

LOGZILLA DOCUMENTATION

Kubernetes API Module

LogZilla API module Kubernetes manifest defining PersistentVolumeClaims and StatefulSet for Gunicorn, Tornado, and Celery worker processes

Kubernetes Deployment Overview · Generated May 3, 2026 · logzilla.ai/docs/kubernetes-deployment/api-module

API module manifest

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: config-pvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 100Mi
---
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: config-logs-pvc
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 100Mi
---
apiVersion: apps/v1
kind: StatefulSet
metadata:
  name: api
spec:
  serviceName: api
  selector:
    matchLabels:
      name: api
  template:
    metadata:
      labels:
        name: api
    spec:
      enableServiceLinks: false
      containers:
        - name: gunicorn
          image: logzilla/runtime:stable
          imagePullPolicy: Always
          resources:
            requests:
              memory: "500Mi"
              cpu: "500m"
            limits:
              memory: "2Gi"
              cpu: "4000m"
```

```
command:
  - /bin/sh
  - -c
  - >
    /usr/lib/logzilla/bin/lz-simple-setup &&
    gunicorn -c /usr/lib/logzilla/lib/logzilla/api/gunicorn.conf.py
    --log-level=info --workers=10 logzilla.api.wsgi:application
ports:
- containerPort: 80
  name: api
  protocol: TCP
startupProbe:
  httpGet:
    path: /ping
    port: api
    scheme: HTTP
  initialDelaySeconds: 20
  periodSeconds: 5
  successThreshold: 1
  failureThreshold: 5
  timeoutSeconds: 30
readinessProbe:
  httpGet:
    path: /ping
    port: api
    scheme: HTTP
  initialDelaySeconds: 5
  periodSeconds: 10
  successThreshold: 1
  failureThreshold: 3
  timeoutSeconds: 30
livenessProbe:
  httpGet:
    path: /ping
    port: api
    scheme: HTTP
  periodSeconds: 10
  successThreshold: 1
  failureThreshold: 3
  timeoutSeconds: 30
envFrom:
  - configMapRef:
      name: postgres-multimap
  - configMapRef:
      name: log-multimap
  - configMapRef:
      name: django-multimap
  - secretRef:
      name: django-secret
  - secretRef:
      name: postgres-secret
  - secretRef:
      name: internal-api-secret
volumeMounts:
```

```
- name: config
  mountPath: /etc/logzilla
  readOnly: false
- name: logs
  mountPath: /var/log/logzilla
  readOnly: false
- name: tornado
  image: logzilla/runtime:stable
  imagePullPolicy: Always
  resources:
    requests:
      memory: "500Mi"
      cpu: "200m"
    limits:
      memory: "1Gi"
      cpu: "1000m"
  command: ["/usr/lib/logzilla/bin/tornado_server"]
  ports:
- containerPort: 8001
  name: ws
  protocol: TCP
  livenessProbe:
    httpGet:
      path: /ping
      port: ws
      scheme: HTTP
    periodSeconds: 10
    successThreshold: 1
    failureThreshold: 3
    timeoutSeconds: 5
  startupProbe:
    httpGet:
      path: /ping
      port: ws
      scheme: HTTP
    initialDelaySeconds: 10
    periodSeconds: 5
    successThreshold: 1
    failureThreshold: 3
    timeoutSeconds: 1
  readinessProbe:
    httpGet:
      path: /ping
      port: ws
      scheme: HTTP
    initialDelaySeconds: 5
    periodSeconds: 10
    successThreshold: 1
    failureThreshold: 3
    timeoutSeconds: 1
  envFrom:
- configMapRef:
  name: postgres-multimap
- configMapRef:
```

```
    name: log-multimap
  - configMapRef:
    name: django-multimap
  - secretRef:
    name: django-secret
  - secretRef:
    name: postgres-secret
  - secretRef:
    name: internal-api-secret
- name: celeryworker
image: logzilla/runtime:stable
imagePullPolicy: Always
resources:
  requests:
    memory: "1Gi"
    cpu: "500m"
  limits:
    memory: "4Gi"
    cpu: "2000m"
command:
  - celery
  - -A
  - logzilla.api.celery_app
  - worker
  - --autoscale
  - 4,2
livenessProbe:
  exec:
    command: ["/usr/lib/logzilla/bin/healthcheck", "celery"]
  periodSeconds: 10
  successThreshold: 1
  failureThreshold: 5
  timeoutSeconds: 5
startupProbe:
  exec:
    command: ["/usr/lib/logzilla/bin/healthcheck", "celery"]
  initialDelaySeconds: 20
  periodSeconds: 10
  successThreshold: 1
  failureThreshold: 3
  timeoutSeconds: 5
envFrom:
  - configMapRef:
    name: postgres-multimap
  - configMapRef:
    name: log-multimap
  - configMapRef:
    name: django-multimap
  - secretRef:
    name: django-secret
  - secretRef:
    name: postgres-secret
  - secretRef:
    name: internal-api-secret
```

```
    volumeMounts:
      - name: config
        mountPath: /etc/logzilla
        readOnly: false
      - name: logs
        mountPath: /var/log/logzilla
        readOnly: false
    volumes:
      - name: config
        persistentVolumeClaim:
          claimName: config-pvc
      - name: logs
        persistentVolumeClaim:
          claimName: config-logs-pvc
---
apiVersion: v1
kind: Service
metadata:
  name: api
  annotations:
    cloud.google.com/neg: '{"ingress": true}'
spec:
  ports:
    - port: 80
      targetPort: 80
      protocol: TCP
      name: api
    - port: 8001
      targetPort: 8001
      protocol: TCP
      name: ws
  clusterIP: None
  selector:
    name: api
```

Notes

- The `api` Service exposes HTTP (80) and WebSocket (8001). On GKE, NEG annotations are used for Ingress backends.
- Ensure secrets and config maps referenced here are created first (see [Common Config and Secrets](#) page).