

Installing PatchMon Server on K8S with Helm

PatchMon Helm Chart Documentation

Helm chart for deploying PatchMon on Kubernetes.

- Chart repository: github.com/RuTHlessBEat200/PatchMon-helm
- Application repository: github.com/PatchMon/PatchMon

Overview

PatchMon v2.0.0 runs as a containerised application made up of four services:

- **Database** -- PostgreSQL 18 (StatefulSet with persistent storage)
- **Redis** -- Redis 8 (used for BullMQ job queues and caching, StatefulSet with persistent storage)
- **Server** -- Single Go binary serving both the backend API and the React frontend static files (StatefulSet with optional HPA)
- **Guacd** -- Apache Guacamole proxy daemon for remote desktop / SSH terminal access (Deployment)

“ **Note:** In v2.0.0 the separate backend and frontend containers have been merged into a single `server` binary. If you are upgrading from v1.x, update your values files accordingly — `backend.*` and `frontend.*` keys no longer exist.

The chart deploys all four components into a single namespace and wires them together automatically using init containers, internal ClusterIP services, and a shared ConfigMap.

Container Images

Component	Image	Default Tag
Server	ghcr.io/patchmon/patchmon-server	2.0.0
Database	docker.io/postgres	18-alpine
Redis	docker.io/redis	8-alpine
Guacd	docker.io/guacamole/guacd	latest

Available Tags (Server)

Tag	Description
latest	Latest stable release
x.y.z	Exact version pin (e.g. 2.0.0)
x.y	Latest patch in a minor series (e.g. 2.0)
x	Latest minor and patch in a major series (e.g. 2)
edge	Latest development build from the main branch -- may be unstable, for testing only

Prerequisites

- Kubernetes 1.19+
- Helm 3.0+
- A PersistentVolume provisioner in the cluster (for database and Redis storage)
- (Optional) An Ingress controller (e.g. NGINX Ingress) for external access
- (Optional) cert-manager for automatic TLS certificate management
- (Optional) Metrics Server for HPA functionality

Quick Start

The quickest way to get PatchMon running is to use the provided [values-quick-start.yaml](#) file. It contains all required secrets inline and sensible defaults so you can install with a single command.

I

Warning: `values-quick-start.yaml` ships with placeholder secrets and is intended for evaluation and testing only. Never use it in production without replacing all secret values.

1. Install the chart

```
wget https://github.com/RuTHlessBEat200/PatchMon-helm/blob/main/values-quick-start.yaml
helm install patchmon oci://ghcr.io/ruthlessbeat200/charts/patchmon \
  --namespace patchmon \
  --create-namespace \
  --values values-quick-start.yaml
```

2. Wait for pods to become ready

```
kubectl get pods -n patchmon -w
```

3. Access PatchMon

If ingress is enabled, open the host you configured (e.g. `https://patchmon-dev.example.com`).

Without ingress, use port-forwarding:

```
kubectl port-forward -n patchmon svc/patchmon-dev-server 3000:3000
```

Then navigate to `http://localhost:3000` and complete the first-time setup to create your admin account.

Production Deployment

For production use, refer to the provided `values-prod.yaml` file as a starting point. It demonstrates how to:

- Use an external secret (e.g. managed by KSOPS, Sealed Secrets, or External Secrets Operator) instead of inline passwords
- Configure HTTPS with cert-manager
- Set the correct server protocol, host, and port for agent communication

1. Create your secrets

The chart does **not** auto-generate secrets. You must supply them yourself.

Required secrets:

Key	Description
<code>postgres-password</code>	PostgreSQL password
<code>redis-password</code>	Redis password
<code>jwt-secret</code>	JWT signing secret for the server
<code>ai-encryption-key</code>	Encryption key for AI provider credentials
<code>oidc-client-secret</code>	OIDC client secret (only if OIDC is enabled)

You can either:

- Set passwords directly in your values file (`database.auth.password`, `redis.auth.password`, `server.jwtSecret`, `server.aiEncryptionKey`), or
- Create a Kubernetes Secret separately and reference it with `existingSecret` / `existingSecretPasswordKey` fields.

Example -- creating a secret manually:

```
kubectl create namespace patchmon

kubectl create secret generic patchmon-secrets \
  --namespace patchmon \
  --from-literal=postgres-password="$(openssl rand -hex 32)" \
  --from-literal=redis-password="$(openssl rand -hex 32)" \
  --from-literal=jwt-secret="$(openssl rand -hex 64)" \
  --from-literal=ai-encryption-key="$(openssl rand -hex 32)"
```

Secret management tools for production:

- [KSOPS](#) -- encrypt secrets in Git using Mozilla SOPS
- [Sealed Secrets](#) -- encrypt secrets that only the cluster can decrypt
- [External Secrets Operator](#) -- sync secrets from external stores (Vault, AWS Secrets Manager, etc.)
- [Vault](#) -- enterprise-grade secret management

2. Create your values file

Start from `values-prod.yaml` and adjust to your environment:

```
global:
  storageClass: "your-storage-class"

fullnameOverride: "patchmon-prod"

server:
  env:
    serverProtocol: https
    serverHost: patchmon.example.com
    serverPort: "443"
    corsOrigin: https://patchmon.example.com
  existingSecret: "patchmon-secrets"
  existingSecretJwtKey: "jwt-secret"
  existingSecretAiEncryptionKey: "ai-encryption-key"

database:
  auth:
    existingSecret: patchmon-secrets
    existingSecretPasswordKey: postgres-password

redis:
  auth:
    existingSecret: patchmon-secrets
    existingSecretPasswordKey: redis-password

secret:
  create: false # Disable chart-managed secret since we use an external one

ingress:
  enabled: true
  className: nginx
  annotations:
    cert-manager.io/cluster-issuer: letsencrypt-prod
    nginx.ingress.kubernetes.io/proxy-body-size: "0"
    nginx.ingress.kubernetes.io/proxy-read-timeout: "3600"
    nginx.ingress.kubernetes.io/proxy-send-timeout: "3600"
    nginx.ingress.kubernetes.io/proxy-connect-timeout: "60"
    nginx.ingress.kubernetes.io/proxy-http-version: "1.1"
```

```

nginx.ingress.kubernetes.io/client-body-buffer-size: "4m"
nginx.ingress.kubernetes.io/websocket-services: "server"
hosts:
  - host: patchmon.example.com
    paths:
      - path: /
        pathType: Prefix
        service:
          name: server
          port: 3000
tls:
  - secretName: patchmon-tls
    hosts:
      - patchmon.example.com

```

3. Install

```

helm install patchmon oci://ghcr.io/ruthlessbeat200/charts/patchmon \
  --namespace patchmon \
  --create-namespace \
  --values values-prod.yaml

```

Configuration Reference

Global Settings

Parameter	Description	Default
<code>global.imageRegistry</code>	Override the image registry for all components	""
<code>global.imageTag</code>	Override the image tag for the server (takes priority over <code>server.image.tag</code> if set)	""
<code>global.imagePullSecrets</code>	Image pull secrets applied to all pods	[]
<code>global.storageClass</code>	Default storage class for all PVCs	""
<code>nameOverride</code>	Override the chart name used in resource names	""

Parameter	Description	Default
<code>fullnameOverride</code>	Override the full resource name prefix	<code>"patchmon-prod"</code>
<code>commonLabels</code>	Labels added to all resources	<code>{}</code>
<code>commonAnnotations</code>	Annotations added to all resources	<code>{}</code>

Database (PostgreSQL)

Parameter	Description	Default
<code>database.enabled</code>	Deploy the PostgreSQL StatefulSet	<code>true</code>
<code>database.image.registry</code>	Image registry	<code>docker.io</code>
<code>database.image.repository</code>	Image repository	<code>postgres</code>
<code>database.image.tag</code>	Image tag	<code>18-alpine</code>
<code>database.image.pullPolicy</code>	Image pull policy	<code>IfNotPresent</code>
<code>database.host</code>	External database host (overrides built-in service discovery when set)	<code>" "</code>
<code>database.port</code>	External database port	<code>" "</code>
<code>database.auth.database</code>	Database name	<code>patchmon_db</code>
<code>database.auth.username</code>	Database user	<code>patchmon_user</code>
<code>database.auth.password</code>	Database password (required if <code>existingSecret</code> is not set)	<code>" "</code>
<code>database.auth.existingSecret</code>	Name of an existing secret containing the password	<code>" "</code>
<code>database.auth.existingSecretPasswordKey</code>	Key inside the existing secret	<code>postgres-password</code>
<code>database.replicaCount</code>	Number of replicas	<code>1</code>
<code>database.updateStrategy.type</code>	StatefulSet update strategy	<code>RollingUpdate</code>
<code>database.persistence.enabled</code>	Enable persistent storage	<code>true</code>
<code>database.persistence.storageClass</code>	Storage class (falls back to <code>global.storageClass</code>)	<code>" "</code>
<code>database.persistence.accessModes</code>	PVC access modes	<code>["ReadWriteOnce"]</code>
<code>database.persistence.size</code>	PVC size	<code>5Gi</code>
<code>database.resources.requests.cpu</code>	CPU request	<code>100m</code>
<code>database.resources.requests.memory</code>	Memory request	<code>128Mi</code>
<code>database.resources.limits.cpu</code>	CPU limit	<code>1000m</code>
<code>database.resources.limits.memory</code>	Memory limit	<code>1Gi</code>

Parameter	Description	Default
<code>database.livenessProbe.enabled</code>	Enable liveness probe	<code>true</code>
<code>database.readinessProbe.enabled</code>	Enable readiness probe	<code>true</code>
<code>database.service.type</code>	Service type	<code>ClusterIP</code>
<code>database.service.port</code>	Service port	<code>5432</code>
<code>database.podAnnotations</code>	Pod annotations	<code>{}</code>
<code>database.podSecurityContext</code>	Pod-level security context	see <code>values.yaml</code>
<code>database.securityContext</code>	Container-level security context	see <code>values.yaml</code>
<code>database.nodeSelector</code>	Node selector	<code>{}</code>
<code>database.tolerations</code>	Tolerations	<code>[]</code>
<code>database.affinity</code>	Affinity rules	<code>{}</code>

Redis

Parameter	Description	Default
<code>redis.enabled</code>	Deploy the Redis StatefulSet	<code>true</code>
<code>redis.image.registry</code>	Image registry	<code>docker.io</code>
<code>redis.image.repository</code>	Image repository	<code>redis</code>
<code>redis.image.tag</code>	Image tag	<code>8-alpine</code>
<code>redis.image.pullPolicy</code>	Image pull policy	<code>IfNotPresent</code>
<code>redis.auth.password</code>	Redis password (required if <code>existingSecret</code> is not set)	<code>""</code>
<code>redis.auth.existingSecret</code>	Name of an existing secret containing the password	<code>""</code>
<code>redis.auth.existingSecretPasswordKey</code>	Key inside the existing secret	<code>redis-password</code>
<code>redis.replicaCount</code>	Number of replicas	<code>1</code>
<code>redis.updateStrategy.type</code>	StatefulSet update strategy	<code>RollingUpdate</code>
<code>redis.persistence.enabled</code>	Enable persistent storage	<code>true</code>
<code>redis.persistence.storageClass</code>	Storage class (falls back to <code>global.storageClass</code>)	<code>""</code>
<code>redis.persistence.accessModes</code>	PVC access modes	<code>["ReadWriteOnce"]</code>
<code>redis.persistence.size</code>	PVC size	<code>5Gi</code>
<code>redis.resources.requests.cpu</code>	CPU request	<code>50m</code>
<code>redis.resources.requests.memory</code>	Memory request	<code>10Mi</code>

Parameter	Description	Default
<code>redis.resources.limits.cpu</code>	CPU limit	<code>500m</code>
<code>redis.resources.limits.memory</code>	Memory limit	<code>512Mi</code>
<code>redis.livenessProbe.enabled</code>	Enable liveness probe	<code>true</code>
<code>redis.readinessProbe.enabled</code>	Enable readiness probe	<code>true</code>
<code>redis.service.type</code>	Service type	<code>ClusterIP</code>
<code>redis.service.port</code>	Service port	<code>6379</code>
<code>redis.podAnnotations</code>	Pod annotations	<code>{}</code>
<code>redis.podSecurityContext</code>	Pod-level security context	see <code>values.yaml</code>
<code>redis.securityContext</code>	Container-level security context	see <code>values.yaml</code>
<code>redis.nodeSelector</code>	Node selector	<code>{}</code>
<code>redis.tolerations</code>	Tolerations	<code>[]</code>
<code>redis.affinity</code>	Affinity rules	<code>{}</code>

Server

The `server` component is a single Go binary that serves both the backend API and the React frontend on port `3000`. It is deployed as a **StatefulSet**.

Parameter	Description	Default
<code>server.enabled</code>	Deploy the server	<code>true</code>
<code>server.image.registry</code>	Image registry	<code>ghcr.io</code>
<code>server.image.repository</code>	Image repository	<code>patchmon/patchmon-server</code>
<code>server.image.tag</code>	Image tag (overridden by <code>global.imageTag</code> if set)	<code>2.0.0</code>
<code>server.image.pullPolicy</code>	Image pull policy	<code>IfNotPresent</code>
<code>server.replicaCount</code>	Number of replicas	<code>1</code>
<code>server.updateStrategy.type</code>	StatefulSet update strategy	<code>RollingUpdate</code>
<code>server.jwtSecret</code>	JWT signing secret (required if <code>existingSecret</code> is not set)	<code>""</code>
<code>server.aiEncryptionKey</code>	AI encryption key (required if <code>existingSecret</code> is not set)	<code>""</code>
<code>server.existingSecret</code>	Name of an existing secret for JWT and AI encryption key	<code>""</code>
<code>server.existingSecretJwtKey</code>	Key for JWT secret inside the existing secret	<code>jwt-secret</code>

Parameter	Description	Default
<code>server.existingSecretAiEncryptionKey</code>	Key for AI encryption key inside the existing secret	<code>ai-encryption-key</code>
<code>server.resources.requests.cpu</code>	CPU request	<code>10m</code>
<code>server.resources.requests.memory</code>	Memory request	<code>256Mi</code>
<code>server.resources.limits.cpu</code>	CPU limit	<code>2000m</code>
<code>server.resources.limits.memory</code>	Memory limit	<code>2Gi</code>
<code>server.autoscaling.enabled</code>	Enable HPA	<code>false</code>
<code>server.autoscaling.minReplicas</code>	Minimum replicas	<code>1</code>
<code>server.autoscaling.maxReplicas</code>	Maximum replicas	<code>10</code>
<code>server.autoscaling.targetCPUUtilizationPercentage</code>	Target CPU utilisation	<code>80</code>
<code>server.autoscaling.targetMemoryUtilizationPercentage</code>	Target memory utilisation	<code>80</code>
<code>server.service.type</code>	Service type	<code>ClusterIP</code>
<code>server.service.port</code>	Service port	<code>3000</code>
<code>server.service.annotations</code>	Service annotations	<code>[]</code>
<code>server.livenessProbe.enabled</code>	Enable liveness probe (TCP on port 3000)	<code>true</code>
<code>server.readinessProbe.enabled</code>	Enable readiness probe (<code>GET /health</code> on port 3000)	<code>true</code>
<code>server.initContainers.waitForDatabase.enabled</code>	Wait for database before starting	<code>true</code>
<code>server.initContainers.waitForRedis.enabled</code>	Wait for Redis before starting	<code>true</code>
<code>server.initContainers.waitForGuacd.enabled</code>	Wait for guacd before starting	<code>true</code>
<code>server.extraEnv</code>	Extra environment variables to inject into the server container	<code>[]</code>
<code>server.extraVolumeMounts</code>	Extra volume mounts for the server container	<code>[]</code>
<code>server.extraVolumes</code>	Extra volumes to add to the server pod	<code>[]</code>
<code>server.podAnnotations</code>	Pod annotations	<code>{}</code>
<code>server.podSecurityContext</code>	Pod-level security context	see <code>values.yaml</code>
<code>server.securityContext</code>	Container-level security context	see <code>values.yaml</code>
<code>server.nodeSelector</code>	Node selector	<code>{}</code>
<code>server.tolerations</code>	Tolerations	<code>[]</code>

Parameter	Description	Default
<code>server.affinity</code>	Affinity rules	<code>{}</code>
<code>server.topologySpreadConstraints</code>	Topology spread constraints	<code>[]</code>

Server Environment Variables

Parameter	Description	Default
<code>server.env.enableLogging</code>	Enable application logging	<code>true</code>
<code>server.env.logLevel</code>	Log level (<code>trace</code> , <code>debug</code> , <code>info</code> , <code>warn</code> , <code>error</code>)	<code>info</code>
<code>server.env.logToConsole</code>	Log to stdout	<code>true</code>
<code>server.env.serverProtocol</code>	Protocol used by agents to reach the server (<code>http</code> or <code>https</code>)	<code>http</code>
<code>server.env.serverHost</code>	Hostname used by agents to reach the server	<code>patchmon.example.com</code>
<code>server.env.serverPort</code>	Port used by agents (<code>80</code> or <code>443</code>)	<code>80</code>
<code>server.env.corsOrigin</code>	CORS allowed origin (should match the URL users access in a browser)	<code>http://patchmon.example.com</code>
<code>server.env.dbConnectionLimit</code>	Database connection pool limit	<code>30</code>
<code>server.env.dbPoolTimeout</code>	Pool timeout in seconds	<code>20</code>
<code>server.env.dbConnectTimeout</code>	Connection timeout in seconds	<code>10</code>
<code>server.env.dbIdleTimeout</code>	Idle connection timeout in seconds	<code>300</code>
<code>server.env.dbMaxLifetime</code>	Max connection lifetime in seconds	<code>1800</code>
<code>server.env.rateLimitWindowMs</code>	General rate limit window (ms)	<code>900000</code>
<code>server.env.rateLimitMax</code>	General rate limit max requests	<code>5000</code>
<code>server.env.authRateLimitWindowMs</code>	Auth rate limit window (ms)	<code>600000</code>
<code>server.env.authRateLimitMax</code>	Auth rate limit max requests	<code>500</code>
<code>server.env.agentRateLimitWindowMs</code>	Agent rate limit window (ms)	<code>60000</code>
<code>server.env.agentRateLimitMax</code>	Agent rate limit max requests	<code>1000</code>
<code>server.env.redisDb</code>	Redis database index	<code>0</code>
<code>server.env.trustProxy</code>	Trust proxy headers -- set to a CIDR range or <code>true</code> when behind a reverse proxy	<code>10.0.0.0/8,172.16.0.0/12,192.168.0.0/16</code>
<code>server.env.enableHsts</code>	Enable HSTS header	<code>false</code>
<code>server.env.defaultUserRole</code>	Default role for new users	<code>user</code>
<code>server.env.autoCreateRolePermissions</code>	Auto-create role permissions	<code>false</code>

OIDC / SSO Configuration

Parameter	Description	Default
<code>server.oidc.enabled</code>	Enable OIDC authentication	<code>false</code>
<code>server.oidc.issuerUrl</code>	OIDC issuer URL	<code>""</code>
<code>server.oidc.clientId</code>	OIDC client ID	<code>""</code>
<code>server.oidc.clientSecret</code>	OIDC client secret (required if <code>existingSecret</code> not set)	<code>""</code>
<code>server.oidc.existingSecret</code>	Existing secret containing the OIDC client secret	<code>""</code>
<code>server.oidc.existingSecretClientSecretKey</code>	Key inside the existing secret	<code>oidc-client-secret</code>
<code>server.oidc.scopes</code>	OIDC scopes	<code>openid profile email</code>
<code>server.oidc.buttonText</code>	Login button text	<code>Login with SSO</code>
<code>server.oidc.autoCreateUsers</code>	Auto-create users on first OIDC login	<code>true</code>
<code>server.oidc.defaultRole</code>	Default role for OIDC-created users	<code>user</code>
<code>server.oidc.syncRoles</code>	Sync roles from OIDC group claims	<code>true</code>
<code>server.oidc.disableLocalAuth</code>	Disable local username/password authentication	<code>false</code>
<code>server.oidc.sessionTtl</code>	OIDC session TTL in seconds	<code>86400</code>
<code>server.oidc.groups.superadmin</code>	OIDC group mapped to the superadmin role	<code>""</code>
<code>server.oidc.groups.admin</code>	OIDC group mapped to the admin role	<code>""</code>
<code>server.oidc.groups.hostManager</code>	OIDC group mapped to the hostManager role	<code>""</code>
<code>server.oidc.groups.user</code>	OIDC group mapped to the user role	<code>""</code>
<code>server.oidc.groups.readonly</code>	OIDC group mapped to the readonly role	<code>""</code>

Guacd

Apache Guacamole proxy daemon used for browser-based SSH and remote desktop sessions.

Parameter	Description	Default
<code>guacd.enabled</code>	Deploy guacd	<code>true</code>
<code>guacd.image.registry</code>	Image registry	<code>docker.io</code>
<code>guacd.image.repository</code>	Image repository	<code>guacamole/guacd</code>

Parameter	Description	Default
<code>guacd.image.tag</code>	Image tag	<code>latest</code>
<code>guacd.image.pullPolicy</code>	Image pull policy	<code>IfNotPresent</code>
<code>guacd.replicaCount</code>	Number of replicas	<code>1</code>
<code>guacd.updateStrategy.type</code>	Deployment update strategy	<code>RollingUpdate</code>
<code>guacd.resources.requests.cpu</code>	CPU request	<code>10m</code>
<code>guacd.resources.requests.memory</code>	Memory request	<code>32Mi</code>
<code>guacd.resources.limits.cpu</code>	CPU limit	<code>1000m</code>
<code>guacd.resources.limits.memory</code>	Memory limit	<code>512Mi</code>
<code>guacd.livenessProbe.enabled</code>	Enable liveness probe (TCP)	<code>true</code>
<code>guacd.readinessProbe.enabled</code>	Enable readiness probe (TCP)	<code>true</code>
<code>guacd.service.type</code>	Service type	<code>ClusterIP</code>
<code>guacd.service.port</code>	Service port	<code>4822</code>
<code>guacd.podAnnotations</code>	Pod annotations	<code>{}</code>
<code>guacd.podSecurityContext</code>	Pod-level security context	see <code>values.yaml</code>
<code>guacd.securityContext</code>	Container-level security context	see <code>values.yaml</code>
<code>guacd.nodeSelector</code>	Node selector	<code>{}</code>
<code>guacd.tolerations</code>	Tolerations	<code>[]</code>
<code>guacd.affinity</code>	Affinity rules	<code>{}</code>

Ingress

Parameter	Description	Default
<code>ingress.enabled</code>	Enable ingress resource	<code>true</code>
<code>ingress.className</code>	Ingress class name	<code>""</code>
<code>ingress.annotations</code>	Ingress annotations	see <code>values.yaml</code>
<code>ingress.hosts</code>	List of ingress host rules	see <code>values.yaml</code>
<code>ingress.tls</code>	TLS configuration	<code>[]</code> (disabled)

The default ingress annotations enable WebSocket support and tune proxy timeouts for agent connections:

```
ingress:
  annotations:
```

```
nginx.ingress.kubernetes.io/proxy-body-size: "0"
nginx.ingress.kubernetes.io/proxy-read-timeout: "3600"
nginx.ingress.kubernetes.io/proxy-send-timeout: "3600"
nginx.ingress.kubernetes.io/proxy-connect-timeout: "60"
nginx.ingress.kubernetes.io/proxy-http-version: "1.1"
nginx.ingress.kubernetes.io/client-body-buffer-size: "4m"
nginx.ingress.kubernetes.io/websocket-services: "server"
```

Other

Parameter	Description	Default
<code>serviceAccount.create</code>	Create a ServiceAccount	<code>false</code>
<code>serviceAccount.annotations</code>	ServiceAccount annotations	<code>{}</code>
<code>serviceAccount.name</code>	ServiceAccount name	<code>""</code>
<code>configMap.create</code>	Create the application ConfigMap	<code>true</code>
<code>configMap.annotations</code>	ConfigMap annotations	<code>{}</code>
<code>secret.create</code>	Create the chart-managed Secret (disable when using an external secret)	<code>true</code>
<code>secret.annotations</code>	Secret annotations	<code>{}</code>

Persistent Volumes

The chart creates the following PersistentVolumeClaims:

PVC	Component	Purpose	Default Size
<code>postgres-data</code>	Database	PostgreSQL data directory	<code>5Gi</code>
<code>redis-data</code>	Redis	Redis data directory	<code>5Gi</code>

All PVCs respect the `global.storageClass` setting unless overridden at the component level.

Updating PatchMon

Using `global.imageTag`

The simplest way to update the server image is to set `global.imageTag`:

```
helm upgrade patchmon oci://ghcr.io/ruthlessbeat200/charts/patchmon \
  -n patchmon \
  -f values-prod.yaml \
  --set global.imageTag=2.1.0
```

When `global.imageTag` is set it overrides `server.image.tag`.

Pinning the server tag individually

```
server:
  image:
    tag: "2.0.0"
```

Upgrading the chart version

```
# Upgrade with new values
helm upgrade patchmon oci://ghcr.io/ruthlessbeat200/charts/patchmon \
  --namespace patchmon \
  --values values-prod.yaml

# Upgrade and wait for rollout
helm upgrade patchmon oci://ghcr.io/ruthlessbeat200/charts/patchmon \
  --namespace patchmon \
  --values values-prod.yaml \
  --wait --timeout 10m
```

Check the [releases page](#) for version-specific changes and migration notes.

Uninstalling

```
# Uninstall the release
helm uninstall patchmon -n patchmon

# Clean up PVCs (optional -- this deletes all data)
kubectl delete pvc -n patchmon -l app.kubernetes.io/instance=patchmon
```

Advanced Configuration

Custom Image Registry

Override the registry for all images (useful for air-gapped environments or private mirrors):

```
global:
  imageRegistry: "registry.example.com"
```

This changes every image pull to use the specified registry:

- `registry.example.com/postgres:18-alpine`
- `registry.example.com/redis:8-alpine`
- `registry.example.com/guacamole/guacd:latest`
- `registry.example.com/patchmon/patchmon-server:2.0.0`
- `registry.example.com/busybox:latest` (init containers)

Without `global.imageRegistry`, components use their default registries (`docker.io` for database/Redis/guacd, `ghcr.io` for server).

Multi-Tenant Deployment

Deploy multiple isolated instances in separate namespaces using `fullnameOverride`:

```
fullnameOverride: "patchmon-tenant-a"

server:
  env:
    serverHost: tenant-a.patchmon.example.com
    corsOrigin: https://tenant-a.patchmon.example.com

ingress:
  hosts:
    - host: tenant-a.patchmon.example.com
  paths:
    - path: /
      pathType: Prefix
  service:
```

```
name: server
port: 3000
```

Horizontal Pod Autoscaling

```
server:
  autoscaling:
    enabled: true
    minReplicas: 2
    maxReplicas: 10
    targetCPUUtilizationPercentage: 70
    targetMemoryUtilizationPercentage: 80
```

Using an External Database

Disable the built-in database and point the server at an external PostgreSQL instance:

```
database:
  enabled: false
  host: "my-postgres.example.com"
  port: 5432
  auth:
    database: patchmon_db
    username: patchmon_user
    existingSecret: patchmon-secrets
    existingSecretPasswordKey: postgres-password
```

Injecting Extra Environment Variables

Use `server.extraEnv` to pass additional environment variables, for example to trust a custom CA for OIDC:

```
server:
  extraEnv:
    - name: NODE_EXTRA_CA_CERTS
      value: /etc/ssl/certs/my-ca.crt
  extraVolumeMounts:
    - name: my-ca
```

```
    mountPath: /etc/ssl/certs/my-ca.crt
    subPath: ca.crt
    readOnly: true
extraVolumes:
  - name: my-ca
    configMap:
      name: my-ca-configmap
```

OIDC / SSO Integration

```
server:
  oidc:
    enabled: true
    issuerUrl: "https://auth.example.com/realms/master"
    clientId: "patchmon"
    clientSecret: "your-client-secret"
    scopes: "openid profile email groups"
    buttonText: "Login with SSO"
    autoCreateUsers: true
    syncRoles: true
  groups:
    superadmin: "patchmon-admins"
    admin: ""
    hostManager: ""
    user: ""
    readonly: ""
```

Troubleshooting

Check pod status

```
kubectl get pods -n patchmon
kubectl describe pod <pod-name> -n patchmon
kubectl logs <pod-name> -n patchmon
```

Check init container logs

```
kubectl logs <pod-name> -n patchmon -c wait-for-database
kubectl logs <pod-name> -n patchmon -c wait-for-redis
kubectl logs <pod-name> -n patchmon -c wait-for-guacd
```

Check service connectivity

```
# Test database connection
kubectl exec -n patchmon -it statefulset/patchmon-prod-server -- nc -zv patchmon-prod-database
5432

# Test Redis connection
kubectl exec -n patchmon -it statefulset/patchmon-prod-server -- nc -zv patchmon-prod-redis
6379

# Test guacd connection
kubectl exec -n patchmon -it statefulset/patchmon-prod-server -- nc -zv patchmon-prod-guacd
4822

# Check server health
kubectl exec -n patchmon -it statefulset/patchmon-prod-server -- wget -q0-
http://localhost:3000/health
```

Common issues

Symptom	Likely cause	Fix
Pods stuck in <code>Init</code> state	Database, Redis, or guacd not yet running	Check StatefulSet/Deployment events: <code>kubectl describe sts -n patchmon</code>
PVC stuck in <code>Pending</code>	No matching StorageClass or no available PV	Verify storage class exists: <code>kubectl get sc</code>
<code>ImagePullBackOff</code>	Registry credentials missing or incorrect image reference	Check <code>imagePullSecrets</code> and image path
Ingress returns 404 / 502	Ingress controller not installed or misconfigured path rules	Verify controller pods and ingress resource: <code>kubectl describe ingress -n patchmon</code>
WebSocket disconnects	Missing WebSocket annotations on ingress	Ensure <code>nginx.ingress.kubernetes.io/websocket-services: "server"</code> and proxy timeout annotations are set

Symptom	Likely cause	Fix
<code>secret ... not found</code>	Required secret was not created before install	Create the secret or set <code>secret.create: true</code> with inline passwords

Development

Lint the chart

```
helm lint .
```

Render templates locally

```
# Render with default values
helm template patchmon . --values values-quick-start.yaml

# Render with production values
helm template patchmon . --values values-prod.yaml

# Debug template rendering
helm template patchmon . --values values-quick-start.yaml --debug
```

Dry-run installation

```
helm install patchmon . \
  --namespace patchmon \
  --dry-run --debug \
  --values values-quick-start.yaml
```

Support

- GitHub Issues: github.com/RuTHlessBEat200/PatchMon-helm/issues
- Application repository: github.com/PatchMon/PatchMon

