

# RabbitMQ Cluster

- [Install RabbitMQ Cluster in Kubernetes](#)

# Install RabbitMQ Cluster in Kubernetes

## Volumes

[cluster-operator/rabbitmq.yaml at main · rabbitmq/cluster-operator · GitHub](#)

we use local path storage provisioner

download latest rabbitmq operator:

```
https://github.com/rabbitmq/cluster-operator/releases/latest/download/cluster-operator.yml  
sed -i 's/rabbitmqoperator\\/cluster-operator: 2.2.0/{REGISTRY_IP}:{PORT}\\rabitmqoperator\\/cl
```

apply this file with editing image:

```
apiVersion: rabbitmq.com/v1beta1  
kind: RabbitmqCluster  
metadata:  
  name: rabbitmq-production-ready  
  namespace: infra  
spec:  
  replicas: 3  
  image: {REGISTRY_IP}:{PORT}/rabbitmq: 3.11.10-management  
  persistence:  
    storageClassName: local-path  
    storage: 10Gi  
  resources:  
    requests:  
      cpu: 500m  
      memory: 1Gi  
    limits:  
      cpu: 20000m  
      memory: 4Gi  
  affinity:  
    nodeAffinity:  
      requiredDuringSchedulingIgnoredDuringExecution:  
        nodeSelectorTerms:  
          - matchExpressions:  
            - key: infra  
              operator: In  
              values:
```

```
- "true"
rabbitmq:
  additionalConfig: |
    cluster_partition_handling = pause_minority
    vm_memory_high_watermark_paging_ratio = 0.99
    disk_free_limit.relative = 1.0
    collect_statistics_interval = 10000
    default_user_tags.administrator = true
    default_user = mavara
    default_pass = vw8Lh-*KZKATNbDb
  service:
    type: LoadBalancer
```

```
apiVersion: policy/v1beta1
kind: PodDisruptionBudget
metadata:
  name: production-ready-rabbitmq
  namespace: infra
spec:
  maxUnavailable: 1
  selector:
    matchLabels:
      app.kubernetes.io/name: rabbitmq-production-ready
```

end