentation and consulting with standards bodies to ensure that terminology is inclusive and aligned with the industry. Our future customer documentation will be updated accordingly.

This document includes Nokia proprietary and confidential information, which may not be distributed or disclosed to any third parties without the prior written consent of Nokia.

This document is intended for use by Nokia's customers ("You"/"Your") in connection with a product purchased or licensed from any company within Nokia Group of Companies. Use this document as agreed. You agree to notify Nokia of any errors you may find in this document; however, should you elect to use this document for any purpose(s) for which it is not intended, You understand and warrant that any determinations You may make or actions You may take will be based upon Your independent judgment and analysis of the content of this document.

Nokia reserves the right to make changes to this document without notice. At all times, the controlling version is the one available on Nokia's site.

No part of this document may be modified.

NO WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY OF AVAILABILITY, ACCURACY, RELIABILITY, TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, IS MADE IN RELATION TO THE CONTENT OF THIS DOCUMENT. IN NO EVENT WILL NOKIA BE LIABLE FOR ANY DAMAGES, INCLUDING BUT NOT LIMITED TO SPECIAL, DIRECT, INDIRECT, INCIDENTAL OR CONSEQUENTIAL OR ANY LOSSES, SUCH AS BUT NOT LIMITED TO LOSS OF PROFIT, REVENUE, BUSINESS INTERRUPTION, BUSINESS OPPORTUNITY OR DATA THAT MAY ARISE FROM THE USE OF THIS DOCUMENT OR THE INFORMATION IN IT, EVEN IN THE CASE OF ERRORS IN OR OMISSIONS FROM THIS DOCUMENT OR ITS CONTENT.

Copyright and trademark: Nokia is a registered trademark of Nokia Corporation. Other product names mentioned in this document may be trademarks of their respective owners.

© 2025 Nokia.

Table of contents

1	About this document.....	5
1.1	What's new.....	5
1.2	Precautionary and information messages.....	5
2	EDA installation overview.....	7
2.1	Deployment models.....	7
2.2	Networking for EDA nodes.....	7
2.3	EDA nodes.....	7
2.4	Requirements for deployment.....	8
2.4.1	Installation platform requirements.....	8
2.4.2	Nokia EDA node requirements.....	9
2.4.3	Virtual IP address requirements.....	9
2.5	Installation process overview.....	9
3	Preparing for installation.....	11
3.1	Downloading the EDA Installation playground.....	11
3.1.1	Downloading the EDAADM tool.....	11
3.1.2	Installing additional tools.....	11
3.1.3	Obtaining the EDA packages.....	12
3.2	Download the Talos machine image.....	12
3.2.1	Downloading the KVM image.....	12
3.2.2	Downloading the VMware OVA image.....	13
4	Setting up the EDA virtual machine nodes.....	14
4.1	Preparing the EDAADM configuration file.....	14
4.1.1	EDAADM configuration file fields.....	14
4.1.2	Example EDAADM configuration file.....	17
4.2	Generating the Talos machine configurations.....	19
4.3	Deploying the Talos virtual machines.....	20
4.3.1	Creating the VM on bridged networks on KVM.....	20
4.3.2	Creating the VM on bridged networks on VMware vSphere.....	22
5	Bootstrap the Talos Kubernetes cluster.....	25
5.1	Bootstrapping Kubernetes on the primary node.....	25

5.2	Obtaining the Kubernetes config file for kubectl.....	25
5.3	Setting up the Rook Ceph storage cluster.....	26
6	Installing the EDA application.....	29
6.1	Customizing the installation file.....	29
6.1.1	The prefs.mk file.....	30
6.2	Installing Nokia EDA.....	31
6.3	Accessing the EDA deployment.....	32

1 About this document

This document provides the information you need to prepare for the installation of the EDA application and provides EDA installation procedures.

This document is intended for network technicians, administrators, operators, service providers, and others who need to understand how to install and provision EDA for deployment.



Note: This document covers the current release and may also contain some content that will be released in later maintenance loads. See the *EDA Release Notes* for information about features supported in each load.

1.1 What's new

This section lists the changes that were made in this release.

Table 1: What's new in EDA 24.12.2

Description	Location
Improvements to procedures	<ul style="list-style-type: none"> • Networking for EDA nodes • Deployment models • Installation platform requirements • Downloading the VMware OVA image • Generating the Talos machine configurations • Creating the VM on bridged networks on KVM • Setting up the Rook Ceph storage cluster • Bootstrap the Talos Kubernetes cluster • Installing Nokia EDA

1.2 Precautionary and information messages

The following are information symbols used in the documentation.



DANGER: Danger warns that the described activity or situation may result in serious personal injury or death. An electric shock hazard could exist. Before you begin work on this equipment, be aware of hazards involving electrical circuitry, be familiar with networking environments, and implement accident prevention procedures.



WARNING: Warning indicates that the described activity or situation may, or will, cause equipment damage, serious performance problems, or loss of data.



Caution: Caution indicates that the described activity or situation may reduce your component or system performance.



Note: Note provides additional operational information.



Tip: Tip provides suggestions for use or best practices.

2 EDA installation overview

This chapter describes the Nokia Event Driven Automation (EDA) components, the requirements for these components, and provides an overview of the installation process.

2.1 Deployment models

Nokia EDA is deployed on one, three, or more nodes (validated for up to six nodes). Nokia EDA is deployed as an application on virtual machine servers.

The nodes (VMs) run a Kubernetes cluster with the following composition:

- One or three Kubernetes master nodes that also function as worker nodes: one, in case a single-VM deployment is used; otherwise, three Kubernetes master nodes.
- Any remaining nodes (in a four or more node deployment) function as worker nodes.
- One, two or more nodes must also be designated as storage nodes. For redundancy, two is the minimum in a three or more node deployment. These nodes still function as worker (and potentially master) nodes as well. Rook-Ceph is used to create a storage cluster across the nodes indicated as storage nodes.

Related topics

[Nokia EDA node requirements](#)

2.2 Networking for EDA nodes

This guide describes the deployment of EDA on a Kubernetes cluster with two networks, where access from users and orchestrators to the UI and API, and access from EDA to the fabric (for example, SR Linux devices) will use different interfaces.

It is possible to use two separate networks for the EDA nodes:

- OAM network
This interface is used to access the UI and the API of Nokia EDA. It is also through this network that the deployment tool reaches the nodes.
- Fabric management network
This interface is used to communicate with the management interfaces of the fabric (for example, SR Linux devices). This interface is where Nokia EDA exposes its DHCP and ZTP services.

2.3 EDA nodes

The Nokia EDA nodes are the VMware vSphere-based or KVM-based virtual machines (VMs) that host the Kubernetes environment on which the Nokia EDA application and Digital Sandbox are run.

These nodes run a hardened Talos Kubernetes environment. Talos is a secure, up-to-date and hardened platform for running Kubernetes.

EDA supports the following deployment models:

- an environment with one node, which hosts only the Nokia EDA application for small scale deployments
- an environment with three nodes or more nodes, which hosts only the Nokia EDA application

2.4 Requirements for deployment

This section describes the platform requirements, node requirements, and virtual IP requirements for deploying EDA.

2.4.1 Installation platform requirements

To execute the installation process, you need access to a Linux environment with the following components installed.



Note:

When installing on a KVM-based hypervisor, the platform must use Red Hat Enterprise Linux or Rocky Linux. The installation steps of the installation platform requirements are for the platform running the edaadm and playground steps, not creating the VMs on KVM (this happens on the RHEL/Rocky Hypervisor directly).

Table 2: Platform requirements

Component	Requirement
Linux environment	Any Linux distribution. The procedures provided in this document are validated on Ubuntu.
Container runtime: Docker	Docker must be running and you should be able to run containers
Tools	<ul style="list-style-type: none"> • make • git • curl • wget • helm <p>The following tools are also helpful. If they are not present, the installation tool downloads them later:</p> <ul style="list-style-type: none"> • kubectl • yq • k9s • kpt
Internet access	Either directly or through a proxy

2.4.2 Nokia EDA node requirements

The Nokia EDA nodes are deployed as virtual machine servers. [Table 3: Node requirements](#) summarizes the requirements of Nokia EDA nodes in KVM and VMware hypervisor.

Table 3: Node requirements

Component	Requirement
CPU	32 vCPU on a modern x86-64 CPU
Memory	64GB
Storage	<ul style="list-style-type: none"> Operating system: 100GB of available SSD-based storage Storage nodes: 300GB of available SSD-based storage on a separate virtual disk
Networking	<ul style="list-style-type: none"> at least one 10 Gbps NIC the configured DNS servers must be reachable, functional, and able to resolve the hostnames used for the Nokia EDA nodes Internet access directly or through a proxy
Virtualization platform	<p>You can run the Nokia EDA nodes as virtual machines using the following virtualization platforms:</p> <ul style="list-style-type: none"> operating system: VMware vSphere 7.0 or 8.0 or RHEL/Rocky hypervisor: ESXi 7.0 or 8.0 or KVM resource reservation for CPU, memory, and disks must be set to 100% for the Nokia EDA node virtual machines

2.4.3 Virtual IP address requirements

The deployment requires two virtual IP addresses in the management network work:

- Kubernetes VIP: the virtual IP address used by all the control plane nodes in the Kubernetes cluster.
- Nokia EDA API/UI VIP: the virtual IP address used by the Nokia EDA API and UI.
- Nokia EDA fabric management VIP address: in a setup with two interfaces, the virtual IP address used by the Nokia EDA API and DHCP relay for the fabric management network.

2.5 Installation process overview

The installation consists of the following high-level tasks:

1. [Downloading the EDA Installation playground](#)

This task describes how to access the EDAADM tool and initiates the EDA installation playground for use during the installation. It also covers how to configure the playground.

2. [Download the Talos machine image](#)

This task describes how to download the Talos base image from the official Talos image factory for your environment.

3. [Preparing the EDAADM configuration file](#)

This task describes the details of the EDAADM configuration file and how to set it up.

4. [Generating the Talos machine configurations](#)

Using the EDA ADM tool and the configuration file, this task generates specific Talos machine configuration files for each Talos VM.

5. [Deploying the Talos virtual machines](#)

This task describes how to use the Talos base image and machine configuration files to deploy the Talos VMs in your KVM or VMware vSphere environment.

6. [Bootstrap the Talos Kubernetes cluster](#)

This task bootstraps the Talos Kubernetes environment using the VMs you have created.

7. [Installing the EDA application](#)

Using the EDA Installation playground, this step installs EDA on the Kubernetes environment in the EDA nodes.

3 Preparing for installation

This section includes the tasks that you need to complete to prepare for the installation of EDA.

3.1 Downloading the EDA Installation playground

Prerequisites

Ensure that your Linux installation environment meets the requirements described in [Installation platform requirements](#). Git must already be installed on your computer.

Procedure

Clone the playground repository locally.

```
git clone https://github.com/nokia-eda/playground && cd playground
```

3.1.1 Downloading the EDAADM tool

Procedure

The EDAADM tool is used to generate configuration files for use while deploying the Talos Linux virtual machines and the Kubernetes environment.

You can download the binary for different platforms from <https://github.com/nokia-eda/edaadm/releases/>. Make sure to download the version that matches your operating system and CPU architecture.

Place the tool in a handy location and make sure you can execute it.

3.1.2 Installing additional tools

Procedure

Step 1. Set up additional tools that can be used during the installation.
You can download these tools to a local folder.

Example

```
make download-tools
```

Step 2. Verify that tools have been downloaded.
The `kind`, `kubectl`, `kpt`, and `yq` utilities are installed in the `~/tools/` directory.

3.1.3 Obtaining the EDA packages

About this task

EDA is packaged using the Kubernetes Package Tool (kpt). EDA uses this package manager tool to install core EDA components. The installer downloads two kpt packages by downloading their relevant git repositories.

Procedure

To obtain the EDA package, enter the following command:

```
make download-pkgs
```

This command downloads the following git repositories to their respective directories:

- EDA kpt package in the `~/eda-kpt` directory
- EDA built-in catalog in the `~/catalog` directory

3.2 Download the Talos machine image

The EDAADM tool provides you with the URL where you can download the latest image VMware OVA image or the image for use with KVM.

To deploy the Talos Kubernetes environment, download the Talos machine image based on the environment in which you want to deploy the VMs.

3.2.1 Downloading the KVM image

Procedure

Step 1. Use the EDAADM tool to display the URL from where you can download the latest image for use with KVM for the supported Talos version.

Example

```
$ edaadm images --mach-type nocloud
Schematic ID is :376567988ad370138ad8b2698212367b8edcb69b5fd68c80be1f2ec7d603b4ba
Asset URLs are:
https://factory.talos.dev/image/
376567988ad370138ad8b2698212367b8edcb69b5fd68c80be1f2ec7d603b4ba/v1.8.3/nocloud-
amd64.iso
https://factory.talos.dev/image/
376567988ad370138ad8b2698212367b8edcb69b5fd68c80be1f2ec7d603b4ba/v1.8.3/nocloud-
amd64.raw.xz
```

Step 2. Download the `nocloud-amd64.iso` image from the ISO URL, `filepath.iso`. You can download using your browser or you can use the **curl** or **wget** commands.

3.2.2 Downloading the VMware OVA image

Procedure

- Step 1.** Use the EDAADM tool to display the URL from where you can download latest image for use with VMware vSphere for the supported Talos version.

Example

```
$ edaadm images --mach-type vmware
Schematic ID is : 903b2da78f99adef03cbbd4df6714563823f63218508800751560d3bc3557e40
Asset URLs are:
https://factory.talos.dev/image/
903b2da78f99adef03cbbd4df6714563823f63218508800751560d3bc3557e40/v1.8.3/vmware-
amd64.iso
https://factory.talos.dev/image/
903b2da78f99adef03cbbd4df6714563823f63218508800751560d3bc3557e40/v1.8.3/vmware-
amd64.ova
```

- Step 2.** Download the `vmware-amd64.iso` image from the OVA URL, `filepath.ova`. You can download using your browser or you can use the **curl** or **wget** commands. You can also use the URL directly with the `ovftool` command to deploy the OVA to your VMware vSphere environment.

4 Setting up the EDA virtual machine nodes

This section describes how to prepare the configurations file, generate the configuration files, and deploy the Talos virtual machines.

4.1 Preparing the EDAADM configuration file

The EDAADM tool helps with the creation of the necessary machine configuration files for the Talos VMs that are part of your deployment.

4.1.1 EDAAM configuration file fields

The EDAADM configuration file is a YAML file that describes your Talos Kubernetes environment. You can use it to configure the different nodes and the general Kubernetes cluster environment.

Table 4: Talos Kubernetes environment parameters

Top-level parameters	Description	
version	The version of the EDA environment to be deployed. Example: 24.12.1	
clusterName	The name of your EDA environment. Example: eda-production-cluster	
machines	A list of Kubernetes nodes. Each Kubernetes node has the following settings:	
	name	The name of a node. Example: eda-node01
	endpoint	The IP address on which the node is reachable for Talos to control. This field must be set to one of the IP addresses configured for the interfaces on the node. Optional.
	interfaces	A list of interfaces present in the node, each with the following settings: <ul style="list-style-type: none"> name: the name of the interface. Example: eth0 dhcp: indicates if DHCP is to be used for the interface. Values: true or false. For production environments, set to false. mtu: the MTU setting for the interface. For an interface used to connect to nodes under

Top-level parameters	Description	
		<p>management, set to 9000 for best practice. Optional.</p> <ul style="list-style-type: none"> • interface: the interface name as it appears in Linux. Typically, eth0, eth1, and so forth. Optional. • addresses: a list of IP addresses; for dual-stack deployments, you can specify both IPv4 and IPv6 addresses. If DHCP is not provided, specify at least one address. • routes: a list of static routes to configure, including the default route. Optional. Routes have the following components: <ul style="list-style-type: none"> – gateway: the next-hop or gateway for the route. – metric: a metric to indicate the priority of the route. Optional. – mtu: a specific MTU for the route. Optional. – network: the destination CIDR of the route. – source: a source interface for the route to apply to. Optional. • deviceSelector: specifies how to select the device associated with this interface. <ul style="list-style-type: none"> – busPath: a PCI buspath that can contain wildcards. Optional. – hardwareAddr: a MAC address that can contain wildcards. Optional.
	disks	<p>Identifies the disks available in the node:</p> <ul style="list-style-type: none"> • os: Specifies which disk to use for the OS. Required setting. Typically /dev/sda or /dev/vda, depending on the hypervisor platform. • storage: Optional disk for use with nodes that are to be part of the storage cluster.
k8s	The Kubernetes-specific configuration. The following parameters define the Kubernetes cluster:	
	stack	Indicates the network stack to support. Values: <code>ipv4</code> , <code>ipv6</code> , or <code>dual</code>
	primaryNode	The first control plane node in the cluster to be used for bootstrapping the Kubernetes cluster. Specify the a the name of a machine.

Top-level parameters	Description	
	endpointUrl	The URL on which to reach the Kubernetes control plane. This setting uses the Kubernetes VIP address. Example: https://192.0.2.10:6443
	allowSchedulingOnControlPlanes	Specifies if workloads can be deployed on the control plane node. Values: true or false. For best practice, set to true.
	control-plane	A list of control plane nodes. Specify a machine name.
	worker	A list of worker nodes. Specify a machine name.
	vip	The VIP addresses used for Kubernetes and the interfaces to which they should be attached in the control plane nodes. Depending on the IP stack in use, some values are required: <ul style="list-style-type: none"> interface: the interface to which the VIP is attached on the nodes. Example: eth0 ipv4: the IPv4 VIP address. Example: 192.0.2.10 ipv6: the IPv6 VIP address
	env	Section that includes the optional proxy settings for the Kubernetes nodes: <ul style="list-style-type: none"> http_proxy: The HTTP proxy URL to use. Example: http://192.0.2.254:808 https_proxy: the HTTPS proxy URL to use. Example: http://192.0.2.254:808 no_proxy: the no proxy setting for IP addresses, IP ranges, and hostnames
	time	Defines NTP settings. <ul style="list-style-type: none"> disabled: Specifies whether NTP is enabled. For production environments, set to false to enable NTP. servers: A list of NTP servers; required for production environments.
	nameservers	A list of DNS servers specified under the following sub-element: servers: the list of DNS servers

Top-level parameters	Description	
	certBundle	An optional set of PEM-formatted certificates that need to be trusted; this setting is used for trust external services.

4.1.2 Example EDAADM configuration file

The following example shows an EDAADM configuration file for a 6-node Kubernetes cluster called `eda-input-6-node.yaml`. This file serves as the reference and example throughout this guide.

In this configuration file, two networks are used. If your deployment only uses one network, the second interface would not be present for each of the machines.

```

version: 24.12.1
clusterName: eda-compute-cluster
machines:
  - name: eda-node01
    endpoint: "192.0.2.11"
    interfaces:
      - name: eth0
        dhcp: false
        interface: eth0
        addresses:
          - 192.0.2.11/24
        routes:
          - network: 0.0.0.0/0
            gateway: 192.0.2.1
        mtu: 9000
      - name: eth1
        dhcp: false
        interface: eth1
        addresses:
          - 203.0.113.11/24
        mtu: 9000
    disks:
      os: /dev/vda
      storage: /dev/vdb
  - name: eda-node02
    endpoint: "192.0.2.12"
    interfaces:
      - name: eth0
        dhcp: false
        interface: eth0
        addresses:
          - 192.0.2.12/24
        routes:
          - network: 0.0.0.0/0
            gateway: 192.0.2.1
        mtu: 9000
      - name: eth1
        dhcp: false
        interface: eth1
        addresses:
          - 203.0.113.12/24
        mtu: 9000
    disks:
      os: /dev/vda
      storage: /dev/vdb
  - name: eda-node03

```

```
endpoint: "192.0.2.13"
interfaces:
  - name: eth0
    dhcp: false
    interface: eth0
    addresses:
      - 192.0.2.13/24
    routes:
      - network: 0.0.0.0/0
        gateway: 192.0.2.1
    mtu: 9000
  - name: eth1
    dhcp: false
    interface: eth1
    addresses:
      - 203.0.113.13/24
    mtu: 9000
disks:
  os: /dev/vda
  storage: /dev/vdb
- name: eda-node04
  endpoint: "192.0.2.14"
  interfaces:
    - name: eth0
      dhcp: false
      interface: eth0
      addresses:
        - 192.0.2.14/24
      routes:
        - network: 0.0.0.0/0
          gateway: 192.0.2.1
      mtu: 9000
    - name: eth1
      dhcp: false
      interface: eth1
      addresses:
        - 203.0.113.14/24
      mtu: 9000
  disks:
    os: /dev/vda
- name: eda-node05
  endpoint: "192.0.2.15"
  interfaces:
    - name: eth0
      dhcp: false
      interface: eth0
      addresses:
        - 192.0.2.15/24
      routes:
        - network: 0.0.0.0/0
          gateway: 192.0.2.1
      mtu: 9000
    - name: eth1
      dhcp: false
      interface: eth1
      addresses:
        - 203.0.113.15/24
      mtu: 9000
  disks:
    os: /dev/vda
- name: eda-node06
  endpoint: "192.0.2.16"
  interfaces:
    - name: eth0
```

```

    dhcp: false
    interface: eth0
    addresses:
      - 192.0.2.16/24
    routes:
      - network: 0.0.0.0/0
        gateway: 192.0.2.1
    mtu: 9000
  - name: eth1
    dhcp: false
    interface: eth1
    addresses:
      - 203.0.113.16/24
    mtu: 9000
  disks:
    os: /dev/vda
k8s:
  stack: ipv4
  primaryNode: eda-node01
  endpointUrl: https://192.0.2.5:6443
  allowSchedulingOnControlPlanes : true
  control-plane:
    - eda-node01
    - eda-node02
    - eda-node03
  worker:
    - eda-node04
    - eda-node05
    - eda-node06
  vip:
    ipv4: 192.0.2.5
    interface: eth0
  env:
    http_proxy: http://192.0.2.254:8080
    https_proxy: http://192.0.2.254:8080
    no_proxy: 192.0.2.0/24,203.0.113.0/24,.domain.tld,172.22.0.0/
16,localhost,127.0.0.1,10.0.1.0/24,0.0.0.0,169.254.116.108
  time:
    disabled: false
    servers:
      - 192.0.2.253
      - 192.0.2.254
  nameservers:
    servers:
      - 192.0.2.253
      - 192.0.2.254

```

4.2 Generating the Talos machine configurations

About this task

After creating the EDAADM configuration file, the next step is to generate all the configuration files that are necessary to deploy the Kubernetes environment using Talos.

Procedure

Use the EDAADM tool to generate the deployment files.

Example

```
$ edaadm generate -c eda-input-6-node.yaml
ConfigFile is eda-input-6-node.yaml
...
[1/4] Validating Machines
[1/4] Validated Machines
[2/4] Validating PrimaryNode
[2/4] Validated PrimaryNode
[3/4] Validating Endpoint URL
[3/4] Validated Endpoint URL
[4/4] Validating Virtual IP
[4/4] Validated Virtual IP
[ OK ] Spec is validated
Generating secrets for eda-compute-cluster
Created eda-compute-cluster/secrets.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node01.yaml
Created eda-compute-cluster/talosconfig.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node02.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node03.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node04.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node05.yaml
generating PKI and tokens
Created eda-compute-cluster/eda-node06.yaml
```

The configuration files created by the EDAADM tool are used in the next steps when you deploy the virtual machines. EDAADM creates a folder with the name of your cluster to store these files in. The output shows where the files are located.



Note: Nokia strongly recommends that you store these files securely and keep a backup.

4.3 Deploying the Talos virtual machines

This section provides the procedures for deploying an EDA node as a virtual machine on KVM or VMware vSphere.

4.3.1 Creating the VM on bridged networks on KVM

About this task

Complete the following steps to deploy an EDA node as a virtual machine on KVM. These steps are executed on the RedHat Enterprise Linux or Rocky Linux hypervisor directly. The steps below assume the deployment of the eda-node01 virtual machine as per the above configuration file. Ensure that you use the correct machine configuration file generated by the EDAADM tool.



Note: This procedure expects two networks to be available on the KVM hypervisors. The OAM network is referred to as br0 and the fabric management network is referred to as br1. Both of

these networks are standard Linux bridge networks. If you use only one interface, adapt Step 7 to only use the br0 network.

Procedure

Step 1. Ensure that the **virt-install** tool is installed on the KVM hypervisor.

If you need to install the tool, use the following command:

```
yum install virt-install
```

Step 2. Verify that the ISO image downloaded in [Downloading the KVM image](#) is available on the hypervisor.

Step 3. Copy the machine configuration file generated for this specific node to a file called `user-data`.

Example

```
cp eda-node01.yaml user-data
```

Step 4. Create a file called `meta-data` for the node.

Use the appropriate `instance-id` and `local-hostname` values.

Example

```
instance-id: eda-node01
local-hostname: eda-node01
```

Step 5. Create a file called `network-config` for the node.

Example

The file should have the following content:

```
version: 2
```

Step 6. Create an ISO file containing the newly created files.

For ease of use, name the ISO file with the name of the node for which you are creating the ISO.

Example

```
mkisofs -o eda-node01-data.iso -V cidata -J -r meta-data network-config user-data
```

Step 7. Create the virtual machine.

This step uses both the newly created ISO file and the ISO file downloaded from the Talos Machine Factory.

Example

```
virt-install -n eda-node01 \
  --description "Talos 1.8.3 vm for node eda-node01" \
  --noautoconsole --os-type=generic \
  --memory 65536 --vcpus 32 --cpu host \
  --disk eda-node01-rootdisk.qcow2,format=qcow2,bus=virtio,size=100 \
  --disk eda-node01-storagedisk.qcow2,format=qcow2,bus=virtio,size=300 \
  --cdrom nocloud-amd64.iso \
  --disk eda-node01-data.iso,device=cdrom \
  --network bridge=br0,model=virtio \
  --network bridge=br1,model=virtio
```



Note: If the node is not a storage node, you can remove the second `--disk` line.

4.3.2 Creating the VM on bridged networks on VMware vSphere

About this task

Complete the following steps to deploy an EDA node as a virtual machine on VMware vSphere. The steps below assume the deployment of the `eda-node01` virtual machine as per the above configuration file. Ensure that you are using the correct machine configuration file generated by the EDAADM tool. You can use one of the following methods to deploy the VM on VMware vSphere:

- the VMware vSphere vCenter or ESXi UI

For instructions, see *Deploy an OVF or OVA Template* in the VMware vSphere documentation.

- the VMware Open Virtualization Format Tool CLI

This procedure provides an example of how to use the VMware OVF Tool CLI.



Note: This procedure uses two networks (portgroups) to be available on the ESXi hypervisors. The OAM network is referred to as OAM and the fabric management network is referred to as FABRIC. Both of these networks can be standard PortGroups or distributed PortGroups. If you only use one network, you do not need to create a second interface on the VM.

Procedure

- Step 1.** Download and install the latest version of the VMware OVF Tool from the VMware Developer website.
- Step 2.** Display details about the OVA image.

Example

```
$ ovftool vmware-amd64.ova
OVF version: 1.0
VirtualApp: false
Name: talos

Download Size: 104.05 MB

Deployment Sizes:
Flat disks: 8.00 GB
Sparse disks: Unknown

Networks:
Name: VM Network
Description: The VM Network network

Virtual Machines:
Name: talos
Operating System: other3xlinux64guest
Virtual Hardware:
Families: vmx-15
Number of CPUs: 2
Cores per socket: automatic
Memory: 2.00 GB
```

```

Disks:
  Index:      0
  Instance ID: 4
  Capacity:   8.00 GB
  Disk Types: SCSI-VirtualSCSI

NICs:
  Adapter Type: VmxNet3
  Connection:   VM Network

Properties:
  Key:          talos.config
  Label:        Talos config data
  Type:         string
  Description:  Inline Talos config

References:
  File:        disk.vmdk

```

Step 3. Create a base64 encoded hash from the Talos machine configuration for the node.

In this example, the output is stored as an environment variable to make it easy to use in the command to deploy the image using the OVF Tool.

Example

```
export NODECONFIG=$(base64 -i eda-node01.yaml)
```

Step 4. Deploy the OVA image using the OVF Tool.

For details about command line arguments, see the OVF Tool documentation from the VMware website.

Example

```

$ ovftool --acceptAllEulas --noSSLVerify \
  -dm=thin \
  -ds=DATASTORE \
  -n=eda-node01 \
  --net:"VM Network=0AM" \
  --prop:talos.config="{NODECONFIG}" \
  vmware-amd64.ova \
  vi://administrator%40vsphere.local@vcenter.domain.tld/My-DC/host/My-Cluster/Resources/
  My-Resource-Group

Opening OVA source: vmware-amd64.ova
The manifest validates
Enter login information for target vi://vcenter.domain.tld/
Username: administrator%40vsphere.local
Password: *****
Opening VI target: vi://administrator%40vsphere.local@vcenter.domain.tld:443/My-DC/
host/My-Cluster/Resources/My-Resource-Group
Deploying to VI: vi://administrator%40vsphere.local@ vcenter.domain.tld:443/My-DC/
host/My-Cluster/Resources/My-Resource-Group
Transfer Completed
Completed successfully

```

Expected outcome

This step deploys the VM with the CPU, memory, disk, and NIC configuration of the default OVA image. The next step updates these settings.

Step 5. In vCenter, edit the VM settings.

Make the following changes:

- Increase the number of vCPU to 32.
- Increase the memory to 64G.
- Increase the main disk size to 100G. On boot, Talos automatically extends the file system.
- Optionally, if this VM is a storage node, add a new disk with a size of 300G.
- Optionally, add a second network interface and connect it to the FABRIC PortGroup.
- Enable 100% resource reservation for the CPU, memory and disk.

Step 6. Power on the virtual machine.

5 Bootstrap the Talos Kubernetes cluster

When all the virtual machines are deployed and running, you can set up the Kubernetes cluster on the virtual machines using Talos.



Note: The procedures in this chapter use the **edadm** command. Ensure that the command is available, as well as the original EDAADM configuration file from which you generated Talos files.

5.1 Bootstrapping Kubernetes on the primary node

About this task

After booting the Talos VMs, you can now bootstrap the Kubernetes cluster using the **talosconfig** command. The following parameter is relevant for this procedure:

-c: Specifies the **edadm** tool.

Procedure

Execute the following command in the folder where the `eda-input-6-node.yaml` is stored, as well as the folder where the Talos configuration files were executed:

```
edadm bootstrap-k8s -c eda-input-6-node.yaml
```

Expected outcome

Wait for several minutes for the Kubernetes cluster to come up and for all the nodes join the cluster. The process should take less than 15 minutes.

5.2 Obtaining the Kubernetes config file for kubectl

About this task

Use the **edadm** command to obtain the Kubernetes configuration file for use with **kubectl**. The following parameter is relevant for this procedure:

-c: Specifies your EDAADM configuration file.

Procedure

Step 1. Obtain the Kubernetes configuration file.

Execute the following command in the folder where the `eda-input-6-node.yaml` is stored, as well as where the folder with the Talos configuration files was executed:

Example

```
edadm get-kubeconfig -c eda-6-node-deployment.yaml
```

Expected outcome

The command creates a `kubeconfig` file in the folder with the name of your cluster according to the `EDAADM` config file.

Step 2. You can configure your environment to use the `kubeconfig` file for use with the `kubectl` command.

Example

```
export KUBECONFIG=eda-compute-cluster/kubeconfig
```

Step 3. Inspect your server and check if all nodes are up and running. You can use the typical `kubectl` commands.

Example

```
kubectl get nodes
```

Expected outcome

When all the nodes are up and Kubernetes is stable, continue with [Setting up the Rook Ceph storage cluster](#).

5.3 Setting up the Rook Ceph storage cluster

About this task

EDA uses Rook Ceph as a secure, distributed, and redundant data store for all the data it stores. Using Ceph guarantees redundancy and high availability of all data by providing multiple copies of all data. The following steps guide you through the configuration and deployment of Rook Ceph.



Note: To ensure that you deploy Rook Ceph on the correct Kubernetes cluster, verify that you are using the `kubeconfig` file created through the `edaadm` tool .

Procedure

Step 1. Add the Rook Ceph Helm chart.

Example

```
helm repo add rook-release https://charts.rook.io/release
```

Step 2. Using the `rook-ceph-operator-values.yaml` file that `edaadm` generated based on the configuration, deploy the Rook Ceph Operator.

Example

```
helm install --create-namespace \
  --namespace rook-ceph \
  -f path/to/rook-ceph-operator-values.yaml \
  rook-ceph rook-release/rook-ceph
```

Step 3. Using the `rook-ceph-cluster-values.yaml` file that the `edaadm` tool generated, deploy the Rook Ceph Cluster.

Example

```
helm install \
  --namespace rook-ceph \
  --set operatorNamespace=rook-ceph \
  -f path/to/rook-ceph-cluster-values.yaml \
  rook-ceph-cluster rook-release/rook-ceph-cluster
```

Expected outcome

The output from this command can report missing CRDs; wait until the Rook Ceph Operator is running in the Kubernetes cluster.

- Step 4.** Using **kubectl** commands, verify that the operator is deployed and the necessary pods are deployed before installing the EDA application.

Example

This example is for a six-node cluster, with six storage nodes.

```
$ kubectl -n rook-ceph get pods
```

NAME	READY	STATUS	RESTARTS
AGE			
csi-cephfsplugin-22rmj 7m6s	2/2	Running	1 (6m32s ago)
csi-cephfsplugin-25p9d 7m6s	2/2	Running	1 (6m30s ago)
csi-cephfsplugin-2gr8v 7m6s	2/2	Running	4 (5m16s ago)
csi-cephfsplugin-48cwk 7m6s	2/2	Running	1 (6m30s ago)
csi-cephfsplugin-fknch 7m6s	2/2	Running	2 (5m32s ago)
csi-cephfsplugin-provisioner-67c8454ddd-mpq4w 7m6s	5/5	Running	1 (6m1s ago)
csi-cephfsplugin-provisioner-67c8454ddd-qmdrq 7m6s	5/5	Running	1 (6m18s ago)
csi-cephfsplugin-vfxnf 7m6s	2/2	Running	1 (6m32s ago)
rook-ceph-mds-ceph-filesystem-a-7c54cdf5bc-lmf6n 2m40s	1/1	Running	0
rook-ceph-mds-ceph-filesystem-b-6dc794b9f4-2lc64 2m37s	1/1	Running	0
rook-ceph-mgr-a-55b449c844-wpps8 4m30s	2/2	Running	0
rook-ceph-mgr-b-5f97fd5746-fzngx 4m30s	2/2	Running	0
rook-ceph-mon-a-76fcb96c4c-vscnc 5m53s	1/1	Running	0
rook-ceph-mon-b-68bf5974bb-p2vnj 4m57s	1/1	Running	0
rook-ceph-mon-c-6d7c64dcb6-phs99 4m47s	1/1	Running	0
rook-ceph-operator-5f4c4bff8d-2fsq2 7m54s	1/1	Running	0
rook-ceph-osd-0-bf89f779-zh4kd 3m49s	1/1	Running	0
rook-ceph-osd-1-64dcd64c5f-7xcbm 3m49s	1/1	Running	0
rook-ceph-osd-2-54ddd95489-5qkdt 3m49s	1/1	Running	0
rook-ceph-osd-3-56cbd54bd6-7mt8w 3m39s	1/1	Running	0

rook-ceph-osd-4-567dcff476-wljll 2m56s	1/1	Running	0
rook-ceph-osd-5-6f69c998b6-2l5wp 2m54s	1/1	Running	0
rook-ceph-osd-prepare-eda-dev-node01-7rfkn 4m8s	0/1	Completed	0
rook-ceph-osd-prepare-eda-dev-node02-rqdkx 4m8s	0/1	Completed	0
rook-ceph-osd-prepare-eda-dev-node03-xtznb 4m8s	0/1	Completed	0
rook-ceph-osd-prepare-eda-dev-node04-db4v8 4m7s	0/1	Completed	0
rook-ceph-osd-prepare-eda-dev-node05-29wmm 4m7s	0/1	Completed	0
rook-ceph-osd-prepare-eda-dev-node06-zxp2x 4m7s	0/1	Completed	0
rook-ceph-tools-b9d78b5d4-8r62p 7m6s	1/1	Running	0



Note: Some of the pods may restart as they initiate Ceph. This behavior is expected.

6 Installing the EDA application

After setting up Kubernetes on the VMs, you can now install Nokia EDA using the playground.

6.1 Customizing the installation file

Update the settings for the parameters in the `prefs.mk` file to control the way EDA is installed.

Table 5: Customizable parameters in the `prefs.mk` file

Parameter	Description
NO_KIND	Specifies not to deploy the lab KIND node. Must be set to 1.
METALLB_VIP	Specifies the VIP address of your EDA deployment. Make sure to use a CIDR format, preferably as a /32 (or /128 for an IPv6 VIP). If you use two networks, this VIP address must be the one used on the fabric management network. If you use a single network, this setting must match the VIP address used for EXT_DOMAIN_NAME FQDN or IP. Example: 203.0.113.10/32
EXT_DOMAIN_NAME	The FQDN that resolves to the EDA VIP or the VIP itself. This setting must refer to the FQDN or VIP address that will be used to access the UI. If you use two networks, this value must be the FQDN or IP address of the OAM network.
EXT_HTTP_PORT	The HTTP port that the EDA UI/API should use to redirect to HTTPS. Set to 80.
EXT_HTTPS_PORT	The HTTPS port on which the EDA UI/API listens. Set to 443.
EXT_IPV4_ADDR	The IPv4 IP address used as the VIP address. If you use two networks, this VIP address must be the one used on the fabric management network. If you use a single network, this VIP address must be the VIP that matches your EXT_DOMAIN_NAME FQDN (or IP).
EXT_IPV6_ADDR	The IPv6 IP address used as the VIP. If you use two networks, this VIP address must be the one used on the fabric management network. If you use a single network, this VIP address must be the VIP that matches your EXT_DOMAIN_NAME FQDN (or IP).

Parameter	Description
HTTPS_PROXY	Optional: The proxy address for the HTTPS proxy.
HTTP_PROXY	Optional: The proxy address for the HTTP proxy.
NO_PROXY	Optional: The list of IP addresses, IP ranges and hostnames that should not be proxied.
https_proxy	Optional: The proxy address for the HTTPS proxy.
http_proxy	Optional: The proxy address for the HTTP proxy.
no_proxy	Optional: The list of IP addresses, IP ranges and hostnames that should not be proxied.
LLM_API_KEY	Optional: The OpenAI API key for the LLM functionality.
SINGLESTACK_SVCS	Optional: Indicates that internal services should be single stack instead of dual stack, if Kubernetes is dual stack. Boolean.
SIMULATE	Specifies if the EDA deployment is to manage simulated workloads (Digital Sandbox) or real hardware. Values: <code>true</code> or <code>false</code> . By default, this parameter is set to <code>true</code> if the parameter is not provided in the file.
KPT_SETTERS_FILE	Advanced configuration file for kpt.

6.1.1 The prefs.mk file

Example: Settings for configurable parameters in the prefs.mk file



Note: The settings for METALLB_VIP and EXT_IPV4_ADDR match the fabric management network.

```
NO_KIND=1
METALLB_VIP=203.0.113.10/32
EXT_DOMAIN_NAME=eda.domain.tld
EXT_HTTP_PORT=80
EXT_HTTPS_PORT=443
EXT_IPV4_ADDR=203.0.113.10
EXT_IPV6_ADDR=""
HTTPS_PROXY=http://192.0.2.254:8080
HTTP_PROXY=http://192.0.2.254:8080
NO_PROXY=192.0.2.0/24,203.0.113.0/24,.domain.tld,172.22.0.0/16,localhost,127.0.0.1,10.0.1.0/24,0.0.0.0,169.254.116.108
https_proxy=http://192.0.2.254:8080
http_proxy=http://192.0.2.254:8080
no_proxy=192.0.2.0/24,203.0.113.0/24,.domain.tld,172.22.0.0/16,localhost,127.0.0.1,10.0.1.0/24,0.0.0.0,169.254.116.108
LLM_API_KEY=...
```

6.2 Installing Nokia EDA

About this task

Follow these steps to install EDA.

Procedure

Step 1. Download the latest tools.

Example

```
make download-tools
```

Step 2. Download the latest packages, including the eda-kpt package.

Example

```
make download-pkgs
```

Step 3. Set up the MetalLB environment for VIP management.

Example

```
make metallb
```

Step 4. Install the necessary external packages.

Example

```
make install-external-packages
```



Note: If this command exits with an error, wait 30 seconds and try again. Sometimes Kubernetes is a bit slower in reconciling the change than the command waits for.

Step 5. Change the eda-git Kubernetes service to a ClusterIP service instead of a LoadBalancer type.

Example

```
kubectl -n eda-system patch service eda-git -p '{"spec": {"type": "ClusterIP"}}'
```

Step 6. Generate the EDA core configuration.

Example

```
make eda-configure-core
```

Step 7. Install EDA core components.

Example

```
make eda-install-core
```



Note: If the command hangs for a long time (>5 minutes) on "reconcile pending" for a workflowdefinition, cancel the command and try again, KPT is designed to handle these cases. This can occasionally happen depending on the Kubernetes cluster

Step 8. Verify that the EDA Config Engine is up and running.

Example

```
make eda-is-core-ready
```

Step 9. Install all the standard EDA apps.

This step can take approximate 5 to 15 minutes, depending on your connectivity.

Example

```
make eda-install-apps
```

Step 10. Bootstrap EDA.

Example

```
make eda-bootstrap
```

Step 11. If your deployment uses two networks, create a second VIP pool for the OAM VIP address.

Example

```
make metallb-configure-pools METALLB_VIP=<OAM VIP> LB_POOL_NAME=pool-nb
```

Step 12. If your deployment uses two networks, create the OAM UI/API service using the new VIP pool.

Example

```
make eda-create-api-lb-svc API_LB_POOL_NAME=pool-nb
```

Step 13. Optional: Deploy an example topology.

Example

```
make topology-load
```


6.3 Accessing the EDA deployment

Procedure

You can now access the new EDA deployment using the following methods:

- use `https://OAM-VIP`
- if an FQDN is configured for the `EXT_DOMAIN_NAME` field, use `https://FQDN`

Customer document and product support



Customer documentation

[Customer documentation welcome page](#)



Technical support

[Product support portal](#)



Documentation feedback

[Customer documentation feedback](#)