

# Clusters

---

---

## Overview

[Overview](#)

---

## Immutable Infrastructure

[Immutable Infrastructure](#)

---

## Cluster Node Planning

[Cluster Node Planning](#)

---

## Creating an On-Premise Cluster

[Creating an On-Premise Cluster](#)

---

## etcd Encryption

[etcd Encryption](#)

---

## How to

[Add External Address for Built-in Registry](#)

[Choosing a Container Runtime](#)

[Updating Public Repository Credentials](#)

---

# Overview

A cluster is the foundational resource collection for running containerized applications, encompassing nodes, load balancers, storage, and other critical components. It is a prerequisite for successfully running containerized applications on the platform. During initial platform installation, a standard Kubernetes cluster, known as the `global` cluster, is created. Subsequently, multiple clusters can be integrated into the `global` cluster for unified management.

## TOC

### Cluster Type

- On-Premises Cluster

- Managed Cluster

### Multi-Cloud and Hybrid Cloud Support

### Implementation Considerations and Limitations

- Version Compatibility

- Network and Security Requirements

### Best Practices for Cluster Management

1. Pre-Implementation Assessment
2. Security and Compliance
3. Monitoring and Observability
4. Backup and Disaster Recovery
5. Continuous Optimization

# Cluster Type

## On-Premises Cluster

On-Premises cluster is Kubernetes clusters directly created by the platform. Users provide virtual or physical machines, and the platform installs and configures Kubernetes clusters on these machines. This approach is suitable for enterprises with existing hardware resources, allowing full utilization of infrastructure.

## Managed Cluster

Managed cluster is Kubernetes clusters provided by cloud service providers, which are integrated into the platform for unified management. Supported integration methods include:

Method	Description	Use Case	Key Characteristics
Import	Integrating existing Kubernetes clusters	Existing clusters with direct network access	<ul style="list-style-type: none"><li>Cluster information submitted to <code>global</code> cluster</li><li><code>global</code> cluster must have network access to the cluster</li></ul>
Register	Integrating clusters with strict security requirements	Clusters with high security constraints	<ul style="list-style-type: none"><li>Specific plugins installed on the target cluster</li><li>Reverse proxy establishes a secure tunnel</li><li>Maintains cluster security while enabling management</li></ul>

Method	Description	Use Case	Key Characteristics
Proxy Create	Creating clusters through cloud service providers	Leveraging public cloud Kubernetes services	<ul style="list-style-type: none"><li>• Cloud service provider credentials required</li><li>• Platform creates Kubernetes clusters using provided credentials</li></ul>

## Multi-Cloud and Hybrid Cloud Support

These cluster management approaches meet enterprise needs in multi-cloud and hybrid cloud scenarios, supporting container transformation at different stages:

- Existing Hardware: Create platform-provided clusters
- Existing Clusters: Import or register into the platform
- Elastic Demands: Quickly create public cloud clusters

## Implementation Considerations and Limitations

### Version Compatibility

- Supported Kubernetes versions: 1.28, 1.29, 1.30, 1.31
- Both On-Premises and Managed clusters must ensure version compatibility
- Version mismatches may result in feature limitations or compatibility issues

### Network and Security Requirements

- Ensure network connectivity between `global` and target clusters
- Implement appropriate firewall and network security policies

- Manage access credentials and authentication mechanisms securely

# Best Practices for Cluster Management

## 1. Pre-Implementation Assessment

- Conduct thorough infrastructure and workload analysis
- Identify specific requirements for each cluster
- Develop a comprehensive migration and integration strategy

## 2. Security and Compliance

- Implement role-based access control (RBAC)
- Use network policies to restrict cluster communication
- Regularly audit and update security configurations
- Ensure compliance with industry standards and regulations

## 3. Monitoring and Observability

- Set up centralized logging and monitoring
- Implement proactive alerting mechanisms
- Use platform-provided observability tools
- Track cluster performance, resource utilization, and health

## 4. Backup and Disaster Recovery

- Establish regular backup procedures
- Create and test disaster recovery plans
- Implement multi-cluster backup strategies
- Ensure minimal downtime and data loss

## 5. Continuous Optimization

- Regularly review cluster configurations
- Optimize resource allocation
- Update to the latest supported Kubernetes versions
- Leverage platform features for automatic updates and scaling

# About Immutable Infrastructure

Immutable Infrastructure uses an immutable operating system to provision Kubernetes clusters. Unlike traditional OS-based clusters, all node configurations are baked into images and remain unchanged after deployment. Cluster upgrades and configuration changes are applied by replacing nodes with new images, ensuring consistency, reliability, and simplified operations throughout the cluster lifecycle.

## Note

Because Immutable Infrastructure releases on a different cadence from Alauda Container Platform, the Immutable Infrastructure documentation is now available as a separate documentation set at [Immutable Infrastructure](#) ↗.



# Cluster Node Planning

A cluster utilizes the Kubernetes node role labels `node-role.kubernetes.io/<role>` to assign different roles to nodes. For convenience of description, we refer to this type of label as a role label.

By default, a cluster contains two types of nodes: control plane nodes and worker nodes, used to host control plane workloads and application workloads, respectively.

In a cluster:

- The control plane nodes are labeled with the role label `node-role.kubernetes.io/control-plane`.

**Note:**

Prior to Kubernetes v1.24, the community also used the label `node-role.kubernetes.io/master` to mark control plane nodes. For backward compatibility, both labels are considered valid for identifying control plane nodes.

- The worker nodes, by default, have no role labels. However, you can explicitly assign the role label `node-role.kubernetes.io/worker` to a worker node if desired.

In addition to these default role labels, you can also define custom role labels on worker nodes to further classify them into different functional types. For example:

- You can add the role label `node-role.kubernetes.io/infra` to designate a node as an infra node, intended for hosting infrastructure components.
- You can add the role label `node-role.kubernetes.io/log` to designate a node as a log node, specialized for hosting logging components.

This document will guide you through creating infra nodes and custom role nodes, and migrating workloads to those nodes.

# TOC

## Creating Infra Nodes on Non-Immutable Cluster

### Adding Infra Nodes

Step 1: Add the Infra Role Label to the Node resources

Step 2: Add a Taint to the Node resources

Step 3: Verify the Label and Taint

## Migrating Pods to Infra Nodes

## Custom Node Planning

### General Steps for Defining Custom Role Nodes

Step 1: Add a Custom Role Label

Step 2: Add a Corresponding Taint

Step 3: Verify the Configuration

### Example: Create A Node Dedicated To Logging Components

Step 1: Add the Log Role Label

Step 2: Add a Taint to the Node

Step 3: Verify the Label and Taint

---

## Creating Infra Nodes on Non-Immutable Cluster

By default, a cluster only includes control plane nodes and worker nodes. If you want to designate certain worker nodes as infra nodes dedicated to hosting infrastructure components, you need to manually add the appropriate role label and taint to those nodes.

### Note:

The operations in this section are only applicable to non-immutable clusters. That is, the following operations are not supported on cloud clusters (such as EKS managed clusters deployed via the `Alauda Container Platform EKS Provider` Cluster Plugin), third-party clusters, or clusters where the nodes use an immutable OS.

# Adding Infra Nodes

## Step 1: Add the Infra Role Label to the Node resources

```
kubectl label nodes 192.168.143.133 node-role.kubernetes.io/infra="" --overwrite
```

This command adds the infra role label to the Node 192.168.143.133: `node-role.kubernetes.io/infra: ""`, indicating that the node is an infra node.

## Step 2: Add a Taint to the Node resources

Add a taint to prevent other workloads from being scheduled onto the infra node.

```
kubectl taint nodes 192.168.143.133 node-role.kubernetes.io/infra=reserved:NoSchedule
```

This command adds the taint `node-role.kubernetes.io/infra=reserved:NoSchedule` to Node 192.168.143.133, indicating that only applications that tolerate this taint can be scheduled onto this node.

## Step 3: Verify the Label and Taint

Check whether the node has been assigned the infra role label and taint:

```
# kubectl describe node 192.168.143.133
Name:          192.168.143.133
Roles:         infra
Labels:        node-role.kubernetes.io/infra=reserved
...
Taints:        node-role.kubernetes.io/infra=reserved:NoSchedule
```

The output indicates that the Node 192.168.143.133 has been configured as an infra node and has been tainted with tainted with `node-role.kubernetes.io/infra=reserved:NoSchedule`.

# Migrating Pods to Infra Nodes

If you want to schedule specific Pod onto infra nodes, you need to make the following configurations:

- A nodeSelector targeting the infra role label.
- Corresponding tolerations for the infra node's taint.

Below is an example Deployment manifest configured to run on the infra node.

```
apiVersion: apps/v1
kind: Pod
metadata:
  name: infra-pod-demo
  namespace: default
spec:
  ...
  nodeSelector:
    node-role.kubernetes.io/infra: ""
  tolerations:
  - effect: NoSchedule
    key: node-role.kubernetes.io/infra
    value: reserved
    operator: Equal
  ...
```

The nodeSelector ensures the Pod is only scheduled on nodes with the label `node-role.kubernetes.io/infra: ""`, the toleration allows the Pod to tolerate the taint `node-role.kubernetes.io/infra=reserved:NoSchedule`.

With these configurations, the Pod will be scheduled onto the infra node.

**Note:**

Moving pods installed via OLM Operators or Cluster Plugins to an infra node is not always possible. The capability to move these pods is depends on the configuration of each Operator or Cluster Plugin.

## Custom Node Planning

Beyond infra nodes, you may want to designate worker nodes for other specialized purposes — such as hosting logging components, storage services, or monitoring agents.

You can achieve this by assigning more custom role labels and corresponding taints to worker nodes, effectively turning them into custom role nodes.

## General Steps for Defining Custom Role Nodes

The process is similar to creating infra nodes.

### Step 1: Add a Custom Role Label

```
kubectl label nodes <node> node-role.kubernetes.io/<role>="" --overwrite
```

Replace <role> with your desired role name, such as monitoring, storage, or log.

### Step 2: Add a Corresponding Taint

```
kubectl taint nodes <node> node-role.kubernetes.io/<role>=<value>:NoSchedule
```

Replace <role> with your custom role name and replace <value> with a meaningful descriptor, such as reserved or dedicated. This value is optional but useful for documentation and clarity.

### Step 3: Verify the Configuration

```
kubectl describe node <node>
```

Ensure the Labels and Taints fields reflect your custom role configuration.

## Example: Create A Node Dedicated To Logging Components

If you want to create a node specifically for installing logging components, you can add the log role. In this case, create the log node as follows.

## Step 1: Add the Log Role Label

```
kubectl label nodes 192.168.143.133 node-role.kubernetes.io/log="" --overwrite
```

This label indicates that the node is designated for log-related workloads.

## Step 2: Add a Taint to the Node

```
kubectl taint nodes 192.168.143.133 node-role.kubernetes.io/log=reserved:NoSchedule
```

This taint prevents unscheduled workloads from being deployed to the node.

## Step 3: Verify the Label and Taint

```
Name:          192.168.143.133
Roles:         log
Labels:        node-role.kubernetes.io/log=reserved
               ...
Taints:        node-role.kubernetes.io/log=reserved:NoSchedule
```

This confirms that the node has been successfully configured as a log node with the appropriate label and taint.

By following the above practices, you can effectively partition your Kubernetes nodes based on their intended purpose, improve workload isolation, and ensure that specific components are deployed onto appropriately configured nodes.

# Creating an On-Premise Cluster

---

## TOC

### Prerequisites

Node Requirements

Load Balancing

Connecting global Cluster and Workload Cluster

Image Registry

Container Networking

### Creation Procedure

Basic Info

Container Network

Node Settings

Extended Parameters

### Post-Creation Steps

Viewing Creation Progress

Associating with Projects

---

## Prerequisites

### Node Requirements

1. If you downloaded a single-architecture installation package from [Download Installation Package](#), ensure your node machines have the same architecture as the package. Otherwise, nodes won't start due to missing architecture-specific images.
2. Verify that your node operating system and kernel are supported. See [Supported OS and Kernels](#) for details.
3. Perform availability checks on node machines. For specific check items, refer to [Node Preprocessing > Node Checks](#).
4. If node machine IPs cannot be directly accessed via SSH, provide a SOCKS5 proxy for the nodes. The `global` cluster will access nodes through this proxy service.

## Load Balancing

For production environments, a load balancer is required for cluster control plane nodes to ensure high availability. You can provide your own hardware load balancer or enable `Self-built VIP`, which provides software load balancing using haproxy + keepalived. We recommend using a hardware load balancer because:

- **Better Performance:** Hardware load balancing performs better than software load balancing.
- **Lower Complexity:** If you're unfamiliar with keepalived, misconfigurations could make the cluster unavailable, leading to lengthy troubleshooting and seriously affecting cluster reliability.

When using your own hardware load balancer, you can use the load balancer's VIP as the `IP Address / Domain` parameter. If you have a domain name that resolves to the load balancer's VIP, you can use that domain as the `IP Address / Domain` parameter. Note:

- The load balancer must correctly forward traffic to ports `6443`, `11780`, and `11781` on all control plane nodes in the cluster.
- If your cluster has only one control plane node and you use that node's IP as the `IP Address / Domain` parameter, the cluster cannot be scaled from a single node to a highly available multi-node setup later. Therefore, we recommend providing a load balancer even for single-node clusters.

When enabling `Self-built VIP`, you need to prepare:

1. An available VRID



2. A host network that supports the VRRP protocol
3. All control plane nodes and the VIP must be on the same subnet, and the VIP must be different from any node IP.

## Connecting `global` Cluster and Workload Cluster

The platform requires mutual access between the `global` cluster and workload clusters. If they're not on the same network, you need to:

1. Provide `External Access` for the workload cluster to ensure the `global` cluster can access it. Network requirements must ensure `global` can access ports `6443` , `11780` , and `11781` on all control plane nodes.
2. Add an additional address to `global` that the workload cluster can access. When creating a workload cluster, add this address to the cluster's annotations with the key `cpaas.io/platform-url` and the value set to the public access address of `global` .

## Image Registry

Cluster images support Platform Built-in, Private Repository, and Public Repository options.

- **Platform Built-in:** Uses the image registry provided by the `global` cluster. If the cluster cannot access `global` , see [Add External Address for Built-in Registry](#).
- **Private Repository:** Uses your own image registry. For details on pushing required images to your registry, contact technical support.
- **Public Repository:** Uses the platform's public image registry. Before using, complete [Updating Public Repository Credentials](#).

## Container Networking

If you plan to use Kube-OVN's Underlay for your cluster, refer to [Preparing Kube-OVN Underlay Physical Network](#).

## Creation Procedure

1. Enter the **Administrator** view, and click **Clusters/Clusters** in the left navigation bar.
2. Click **Create Cluster**.
3. Configure the following sections according to the instructions below: Basic Info, Container Network, Node Settings, and Extended Parameters.

## Basic Info

Parameter	Description
<b>Kubernetes Version</b>	<p>All optional versions are rigorously tested for stability and compatibility.</p> <p><b>Recommendation:</b> Choose the latest version for optimal features and support.</p>
<b>Container Runtime</b>	<p>Containerd is provided as the default container runtime.</p> <p>If you prefer using Docker as the container runtime, please refer to <a href="#">Choosing a Container Runtime</a>.</p>
<b>Cluster Network Protocol</b>	<p>Supports three modes: IPv4 single stack, IPv6 single stack, IPv4/IPv6 dual stack.</p> <p><b>Note:</b> If you select dual stack mode, ensure all nodes have correctly configured IPv6 addresses; the network protocol cannot be changed after setting.</p>

## Cluster Endpoint

**IP Address / Domain** : Enter the pre-prepared domain name or VIP if no domain name is available.

**Self-Built VIP** : Disabled by default. Only enable if you haven't provided a LoadBalancer. When enabled, the installer will automatically deploy **keepalived** for software load balancing support.

**External Access** : Enter the externally accessible address prepared for the cluster when it's not in the same network environment as the **global** cluster.

## Container Network

### Kube-OVN

An enterprise-grade Cloud Native Kubernetes container network orchestration system developed by Alauda. It brings mature networking capabilities from the OpenStack domain to Kubernetes, supporting cross-cloud network management, traditional network architecture and infrastructure interconnection, and edge cluster deployment scenarios, while greatly enhancing Kubernetes container network security, management efficiency, and performance.

Parameter	Description
<b>Subnet</b>	Also known as Cluster CIDR, represents the <b>default subnet</b> segment. After cluster creation, additional subnets can be added.
<b>Transmit Mode</b>	<b>Overlay</b> : A virtual network abstracted over the infrastructure that doesn't consume physical network resources. When creating an

	<p>Overlay default subnet, all Overlay subnets in the cluster use the same cluster NIC and node NIC configuration.</p> <p><b>Underlay:</b> This transmission method relies on physical network devices. It can directly allocate physical network addresses to Pods, ensuring better performance and connectivity with the physical network. Nodes in an Underlay subnet must have multiple NICs, and the NIC used for bridge networking must be exclusively used by Underlay and not carry other traffic like SSH. When creating an Underlay default subnet, the cluster NIC is actually a default NIC for bridge networking, and the node NIC is the node NIC configuration in the bridge network.</p> <ul style="list-style-type: none"> <li>• <b>Default Gateway:</b> The physical network gateway address, which is the gateway address for the Cluster CIDR segment (must be within the Cluster CIDR address range).</li> <li>• <b>VLAN ID:</b> Virtual LAN identifier (VLAN number), e.g., 0.</li> <li>• <b>Reserved IPs:</b> Set reserved IPs that won't be automatically allocated, such as IPs in the subnet that are already used by other devices.</li> </ul>
<b>Service CIDR</b>	IP address range used by Kubernetes Services of type ClusterIP. Cannot overlap with the default subnet range.
<b>Join CIDR</b>	In Overlay transmission mode, this is the IP address range used for communication between nodes and pods. Cannot overlap with the default subnet or Service CIDR.

### Calico

Calico is a layer 3 networking solution that provides secure network connections for containers.

Parameter	Description
-----------	-------------

<b>Default Subnet</b>	Also known as Cluster CIDR, represents the <b>default subnet</b> segment. After cluster creation, additional subnets can be added.
<b>Service CIDR</b>	IP address range used by Kubernetes Services of type ClusterIP. Cannot overlap with the default subnet range.

## Flannel

Flannel provides a flat network environment for all containers in the cluster, giving containers created on different node hosts a unique virtual IP address across the entire cluster. The pod subnet is divided evenly among the cluster nodes according to the mask, and pods on each node are assigned IP addresses from the segment allocated to that node. This improves communication efficiency between containers without having to consider IP translation issues.

Parameter	Description
<b>Cluster CIDR</b>	<p>IP address range used by pods created when the cluster starts. Supports setting the maximum number of IP addresses that can be allocated to pods on each node under the current container network.</p> <p><b>Note:</b> The platform will automatically calculate the maximum number of nodes the cluster can accommodate based on the above configuration and display it in the hint below the input field.</p> <p><b>Important:</b> After cluster creation, the cluster network cannot be changed, so please plan the network carefully.</p>

**Service  
CIDR**

IP address range used by Kubernetes Services of type ClusterIP.  
Cannot overlap with the container subnet range.

**Custom**

If you need to install other network plugins, select **Custom** mode. You can manually install network plugins after the cluster is successfully created.

Parameter	Description
<b>Cluster CIDR</b>	IP address range used by pods created when the cluster starts.
<b>Service CIDR</b>	IP address range used by Kubernetes Services of type ClusterIP. Cannot overlap with the container subnet range.

## Node Settings

Parameter	Description
<b>Network Interface Card</b>	<p>The name of the host network interface device used by the cluster network plugin.</p> <p><b>Note:</b></p> <ul style="list-style-type: none"><li>When selecting Underlay transmission mode for the Kube-OVN default subnet, you must specify the network interface name, which will be the default NIC for bridge networking.</li><li>The platform's network interface traffic monitoring by default recognizes traffic on interfaces named like <code>eth. en. wl. ww.</code> . If you use interfaces with different naming conventions, please refer to</li></ul>

	<p><a href="#">Collect Network Data from Custom-Named Network Interfaces</a> after cluster onboarding to modify the relevant resources and ensure the platform can properly monitor network interface traffic.</p>
<b>Node Name</b>	<p>You can choose to use either the node IP or hostname as the node name on the platform.</p> <p><b>Note:</b> When choosing to use hostname as the node name, ensure that the hostnames of nodes added to the cluster are unique.</p>
<b>Nodes</b>	<p><b>Add nodes</b> to the cluster, or <b>Recovery from draft</b> temporarily saved node information. See the detailed parameter descriptions for adding nodes below.</p>
<b>Monitoring Type</b>	<p>Supports <b>Prometheus</b> and <b>VictoriaMetrics</b>.</p> <p>When selecting <b>VictoriaMetrics</b> as the monitoring component, you must configure the <b>Deploy Type</b>:</p> <ul style="list-style-type: none"><li>- <b>Deploy VictoriaMetrics</b>: Deploys all related components, including <b>VMStorage</b>, <b>VMAAlert</b>, <b>VMAgent</b>, etc.</li><li>- <b>Deploy VictoriaMetrics Agent</b>: Only deploys the log collection component, <b>VMAgent</b>. When using this deployment method, you need to associate with a VictoriaMetrics instance already deployed on another cluster in the platform to provide monitoring services for the cluster.</li></ul>
<b>Monitoring Nodes</b>	<p>Select nodes for deploying cluster monitoring components. Supports selecting compute nodes and control plane nodes that allow application deployment.</p>

To avoid affecting cluster performance, it's recommended to prioritize compute nodes. After the cluster is successfully created, monitoring components with storage type **Local Volume** will be deployed on the selected nodes.

## Node Addition Parameters

Parameter	Description
Type	<p><b>Control Plane Node:</b> Responsible for running components such as kube-apiserver, kube-scheduler, kube-controller-manager, etcd, container network, and some platform management components in the cluster. When <b>Application Deployable</b> is enabled, control plane nodes can also be used as compute nodes.</p> <p><b>Worker Node:</b> Responsible for hosting business pods running on the cluster.</p>
IPv4 Address	The IPv4 address of the node. For clusters created in internal network mode, enter the node's <b>private IP</b> .
IPv6 Address	Valid when the cluster has IPv4/IPv6 dual stack enabled. The IPv6 address of the node.
Application Deployable	Valid when <b>Node Type</b> is <b>Control Plane Node</b> . Whether to allow business applications to be deployed on this control plane node, scheduling business-related pods to this node.
Display Name	The display name of the node.



<b>SSH Connection IP</b>	<p>The IP address that can connect to the node when accessing it via SSH service.</p> <p>If you can log in to the node using <code>ssh &lt;username&gt;@&lt;node's IPv4 address&gt;</code> , this parameter is not required; otherwise, enter the node's public IP or NAT external IP to ensure the <code>global</code> cluster and proxy can connect to the node via this IP.</p>
<b>Network Interface Card</b>	<p>Enter the name of the network interface used by the node. The priority of network interface configuration effectiveness is as follows (from left to right, in descending order):</p> <p><b>Kube-OVN Underlay:</b> Node NIC &gt; Cluster NIC</p> <p><b>Kube-OVN Overlay:</b> Node NIC &gt; Cluster NIC &gt; NIC corresponding to the node's default route</p> <p><b>Calico:</b> Cluster NIC &gt; NIC corresponding to the node's default route</p> <p><b>Flannel:</b> Cluster NIC &gt; NIC corresponding to the node's default route</p>
<b>Associated Bridge Network</b>	<p><b>Note:</b> When creating a cluster, bridge network configuration is not supported; this option is only available when <b>adding nodes</b> to a cluster that already has Underlay subnets created.</p>

	<p>Select an existing <a href="#">Add Bridge Network</a>. If you don't want to use the bridge network's default NIC, you can configure the node NIC separately.</p>
<b>SSH Port</b>	SSH service port number, e.g., <code>22</code> .
<b>SSH Username</b>	SSH username, needs to be a user with root privileges, e.g., <code>root</code> .
<b>Proxy</b>	<p>Whether to access the node's SSH port through a proxy. When the <code>global</code> cluster cannot directly access the node to be added via SSH (e.g., the <code>global</code> cluster and workload cluster are not in the same subnet; the node IP is an internal IP that the <code>global</code> cluster cannot directly access), this switch needs to be turned on and proxy-related parameters configured. After configuring the proxy, node access and deployment can be achieved through the proxy.</p> <p><b>Note:</b> Currently, only SOCKS5 proxy is supported.</p> <p><b>Access URL:</b> Proxy server address, e.g., <code>192.168.1.1:1080</code>.</p> <p><b>Username:</b> Username for accessing the proxy server.</p> <p><b>Password:</b> Password for accessing the proxy server.</p>
<b>SSH Authentication</b>	Authentication method and corresponding authentication information for logging into the added node. Options include:

	<p><b>Password:</b> Requires a username with root privileges and the corresponding <b>SSH password</b>.</p> <p><b>Key:</b> Requires a <b>private key</b> with root privileges and the <b>private key password</b> .</p>
<b>Save Draft</b>	<p>Saves the currently configured data in the dialog as a draft and closes the <b>Add Node</b> dialog.</p> <p>Without leaving the <b>Create Cluster</b> page, you can select <b>Restore from draft</b> to open the <b>Add Node</b> dialog and restore the configuration data saved as a draft.</p> <p><b>Note:</b> The data restored from the draft is the most recently saved draft data.</p>

## Extended Parameters

### Note:

- Apart from required configurations, it's not recommended to set extended parameters, as incorrect settings may make the cluster unavailable and cannot be modified after cluster creation.
- If a entered **Key** duplicates a default parameter **Key**, it will override the default configuration.

### Procedure

1. Click **Extended Parameters** to expand the extended parameter configuration area. You can optionally set the following extended parameters for the cluster:

Parameter	Description
-----------	-------------

<b>Docker Parameters</b>	<code>dockerExtraArgs</code> , additional configuration parameters for Docker, which will be written to <code>/etc/sysconfig/docker</code> . Modification is not recommended. To configure Docker through the <code>daemon.json</code> file, it must be configured as key-value pairs.
<b>Kubelet Parameters</b>	<p><code>kubeletExtraArgs</code> , additional configuration parameters for Kubelet.</p> <p><b>Note:</b> When the <b>Container Network's Node IP Count</b> parameter is entered, a default <b>Kubelet Parameter</b> configuration with the key <code>max-pods</code> and a value of <b>Node IP Count</b> is automatically generated. This sets the maximum number of pods that can run on any node in the cluster. This configuration is not displayed in the interface.</p> <p>Adding a new <code>max-pods: maximum number of runnable pods</code> key-value pair in the <b>Kubelet Parameters</b> area will override the default value. Any positive integer is allowed, but it's recommended to use the default value (Node IP Count) or enter a value not exceeding <code>256</code> .</p>
<b>Controller Manager Parameters</b>	<code>controllerManagerExtraArgs</code> , additional configuration parameters for the Controller Manager.
<b>Scheduler Parameters</b>	<code>schedulerExtraArgs</code> , additional configuration parameters for the Scheduler.
<b>APIServer Parameters</b>	<code>apiServerExtraArgs</code> , additional configuration parameters for the APIServer.

<b>APIServer URL</b>	<code>publicAlternativeNames</code> , APIServer access addresses issued in the certificate. Only IPs or domain names can be entered, with a maximum of 253 characters.
<b>Cluster Annotations</b>	Cluster annotation information, marking cluster characteristics in metadata in the form of key-value pairs for platform components or business components to obtain relevant information.

4. Click **Create**. You'll return to the cluster list page where the cluster will be in the **Creating** state.

## Post-Creation Steps

### Viewing Creation Progress

On the cluster list page, you can view the list of created clusters. For clusters in the **Creating** state, you can check the execution progress.

#### Procedure

1. Click the small icon **View Execution Progress** to the right of the cluster status.
2. In the execution progress dialog that appears, you can view the cluster's execution progress (status.conditions).

**Tip:** When a certain type is in progress or in a failed state with a reason, hover your cursor over the corresponding reason (shown in blue text) to view detailed information about the reason (status.conditions.reason).

### Associating with Projects

After the cluster is created, you can add it to projects in the project management view.

# etcd Encryption

This guide helps you install, understand, and operate the etcd Encryption Manager in ACP to automate etcd data encryption key rotation within your clusters.

It ensures that sensitive data stored in etcd, such as secrets and configmaps, is encrypted using a secure algorithm, enhancing your cluster's security.

---

## TOC

Installation

How it Works

Default Configuration

Operations Guide

Configuration Files

Checking Status

---

## Installation

See [Cluster Plugin](#) for installation instructions.

### Note:

- Currently supported:
  - On-Premises clusters

- DCS clusters
- Not supported:
  - `global cluster`

## How it Works

Upon installation, an `etcd-encryption-manager` controller is deployed in the `kube-system` namespace, which:

- Periodically rotates etcd data encryption keys.
- Retains the 8 most recent keys for rollback compatibility.
- Updates encryption configurations on all control nodes.
- Triggers `kube-apiserver` to hot reload new keys.
- Automatically migrates resources to re-encrypt data with new keys.

Cluster stability is maintained throughout these operations.

## Default Configuration

Parameter	Value
Encrypted resources	secrets, configmaps
Encryption algorithm	256-bit AES-GCM
Rotation interval	168 hours (7 days)

## Operations Guide

## Configuration Files

Path	Content
<code>/etc/kubernetes/encryption-provider.conf</code>	Current encryption configuration
<code>/etc/kubernetes/encryption-provider-history.bak</code>	Historical key records (for recovery)
<code>/etc/kubernetes/encryption-provider-bak/</code>	Expired encryption configuration versions

## Checking Status

Run the following command to check the current rotation status:

```
kubectl get EtcdEncryptionConfig default -o yaml
```

Example output:



```
apiVersion: cluster.alauda.io/v1alpha1
kind: EtcdEncryptionConfig
metadata:
  name: default
spec:
  resources:
    - secrets
    - configmaps
  rotationInterval: 168h0m0s
  type: aesgcm
status:
  deployStatus:
    192.168.100.1:
      revision: 3
      state: Success
    192.168.100.2:
      revision: 3
      state: Success
    192.168.100.3:
      revision: 3
      state: Success
  migration:
    completeTimestamp: "2025-05-27T05:47:01Z"
    resources:
      - secrets
      - configmaps
    revision: 3
    state: Success
  revision: 3
```

# How to

---

[Add External Address for Built-in Registry](#)

[Choosing a Container Runtime](#)

[Updating Public Repository Credentials](#)

# Add External Address for Built-in Registry

## TOC

[Overview](#)[Prerequisites](#)[Procedure](#)[Configure Certificate and Routing Rules for the Platform Registry](#)

## Overview

When the `global` cluster uses the `Platform Built-in` registry, workload clusters typically also use this registry to pull images. The registry not only serves components within the `global` cluster but must also be accessible to workload cluster nodes.

In certain scenarios, workload cluster nodes cannot directly access the `global` cluster's registry address - for example, when the `global` cluster is in a private data center while workload clusters are in public clouds or edge environments.

This guide explains how to configure an externally accessible address for the platform's default registry to allow workload clusters to pull images.

## Prerequisites

Before you begin, prepare the following:

- A domain name accessible by workload cluster nodes
- The IP address that the domain name points to
- A valid SSL certificate for the domain name

### WARNING

- The domain name must be different from the platform access address
- Ensure the domain's IP address can forward traffic to all control plane nodes of the `global` cluster

## Procedure

### Configure Certificate and Routing Rules for the Platform Registry

1. Copy the domain's valid certificate to any control plane node of the `global` cluster
2. Create a TLS secret containing the domain certificate:

```
kubectl create secret tls registry-address.tls --cert=<certificate-filename> --key=<key-filename> -n kube-system
```

Example:

```
kubectl create secret tls registry-address.tls --cert=custom.crt --key=custom.key -n kube-system
```

**Note:** After creating the certificate, monitor the expiration date of the **registry-address.tls** secret in the **kube-system** namespace of the `global` cluster. Replace the certificate before it expires.

3. Create ingress rules on any control plane node of the `global` cluster:

```

REGISTRY_DOMAIN_NAME=<www.registry.com> # Replace with your accessible domain name
cat << EOF | kubectl create -f -
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  annotations:
    nginx.ingress.kubernetes.io/backend-protocol: HTTPS
  name: registry-address
  namespace: kube-system
  labels:
    service_name: registry
spec:
  rules:
    - host: $REGISTRY_DOMAIN_NAME
      http:
        paths:
          - backend:
              service:
                name: registry
                port:
                  number: 443
              path: /v2/
              pathType: ImplementationSpecific
          - backend:
              service:
                name: registry
                port:
                  number: 443
              path: /v2/_catalog
              pathType: ImplementationSpecific
          - backend:
              service:
                name: registry
                port:
                  number: 443
              path: /v2/./tags/list
              pathType: ImplementationSpecific
          - backend:
              service:
                name: registry
                port:
                  number: 443
              path: /v2/./manifests/[A-Za-z0-9_+.-:]+

```

```

    pathType: ImplementationSpecific
  - backend:
      service:
        name: registry
        port:
          number: 443
      path: /v2/./blobs/[A-Za-z0-9-:]+
      pathType: ImplementationSpecific
  - backend:
      service:
        name: registry
        port:
          number: 443
      path: /v2/./blobs/uploads/[A-Za-z0-9-:]+
      pathType: ImplementationSpecific
  - backend:
      service:
        name: registry
        port:
          number: 443
      path: /auth/token
      pathType: ImplementationSpecific
tls:
  - secretName: registry-address.tls
  hosts:
    - $REGISTRY_DOMAIN_NAME
EOF

```

A response similar to `... created` indicates successful ingress creation.

#### 4. Check if a Registry Service resource exists:

```
kubectl -n kube-system get svc | grep registry
```

If the Service doesn't exist, create it with:

```
cat << EOF | kubectl create -f -
apiVersion: v1
kind: Service
metadata:
  labels:
    name: registry
    service_name: registry
  name: registry
  namespace: kube-system
spec:
  ports:
    - protocol: TCP
      port: 443
      targetPort: 60080
  selector:
    component: registry
  type: ClusterIP
EOF
```

5. Test the configuration by pulling an image from the registry using the domain name:

```
cricctl pull <registry-domain-name>/automation/qaimages:helloworld
```

Or

```
docker pull <registry-domain-name>/automation/qaimages:helloworld
```

# Choosing a Container Runtime

## TOC

[Overview](#)[Quick Selection Guide](#)[Differences Between Docker and Containerd](#)[Common Commands](#)[Call Chain Differences](#)[Log and Parameter Comparison](#)[CNI Network Comparison](#)

## Overview

Container Runtime is a core component of Kubernetes, responsible for managing the lifecycle of images and containers.

When creating clusters through the platform, you can choose either Containerd or Docker as your runtime component.

**Note:** Kubernetes version 1.24 and above no longer officially supports Docker runtime. The officially recommended runtime is Containerd. If you still need to use Docker runtime, you must first enable `cri-docker` in the feature gate before you can select Docker as the runtime component when creating a cluster. For details on using feature gates, see [Feature Gate Configuration](#).



# Quick Selection Guide

Choose Containerd	Choose Docker
<ul style="list-style-type: none"><li>• Shorter call chain</li><li>• Fewer components</li><li>• More stable</li><li>• Consumes fewer node resources</li></ul>	<ul style="list-style-type: none"><li>• Supports docker-in-docker</li><li>• Allows use of <code>docker build/push/save/load</code> commands on nodes</li><li>• Can call Docker API</li><li>• Supports docker compose or docker swarm</li></ul>

## Differences Between Docker and Containerd

### Common Commands

Containerd	Docker	Description
<b>crictl ps</b>	docker ps	View running containers
<b>crictl inspect</b>	docker inspect	View container details
<b>crictl logs</b>	docker logs	View container logs
<b>crictl exec</b>	docker exec	Execute commands in container
<b>crictl attach</b>	docker attach	Attach to container
<b>crictl stats</b>	docker stats	Display container resource usage
<b>crictl create</b>	docker create	Create container
<b>crictl start</b>	docker start	Start container
<b>crictl stop</b>	docker stop	Stop container
<b>crictl rm</b>	docker rm	Remove container

Containerd	Docker	Description
<b>crictl images</b>	docker images	View image list
<b>crictl pull</b>	docker pull	Pull image
<b>None</b>	docker push	Push image
<b>crictl rmi</b>	docker rmi	Delete image
<b>crictl pods</b>	None	View pod list
<b>crictl inspectp</b>	None	View pod details
<b>crictl runp</b>	None	Start pod
<b>crictl stopp</b>	docker images	View images
<b>ctr images ls</b>	None	Stop pod
<b>crictl stopp</b>	docker load/save	Import/export images
<b>ctr images import/export</b>	None	Stop pod
<b>ctr images pull/push</b>	docker pull/push	Pull/push images
<b>ctr images tag</b>	docker tag	Tag images

## Call Chain Differences

- Docker as Kubernetes container runtime has the following call relationship:

kubelet > cri-dockerd > dockerd > containerd > runC

- Containerd as Kubernetes container runtime has the following call relationship:

kubelet > cri plugin (in containerd process) > containerd > runC

**Summary:** Although dockerd adds features like swarm cluster, docker build, and Docker API, it can introduce bugs and adds an extra layer in the call chain. Containerd has a shorter call chain, fewer components, greater stability, and consumes fewer node resources.

## Log and Parameter Comparison

Comparison	Docker	Containerd
<b>Storage Path</b>	<p>When Docker serves as the Kubernetes container runtime, container logs are stored by Docker in directories like <code>/var/lib/docker/containers/\$CONTAINERID</code>.</p> <p>Kubelet creates symbolic links in <code>/var/log/pods</code> and <code>/var/log/containers</code> pointing to the container log files in this directory.</p>	<p>When Containerd serves as the Kubernetes container runtime, container logs are stored by Kubelet in the <code>/var/log/pods/\$CONTAINER_NAME</code> directory, with symbolic link created in the <code>/var/log/containers</code> directory pointing to the log files.</p>
<b>Configuration Parameters</b>	<p>Specified in the Docker configuration file:</p> <pre>"log-driver": "json-file", "log-opts": {"max-size": "100m", "max-file": "5"}</pre>	<p><i>Method 1:</i> Specified in kubelet parameters:</p> <pre>--container-log-max-files=5 --container-log-max-size="100Mi"</pre> <p><i>Method 2:</i> Specified in KubeletConfiguration:</p> <pre>"containerLogMaxSize": "100Mi", "containerLogMaxFiles": 5,</pre>
<b>Saving Container Logs to Data Disk</b>	<p>Mount the data disk to "data-root" (default is <code>/var/lib/docker</code>).</p>	<p>Create a symbolic link <code>/var/log/pods</code> pointing to a directory under the data disk mount point.</p>

## CNI Network Comparison

Comparison	Docker	Containerd
<b>Who Calls CNI</b>	cri-dockerd	cri-plugin built into Containerd (after containerd 1.1)

Comparison	Docker	Containerd
How to Configure CNI	cri-dockerd parameters <code>--cni-conf-dir</code> <code>--cni-bin-dir</code> <code>--cni-cache-dir</code>	Containerd configuration file (toml): <code>[plugins.cri.cni]</code> <code>bin_dir = "/opt/cni/bin"</code> <code>conf_dir = "/etc/cni/net.d"</code>

# Updating Public Repository Credentials

## TOC

[Overview](#)[Procedure](#)

## Overview

The `Public Repository` is a platform-provided image registry service available on the public internet. When you want your clusters to use the `Public Repository` as their image registry, you need to update the built-in `public-registry-credential` Cloud Credentials. This ensures your platform has permission to pull images from the public registry.

## Procedure

1. Log in to the **Customer Portal** and download your organization's authentication file from the **Enterprise Management** section located in the **User Information** dropdown in the upper right corner.
2. Navigate to **Clusters > Cloud Credential** in the left navigation bar of the **Administrator** console.
3. Locate the cloud credential named `public-registry-credential` and click **Update** from the dropdown menu on the right.

4. In the **Upload Public Repository Address** section, upload the authentication file you downloaded from the **Customer Portal**.
5. Click **Update** to apply the changes.