
How to Deploy a Proxy in Kubernetes

<https://campus.barracuda.com/doc/93201597/>

Before You Begin

Prerequisites:

- Running [kubernetes cluster](#) or local [minikube](#) installation
- Configured and installed [kubect](#)
- [CloudGen Access Proxy enrollment link](#)

The steps described assume familiarity with kubernetes. The required images are available in the Dockerhub registry under the organization [FydeInc](#).

Helm Chart

The helm chart is available at [Artifactory](#) and includes:

- Orchestrator
- Envoy
- Redis (from [dandy-developer chart](#))
- HTTP resource (to test access)

Check the Artifactory link or `values.yaml` file for all the configuration parameters.

Deployment

Create a `custom-values.yaml` file with the desired values.

Minimum required configuration for a successful deployment:

- 'orchestrator.enrollmentToken.existingSecret.*' or 'orchestrator.enrollmentToken.newSecret'

Minimum required configuration to allow external access:

- 'envoy.loadBalancer.enabled'
- 'envoy.loadBalancer.annotations'

Add the helm repo and install the chart:

```
<code sh>
helm repo add barracuda-cloudgen-access
https://barracuda-cloudgen-access.github.io/helm-charts
helm install my-release barracuda-cloudgen-access/cga-proxy --namespace my-
namespace --values custom-values.yaml
</code sh>
```

Optionally, instead of using "helm install", export the yaml files and deploy using your preferred method:

```
<code sh>
helm template barracuda-cloudgen-access/cga-proxy --values custom-values.yaml
</code sh>
```

Example

This extended example for a high availability deployment uses AWS NLB and Prometheus service monitors:

```
<code yaml>
nameOverride: cga-proxy-my-deploy
serviceMonitor: true
priorityClassName: high-priority

orchestrator:
  enrollmentToken:
```

```
existingSecret:
  name: cga-orchestrator-secret
  key: enrollment-token
highAvailability:
  enabled: true
resources:
  limits:
    cpu: 500m
    memory: 512Mi
  requests:
    cpu: 100m
    memory: 128Mi

envoy:
  replicaCount: 3
  loadBalancer:
    enabled: true
    annotations:
      service.beta.kubernetes.io/aws-load-balancer-additional-resource-tags:
        app=cga-proxy,service=envoy-proxy-external,namespace=cga-proxy-my-
deploy
  service.beta.kubernetes.io/aws-load-balancer-type: nlb
  externalTrafficPolicy: Local
  port: 443

resources:
  limits:
    cpu: 500m
    memory: 512Mi
  requests:
    cpu: 100m
    memory: 128Mi

redis-ha:
  nameOverride: redis
  priorityClassName: high-priority
  redis:
    resources:
      requests:
        cpu: 100m
        memory: 200Mi
      limits:
        cpu: 100m
        memory: 200Mi
  sentinel:
    resources:
```

```
    requests:
      cpu: 100m
      memory: 200Mi
    limits:
      cpu: 100m
      memory: 200Mi
  hardAntiAffinity: true
  exporter:
    enabled: true
    resources:
      requests:
        cpu: 100m
        memory: 200Mi
      limits:
        cpu: 100m
        memory: 200Mi
  serviceMonitor:
    enabled: true
  podDisruptionBudget:
    minAvailable: 2
  persistentVolume:
    enabled: false
  http-test:
    enabled: true
</code yml>
```

© Barracuda Networks Inc., 2024 The information contained within this document is confidential and proprietary to Barracuda Networks Inc. No portion of this document may be copied, distributed, publicized or used for other than internal documentary purposes without the written consent of an official representative of Barracuda Networks Inc. All specifications are subject to change without notice. Barracuda Networks Inc. assumes no responsibility for any inaccuracies in this document. Barracuda Networks Inc. reserves the right to change, modify, transfer, or otherwise revise this publication without notice.