

# EMCO Overview

EMCO is a Geo-distributed application orchestrator that intelligently places a tenant's workload onto one or more clusters. These clusters can be a public cloud cluster, an enterprise's private cloud cluster, an IoT edge site cluster, a 5G cell cluster, or a telco's CO cluster. The workload can be a simple app or a complicated app that is composed of multiple simple apps, or it can be a network function in the form of container/VM, etc.

EMCO operates at a higher level than kubernetes. It makes a decision on which clusters a workload should run, then it interacts with the Kubernetes API server and hands over that workload to the Kubernetes control plane. EMCO Supports multiple placement constraints: Affinity and Anti-Affinity; Platform capabilities; Latency; Cost. Through auto-configuring service mesh (ISTIO) and security policy (NAT, firewall), it enables the communication between deployed apps or between a deployed app and an external service within or across clusters with/without mutual TLS. It provides application life cycle management including upgrades as well as comprehensive status monitoring of deployed applications.

EMCO is an open source project in Linux foundation and there is no vendor lock. In EMCO cluster onboarding is done through cluster registration. You can specify the deployment intent of your workload and use EMCO to deploy your workload to a specific provider's cluster.

EMCO can run independently as a universal multi-cloud orchestrator which manages workload deployment across geo-distributed cloud clusters and edge clusters. It can also integrate with CSP's multi-cloud orchestrator to deploy a complicated application to multiple CSP's cloud clusters.

See: [\[EMCO Architecture and Design\]](#)