

Keel by Infranexis

AI-Native Log Analytics You Own

Self-hosted log analytics for teams that need their log data to stay on infrastructure they own — in your cloud or your datacenter. Ask questions in plain English, get a root cause on every incident, and use the AI model you already trust.

Executive Overview

Log volume keeps climbing, and most teams end up paying for it twice: first in licensing that scales with every new service, then in the engineering time it takes to keep a self-managed stack running. In regulated industries there's a third cost — once logs leave for a vendor's cloud, you can be out of compliance before you've investigated anything.

Keel is an appliance you run yourself, in your own cloud or datacenter. Ingest, storage, search, alerting, incident tracking, and AI arrive in one VM image — priced per appliance rather than per gigabyte, and able to run fully air-gapped.

In practice, that lets your engineering and security teams:

-  **Ask their logs anything in plain English**
-  **Get root cause automatically on every incident**
-  **Bring the AI model they trust**
-  **Keep their log data on infrastructure they own**

The point is simple: spend less time digging through logs by hand, and keep your data where it belongs.



Keel

-  **Ask Keel Anything**
-  **Dynamic Charts & Dashboards**
-  **Keel MCP Server for AI agents**
-  **Automatic Root-Cause Analysis**
-  **Detection & Incident Management**
-  **Bring-Your-Own-Model AI**



Ingest. Store. Ask. Act.

Keel runs as one pipeline — from a raw log line to an answered question or a closed incident. There are no separate clusters to operate and no plugins to keep patched.



Ingest

- OTLP gRPC & HTTP, syslog, and Fluentd out of the box, plus a REST API
- Any logs, IT or OT — apps, servers, network & security gear, and OT/ICS systems
- Connectors for AWS, Azure, GCP, Datadog and more; no proprietary agent lock-in



Store

- High-compression column store indexed for millisecond search
- Hot local storage plus an optional S3-compatible cold tier that stays queryable
- Retention is configurable and limited only by your own disk



Ask

- Ask anything in plain English — get a result, a chart, or a summary
- Save queries and pin them to dashboards; read-only SQL in the same window
- Connect your own AI agents over MCP — read-only and scoped



Act

- When a detection fires, Keel investigates on its own and returns a root cause
- Incidents carry a full timeline and the log entries behind them
- Routing to Slack, PagerDuty, email, SNMP traps, or a webhook

What Makes Keel Different

Flat Pricing, Not Per-Gigabyte

You license Keel per appliance, so the bill stays flat as log volume grows. Teams moving off volume-based platforms typically cut 80–90% of their logging spend.

Bring Your Own AI Model

Point Keel at your own Anthropic, OpenAI, or Google Gemini key, any OpenAI-compatible endpoint, or a local model — with automatic failover. Only the context for a request leaves your network; run local and nothing leaves at all.

Your Own AI Agents, Over MCP

Keel ships a Model Context Protocol server, so the agents your team already uses — Claude, Cursor, your own — can query your logs directly. Read-only, scoped per key, rate-limited. It runs inside your network, so your agents reach your logs without your logs reaching anyone's cloud.

One Appliance, Not a Stack

Ingest, storage, search, AI, alerting, and HA ship in a single VM image. There's no cluster to tune, no plugins to patch, and no need for a full-time platform engineer to keep it alive.

It Investigates, Not Just Alerts

Most tools tell you something broke. Keel tells you why: when a detection fires, it correlates the surrounding logs and your runbooks and writes up the likely root cause — no analyst required.

Self-Hosted — Cloud or On-Prem

Runs on infrastructure you own: your datacenter, your private cloud, or your own AWS, Azure, or GCP account. Air-gap capable, with no phone-home and no vendor cloud holding your logs.

Core Capabilities



Ask Keel Anything

Ask your logs a question in plain English and get an answer — no query language to learn first.

- Plain-English questions become a result, a chart, or a summary
- Save a question and pin it to a dashboard
- Read-only SQL in the same window for power users

Every answer shows the log lines behind it, so your engineers can check the work.



Detection Engine

Turn the patterns that matter into incidents automatically — anomalies, missing logs, thresholds, and more.

- Plain-English rules, with a recommended pack to start from
- AI compiles each rule once; evaluation runs with no LLM in the hot path
- Host-scoped incidents, one root-cause analysis per episode

Nothing calls a model in the hot path, so cost and latency stay predictable at volume.



Keel MCP Server

Connect the AI agents your team already uses straight to your logs and Keel's intelligence.

- Works with Claude, Cursor, or your own agents over MCP
- Read-only, scoped per key, and rate-limited
- Search logs, incidents, and RCA — never change anything

Your agents reach your logs without your logs reaching anyone's cloud.



Ingest — IT and OT

Logs arrive over open standards, so existing instrumentation keeps working and nothing is tied to a proprietary agent.

- OTLP gRPC / HTTP, syslog, Fluentd, and a REST API
- Apps, servers, network & security gear, and OT/ICS systems
- Connectors for AWS, Azure, GCP; a generic webhook for the rest

Each source is isolated, so adding or removing one during a migration never disturbs the others.



BI Connector

Point your existing BI stack at Keel over a read-only, ClickHouse-compatible SQL endpoint.

- Works with Grafana, Metabase, Superset, and Power BI
- Dedicated read-only user with rotatable credentials
- Queries run on your live data — no export, no second copy

Your tools and your dashboards — and the data never leaves Keel.



High Availability

A three-node cluster keeps ingest and search running when a node fails, with no manual intervention.

- 3-node HA cluster with automatic failover
- Virtual IP switchover, quorum maintained



Automatic Root-Cause Analysis

When an incident fires, Keel investigates on its own and returns a written root cause.

- Correlates logs across services and reads your runbooks
- Plain-English output to Slack, PagerDuty, or email
- One analysis per incident — no noise, no repeat spend

Because it reads your own documentation, first response gets faster without more headcount.



Bring Your Own AI Model

Keel ships no bundled model — you point every AI feature at the model you trust.

- Your own Anthropic, OpenAI, or Gemini key, any OpenAI-compatible endpoint, or a local model
- A primary backend plus automatic failover
- Only the request's context leaves — or nothing at all, in local / air-gapped mode

You get modern AI and data residency at the same time, on your terms rather than a vendor's.



Knowledge Base & Assistant

Turn your internal documentation into a source Keel can draw on for every investigation.

- Upload runbooks, architecture diagrams, and SOPs
- Used automatically to ground every root-cause analysis
- A chat assistant that investigates — and acts only with your confirmation

Knowledge stays useful as teams turn over, and a human signs off on every change.



Log Pipelines

Shape data on the way in with rules defined per source rather than globally.

- Parse, filter, mask PII, enrich, deduplicate, sample, and classify
- Rules scoped to individual sources
- Noise dropped before it consumes storage

Dropping noise at ingest keeps retention costs down and search results clean.



Security & Access Control

Enterprise identity and audit controls, managed from the dashboard with no config files to hand-edit.

- Role-based access control with teams and roles
- SSO via LDAP/RADIUS, SAML, and OIDC, plus MFA
- Tamper-evident audit, your-own-TLS, and firewall management

Every action is logged, which matters when the platform itself is in audit scope.



Appliance Operations

Day-to-day administration happens in the web dashboard, so keeping the platform current isn't a project.

- Guided first-run setup wizard
- Signed in-place updates with rollback, backup and restore

Built For

- Platform & DevOps
- Security & SOC
- MSPs & MSSPs
- Mid-market & Enterprise
- Regulated industries
- Air-Gapped & OT

Where Teams Use It

**Incident Investigation & Triage**

When an alert fires, Keel has already correlated the surrounding logs and your runbooks — so it arrives with a likely root cause, not just a notification.

**Log Platform Consolidation**

Replace volume-based licensing and self-managed stacks with one appliance. Run it alongside your current tool and migrate sources over at your own pace.

**Compliance & Audit Retention**

Keep logs as long as regulation requires, on storage you provision, with RBAC, SSO, and a full audit trail of who searched what.

**Sovereign & Air-Gapped Operations**

Run full AI-assisted analysis inside isolated networks using a local model — no phone-home, delivered on USB or secure transfer.

How Keel Compares

How Keel stacks up against the log platforms most teams run today — whether that's a commercial SaaS tool or a self-managed open-source stack.

Capability	Keel	Typical Log Platforms
Self-hosted & air-gapped deployment	Built for it	Cloud-only or heavy setup
Cost model	Flat, per appliance	Scales with ingest volume
Packaging & high availability	One VM image, HA included	Multi-tier or DIY assembly
Ask anything in plain English	Native	Add-on or unavailable
Where the AI model runs	Your key, or fully local	Vendor-hosted cloud model
Your own AI agents (MCP)	Read-only, scoped	Not available
OpenTelemetry-native ingest	Yes, no proprietary agent	Partial or agent-locked
Retention limit	Your disk, your call	Capped or billed per GB

Runs Anywhere You Run VMs

It's a pre-configured image: import it, power on, and run the guided setup wizard. Our team handles the initial configuration during your POC.

**VMware vSphere / ESXi**

OVA via Deploy OVF Template

**KVM / QEMU / Proxmox**

QCOW2 via virt-install or UI

**AWS / Azure / GCP**

Your own account (image import; marketplace on the roadmap)

**OpenStack**

QCOW2 to Glance, boot via Nova

**OpenShift Virtualization**

Import via KubeVirt / CDI

**Air-Gapped Delivery**

USB or secure transfer