



Infrastructure management through API

BigDataStack Software Component developed by Red Hat

Description

Unified API for infrastructure resources to make infrastructure management easy and abstracted from the real infrastructure.

Features

Extension of Kubernetes Cluster API project to have an OpenStack abstraction together with its own operator/actuator, named Cluster API Provider OpenStack. This allows automating creation/scaling actions regarding OpenShift/Kubernetes nodes when running on top of OpenStack. Therefore, helping to manage the infrastructure (OpenShift/Kubernetes nodes) through the abstracted API.

Areas of Application

Deploying OpenShift on top of OpenStack. Allows to easily adapt the size of the OpenShift cluster to the current needs. Also, it allows having upper layers managing that, like the OpenShift autoscaler (based on CPU and Memory) or the Big Data Stack components with more advance mechanisms.



Infrastructure management through API

Market trends & opportunities

Managing Infrastructure (as well as applications) in a declarative manner is the current trend, and this is enabling this direction for Infrastructure management on top of OpenStack.

Customer benefits

Easy to scale up/down the Kubernetes/OpenShift clusters through the Kubernetes API, as well as improved fault tolerance and resilience support. Managing Infrastructure through APIs.

Technological novelty

N/A.

TRL level: 9

Find the Open Source code here:

This involves different code in different repos:

- <https://github.com/openshift/cluster-network-operator>
- <https://github.com/openshift/installer>
- <https://github.com/kubernetes-sigs/cluster-api-provider-openstack>
- <https://github.com/openshift/machine-config-operator>
- <https://github.com/openshift/machine-api-operator>