FreeOS ShadowOps — Executive Summary

Autonomous AI DevOps for Sovereign Infrastructure

eHealthBrains ApS

March 2026

FreeOS ShadowOps

Autonomous AI DevOps for Sovereign Infrastructure

The Problem

Modern infrastructure is complex, fragile, and expensive to operate. Organisations running private cloud, healthcare IT, or regulated workloads face three compounding pressures:

1. The talent gap. Experienced DevOps engineers are scarce and expensive. 24/7 on-call coverage requires at minimum two engineers — often more. Hiring, retaining, and developing this talent in a competitive market costs €150,000–250,000 per year per FTE, not counting burnout, turnover, and knowledge loss.

2. The tooling tax. Commercial monitoring platforms — Datadog, PagerDuty, Dynatrace — promise visibility but deliver alert storms. They are cloud-hosted, expensive, and fundamentally reactive. They tell you something broke. They do not fix it.

3. The sovereignty gap. Hyperscaler managed services (AWS EKS, Azure AKS, GCP GKE) abstract away operational complexity — but at the cost of data sovereignty, fiscal unpredictability, and vendor lock-in. For healthcare, finance, and government workloads, this is not acceptable.

FreeOS ShadowOps was built to solve all three at once.

What FreeOS ShadowOps Is

FreeOS ShadowOps is an **AI-native autonomous DevOps platform** that operates continuously alongside any Kubernetes or Docker infrastructure. It observes, reasons, acts, escalates, and heals — without requiring a human engineer to be on-call.

The platform consists of three layers:

Layer	Component	Role
Master	FreeOSBot	Central AI agent — reasoning, decision-making, GitOps execution, human interface
Sensor	Watchmen	Domain-specific monitoring agents — one per concern, each with its own on-prem AI

Layer	Component	Role
Resilience	FreeOSBot-Twin	Standby HA master — auto-promotes on primary failure, loads full context in seconds

Every component runs **entirely on-premises**. No data leaves your infrastructure. No hyperscaler dependency. No SaaS subscription required beyond the AI model API (and even that is replaceable with a local model).

Key Capabilities

Continuous Autonomous Monitoring Watchmen poll the cluster every 60 seconds across five domains: cluster health, infrastructure, certificates, backup integrity, and security posture. Every alert is enriched by a local AI model before it reaches the operator.

AI-Enriched Triage When a Watchman detects an anomaly, it asks its local Ollama AI instance for a root cause hypothesis and recommended action. The operator receives not just "pod CrashLoopBackOff" but a structured analysis of likely cause, urgency, and recommended action — before they open a terminal.

Autonomous Remediation — Within Boundaries FreeOSBot executes routine remediation without waking the operator: restarting crashed pods, re-syncing ArgoCD applications, rotating short-lived secrets. For irreversible or high-risk actions, it escalates immediately with full context.

High-Availability Master Agent FreeOSBot-Twin runs in standby in the cluster. If the primary fails, the twin detects the stale heartbeat, collects crash diagnostics via SSH, generates an AI root cause analysis, and presents itself to the operator on a dedicated Telegram bot — fully context-loaded from the last session export. Maximum context gap: 15 minutes.

Sovereign and On-Premises All AI inference runs on-prem via Ollama. All configuration lives in a private GitOps repository. All secrets remain in Kubernetes Secrets. No telemetry, no cloud dependency, no data exfiltration.

GitOps Native Every infrastructure change is committed to Git and applied via ArgoCD. FreeOSBot never makes ad-hoc changes. The entire platform state is reproducible from the repository.

The Platform Has Eyes, Ears, and a Voice

> *"The platform now has eyes, ears, and a voice — and it knows when to wake up its human."*

FreeOSBot watches your infrastructure continuously. It hears every alert the Watchmen raise. It speaks to you on Telegram — briefly when everything is fine, with full diagnostic detail when something needs your attention. And when it cannot handle something alone, it says so clearly, provides everything you need to act, and waits for your instruction.

Headline Numbers

Metric	Value
DevOps FTE displaced	1.0–1.5
Commercial monitoring tooling replaced	Datadog, PagerDuty, custom alert scripts
AI triage coverage	100% of alerts enriched before delivery
Mean time to operator notification	< 2 minutes
HA failover time	< 10 minutes
Context gap on failover	< 15 minutes
Data sovereignty	100% on-premises
Infrastructure cost model	Flat monthly — hardware + AI model API only

---

Call to Action

FreeOS ShadowOps is production-ready and actively operating the FreeOSCloud HealthCloud platform. It is designed to be deployed on any sovereign Kubernetes or Docker environment in under one business day.

Henrik L. Ibsen, CEO eHealthBrains ApS | www.ehealthbrains.com | hibsen@ehealthbrains.com