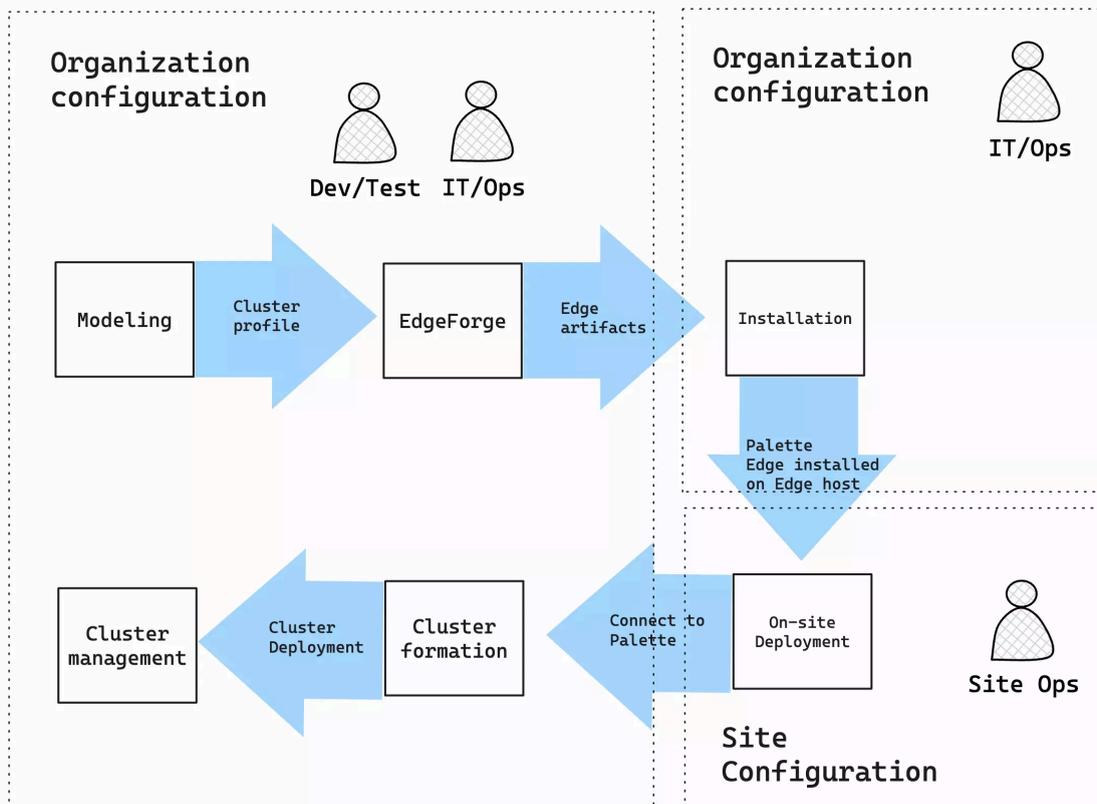


Edge Deployment Lifecycle

The typical end-to-end lifecycle of deploying clusters at edge locations involves several distinct phases in which different organizations or teams need to perform specific tasks.

Central Clusters

A central cluster is a cluster that is managed by a Palette instance. The Palette instance could be the public Palette SaaS or a self-hosted instance. The following diagram represents the deployment lifecycle of a central cluster.



Feedback

1. **Modeling.** App owners build and test the applications in test environments and model application profiles in Palette for a cluster installation. For more information, refer to [Model Cluster Profile](#).
2. **EdgeForge.** IT/Ops teams build the Palette Edge Installer and the provider images, which contain the OS and Kubernetes layer of their cluster. The Edge Installer build process expects a configuration file called **user-data**, which specifies properties such as Palette

endpoints, registration tokens, default OS users, and default network settings. For more information, refer to [EdgeForge Workflow](#).

3. **Installation.** IT/Ops use the installer prepared in the EdgeForge phase and installs Palette Edge onto the Edge hosts. The Edge hosts are then shipped to edge sites. For more information, refer to [Installation](#).
4. **On-site deployment.** Edge hosts need to be registered with Palette. In this phase, the field technician powers on the Edge host, which will boot to registration. The field technician also applies site-specific properties such as static IP address, network proxy, and SSL certificates with site user data or Palette Terminal User Interface (TUI). For more information, refer to [Deploy Edge Hosts On-Site](#).
5. **Cluster formation.** Once an Edge host is registered with Palette, the Edge host can be used to form a new cluster or be added to an existing cluster. Each cluster uses a cluster profile modeled in the **Modeling** stage. For more information, refer to [Cluster Definition](#).
6. **Cluster management.** In this phase, the IT/Ops team performs maintenance and update activities on your active clusters. Using versioned cluster profiles, Palette allows you to centrally managed all your Edge clusters and perform rolling upgrades with zero downtime. For more information about cluster upgrades, refer to [Update a Cluster](#).

If the edge location configuration is known and predictable, then the IT/Ops team can combine EdgeForge, installation, and registration into one step and ship the fully configured Edge hosts to the edge location. The field technician at the edge location only needs to hook up the power and network cables without further device configuration. The Edge host will be ready to be centrally managed for cluster formation.

Local Clusters <#>

A local cluster is a cluster that does not have a connection to a Palette instance. The deployment lifecycle for local Edge hosts is slightly different from central Edge hosts. Since there is no Palette instance and each Edge host is managed locally, there is no Edge host registration during the on-site deployment phase.

In addition, if your Edge host does not have internet access, you need to provide the Edge host with the images that are required for cluster deployment through a content bundle or through an external registry. After a cluster is formed, you manage the cluster locally with Local UI instead of using through a Palette instance.

The following are the phases of the deployment lifecycle of local Edge clusters:

1. **Modeling.** App owners build and test the applications in test environments and model application profiles in Palette for a cluster installation.
2. **EdgeForge.** IT/Ops teams build the Palette Edge Installer and the provider images, which contain the OS and Kubernetes layer of their cluster. Through the installer user data configuration, IT/Ops teams can configure the Edge host to download images from a private external registry. For more information, refer to [Deploy Cluster with Private External Registry](#).

The IT/Ops teams can also create a content bundle containing all the necessary assets needed to provision their cluster and a cluster definition that reflects the cluster profile designed in the **Modeling**

stage. Both the content bundle and the cluster definition can be included in the installer. The installer is then exported to all the site locations. For more information, refer to [Build Content Bundle](#) and [Export Cluster Definition](#).

3. **Installation.** IT/Ops use the installer prepared in the EdgeForge phase and installs Palette Edge onto the Edge hosts. The Edge hosts are then shipped to edge sites.
4. **On-site deployment.** Local Edge hosts do not have a connection to Palette and therefore do not require registration. Site operator powers on the Edge host, which will boot up and serve Local UI and the Edge Management API instead of to the registration screen. The field technician also applies site-specific properties such as static IP address, network proxy, and SSL certificates with site user data or Palette Terminal User Interface (TUI). For more information, refer to [Access Local UI](#).
5. **Cluster formation.** If the Edge host is installed with a content bundle and cluster definition, the field technician can power on the Edge host and create a cluster using Local UI. For more information, refer to [Create a Cluster with Local UI](#).

If the Edge host is missing either the content bundle or the cluster definition, you can still build a content bundle or export the cluster definition and upload them to the Edge host through Local UI. For more information, refer to [Upload Content Bundle](#). If the Edge host is configured to pull images from a private external registry, then cluster formation does not require a content bundle, but it will still require a cluster definition.

6. **Cluster management.** In this phase, the IT/Ops teams perform maintenance and update activities on your active clusters. Using a in-cluster primary registry, you can provide images that can be used to upgrade your cluster. For more information, refer to [Deploy Cluster with Primary Registry](#).

Next Steps <#>

Now that you have an understanding of the deployment lifecycle, start the deployment of your Edge host by modeling a cluster profile for your Edge

cluster to use and start building the required Edge artifacts for deployment.

- [Model Cluster Profile](#)
- [EdgeForge](#)

Tags: [edge](#)

*Last updated on **May 31, 2025***