KLUTCH
sponsored by anynines

A centralized control plane for data services across Kubernetes clusters

HOW KLUTCH WORKS

Klutch gives platform teams centralized governance over databases, so developers get self-service without ops losing control.

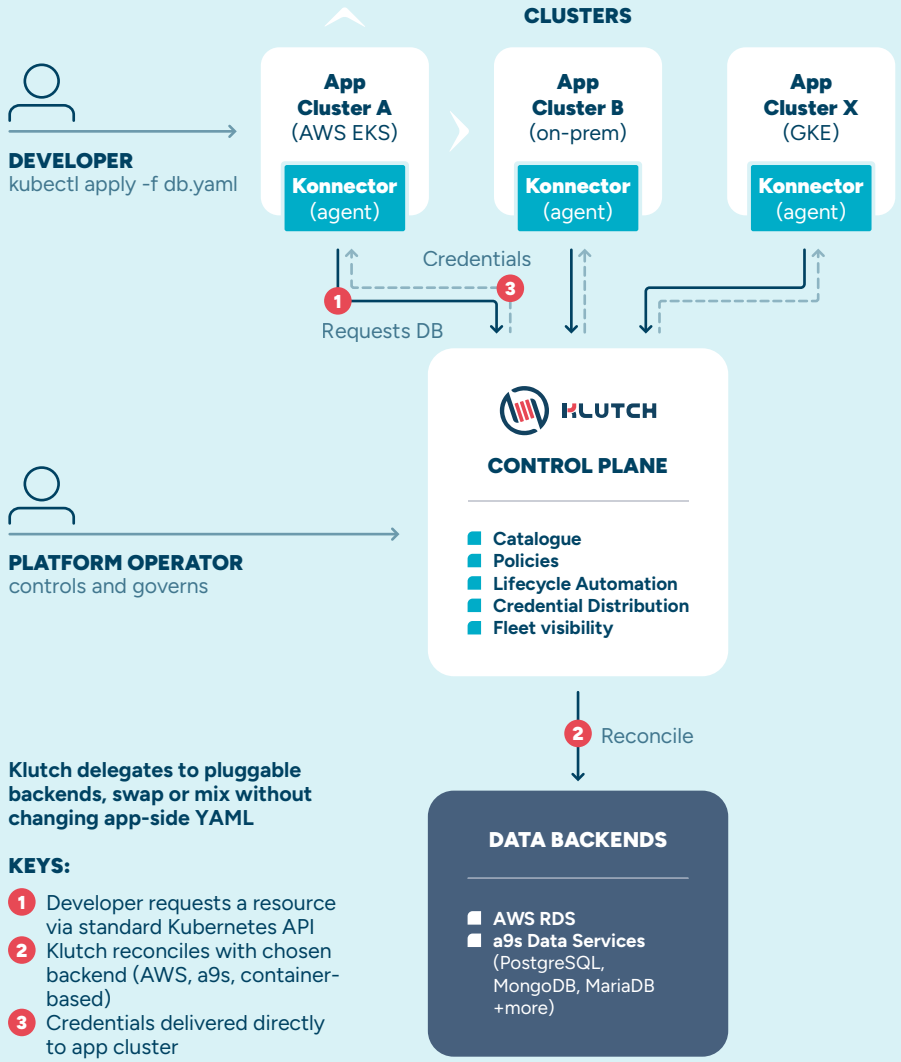

Klutch delegates to pluggable backends, swap or mix without changing app-side YAML

KEYS:

1. Developer requests a resource via standard Kubernetes API
2. Klutch reconciles with chosen backend (AWS, a9s, container-based)
3. Credentials delivered directly to app cluster

WHAT CHANGES FOR YOUR TEAM

Self-Service

Developer writes a YAML spec for a PostgreSQL instance. No tickets, no waiting for ops.

Central Control

Klutch's control plane provisions it in the cloud or on-prem, against your approved policies. Backups, encryption, and resource limits are already baked in.

Transparency

Platform team sees every instance fleet-wide, with one observability surface. Lifecycle tasks, upgrades, restores, are automated.

HOW TO GET STARTED

1 **Download Klutch**

Get the latest release from GitHub
github.com/anynines/klutchio

2 **Deploy the control plane**

Install into your management Kubernetes cluster

3 **Connect your clusters**

Deploy agents and configure your data service backends

4 **Start managing databases**

Provision, scale, and govern with full fleet visibility

BUILT FOR BOTH SIDES OF THE PLATFORM

FOR APPLICATION DEVELOPERS

- Self-service DB provisioning via simple YAML specs
- Kubernetes-native, works with kubectl and GitOps
- Apply in your cluster, The DB runs fully managed elsewhere
- Consistent experience across cloud providers

FOR PLATFORM TEAMS

- Central control plane, full fleet visibility
- Pluggable automation backends for any infrastructure
- Fine-grained access control and audit trail
- Lifecycle automation, backups, scaling, upgrades

	WITHOUT KLUTCH	WITH KLUTCH
Policy & governance	Compliance lives in docs, not infrastructure, drift is invisible until it's an incident	Configuration enforced at the platform level, not left to individual cluster admins
Audit & access	No central record, compliance reviews mean chasing logs across clusters	All provisioning through one control plane, one audit trail, access enforced per cluster
Developer self-service	Every DB request goes through a ticket queue, platform team becomes the bottleneck	Developer teams use kubectl or GitOps best practices, infrastructure provisioned automatically, no ticket, no wait
Cluster onboarding	Every new Kubernetes cluster needs full infrastructure replication before teams can use shared services	One agent, one bind command, full approved service catalog inherited instantly
Credential delivery	Secrets provisioned manually per cluster or wired up with separate tooling	Credentials auto-generated and synced into the app cluster namespace, no extra tooling
Service rollout	New service means touching every cluster, inconsistency scales with you	Define once, choose who gets it, bound clusters sync automatically

HOW IS IT BUILT

1. **Klutch control plane**
   Centralized database orchestration and governance, sits above all app clusters

2. **Cluster agents (Konnector)**
   Local coordinators within each app cluster, pull-based, never expose the control plane

3. **Service registry**
   Tracks where every database is running and which application is consuming it

4. **Automation backends**
   Bring your own automation backend or use one of the turnkey solutions provided by the a9s Hub: a9s (VM), a8s (container), AWS RDS

PRODUCTION GRADE AUTOMATION BACKENDS BY A9S HUB— same kubernetes api regardless of where data lives.

a9s Data Services (VM-based)

- PostgreSQL
- MongoDB
- MariaDB
- KeyValue (Redis-compatible)
- Messaging (Kafka/RabbitMQ)
- Search (Elasticsearch-compatible)
- Prometheus · LogMe2

a8s Data Services (Container-native)
PostgreSQL (Kubernetes operator, HA)
Additional a8s services via a9s Hub

AWS managed (Cloud-native)
Amazon RDS and S3
Well-architected framework

NOT SEEING YOUR SERVICE? Klutch supports custom backends via Crossplane providers, integrate any data service your organisation runs.

PRODUCTION USE | A9S HUB

Need more than open source?

- Professional support and SLAs
- Additional security enhancements
- Custom automation backend integration
- Advanced networking
- Team-based tenancy
- CI/CD-compatible Klutch API binding
- Usage tracking and detailed reporting

KLUTCH FOR ENTERPRISES
anynines.com/products/klutch

Join the community

Klutch is built on open-source collaboration. We invite engineers passionate about cloud-native data management to try it, contribute, and shape its roadmap.

GITHUB
github.com/anynines/klutchio

DOCS
klutch.io/docs

DISCUSSIONS
github.com/anynines/klutchio/discussions

10+ Data service types

3 backend types: VM | container | cloud

v1.5.x Actively maintained

100% open source Apache 2.0

KLUTCH sponsored by anynines

Open Source via Apache 2.0

klutch.io
info@klutch.io
/company/klutch-io