


● APPLICATION & AI CONTROL CENTRE

# Think *It.* Build *It.* Protect *It.*

Loglevel Sentinel gives you the tools to build, run and secure custom software.

It turns ideas into working business systems in minutes, built by your own people. It brings in the systems you already have — and keeps every one of them, and every AI request they make, under your control. On your own hardware, or in the cloud of your choice.

[sentinel.loglevel.com](https://sentinel.loglevel.com)



● AI SPEND TODAY LIVE  
**\$149.61**  
1,288 AI requests · 4 stopped


● AI request scored  
Safe · 12 / 100 · Sales **\$0.1842**

● Stopped before sending  
Prompt injection · 84 / 100 **Stopped**

● Change package ready  
Training dashboard · 3 steps

● AutoFix applied  
Nordic Comply · error resolved

● AI request · local model Free  
HR · Policy search



## — WHAT IS SENTINEL?

# Every team needs its own customized system. Sentinel builds them — and keeps them all in one place.

Sales wants deal approvals. People wants onboarding. Operations wants shift planning. Hundreds of useful applications can quickly become a burden for IT. Sentinel is the control platform that makes them easy: it builds new systems, imports the ones you already have, runs each in its own container, and watches over all of them — including every question they send to AI.



01

## Build

Describe the system you need in plain words. Sentinel plans it, writes it, checks it and deploys it.

02

## Monitor

See who is using what right now, how every system is doing, and what AI is being asked and what it costs.

03

## Protect

A firewall that knows your systems, an AI gateway that scores every question, and lockdown when it matters.

## 01 THINK IT

# Say what you need, *in plain words.*

A team needs a training dashboard, a compliance register or a smarter search. Write it the way you would explain it to a colleague. Sentinel turns the idea into a change package — a short, focused plan you can read before a single line is written.

- ✓ Requests from users and administrators flow straight into the Builder
- ✓ Big ideas are split into small, reviewable steps
- ✓ Nothing changes until an administrator says go

The screenshot shows a web application window titled "Builder · from request to change package". The main heading is "Start a Change Package" with a subtext: "Choose the type and give the package a short task. The package groups several small development steps into one accepted change." There are two main selection boxes: "Change / bug fix" (FIX OR IMPROVE EXISTING FUNCTIONALITY.) and "New function / module" (BUILD A NEW, WELL-SCOPED PART OF THE SYSTEM.). Below these are input fields for "Package task" (containing "Add a training dashboard that shows each instructor's upcoming courses") and "Module" (containing "Choose the module this package works on"). A note states: "One administrator per module at a time. Other modules can be worked on in parallel." A "Best practices" section lists four bullet points: "Keep each package focused on one clear change.", "Break large changes or modules into small, related tasks.", "Complete and review each task before moving on.", and "Start a new package for unrelated work." At the bottom right are "Cancel" and "Start package" buttons.

Builder · from request to change package

**SENTINEL**

### Start a Change Package

Choose the type and give the package a short task. The package groups several small development steps into one accepted change.

**Change / bug fix**  
FIX OR IMPROVE EXISTING FUNCTIONALITY.

**New function / module**  
BUILD A NEW, WELL-SCOPED PART OF THE SYSTEM.

Package task

Add a training dashboard that shows each instructor's upcoming courses

Module

Choose the module this package works on

One administrator per module at a time. Other modules can be worked on in parallel.

**Best practices**



- Keep each package focused on one clear change.
- Break large changes or modules into small, related tasks.
- Complete and review each task before moving on.
- Start a new package for unrelated work.

Cancel Start package

# Sentinel builds it, *checks* *it and ships it.*

In minutes, the Builder writes the code, runs the checks and deploys the result into the system's own container. Every change is saved as a snapshot you can roll back. Every file is readable in the codebase browser, and every database can be opened, explored and queried from one screen.


- ✓ New systems start from templates and run in standard Docker containers
- ✓ Git history, snapshots and a full audit trail for every change
- ✓ Built-in database explorer with a SQL console — read-only by default



# Then it watches over *everything.*

Who is signed in, which systems are busy, how fast they answer, what AI was asked and what it cost — live. The firewall shows where traffic comes from and keeps out what shouldn't reach you. And if something ever goes wrong, one deliberate move locks every system down.

- ✓ Live requests, response times and people online, per system
- ✓ Access per system: everyone, approved networks only, or closed
- ✓ Lockdown with database dumps and a full incident record




— BUILD AND CHANGE IT YOURSELF

# *Minutes,* ~~not months.~~

A new system, a new field, a new report or a whole new module. The people who need it describe it in plain words, and Sentinel builds, checks and ships it while they wait. No development project, no waiting list, no consultants: **your own team is in charge of its own systems.**

## THE USUAL WAY

Weeks to months



## WITH SENTINEL

Minutes



### Your own people build it

Anyone who knows the work can ask for a change in plain words. No programming, no developers, no tickets.



### Change it when the business changes

A new field today, a new report tomorrow. Systems keep up with you instead of waiting for a supplier's next release.



### Safe to try

An administrator approves with one click, and every change keeps a snapshot you can roll back in seconds.



### No licensing fees

The systems you build are yours, source code included. No per-seat or per-system licence fees for anything you create.

# Hundreds of forgotten systems. Bring them all in.

A mid-sized company runs hundreds of smaller systems: a phone list, a booking tool, a supplier register that someone built years ago. Many are no longer maintained, and nobody watches them. Sentinel imports them, converts them to run in its own containers, bakes in AI where it helps, and from then on keeps them updated, monitored and secure, just like the systems it builds.



**What you need: the original source code.** Sentinel imports and converts each system from its source code, so you must have access to the original code for every system you bring in.

01

## Import

Bring in an existing system with its original source code and its data. Nothing has to be rebuilt from scratch.

02

## Convert

Sentinel converts it to run in a standard Docker container, with backups and a snapshot of every version.

03

## Monitor

Live traffic, errors, users, AI use and firewall rules, on the same screens as every other system.

04

## Maintain

Updates, new features and AutoFix from the Builder. The original developer is no longer needed.



NEW LIFE, NEW ABILITIES

## Bake AI into systems that never had it.

Once a system is in Sentinel, you can add AI exactly where it helps the people who use it: search that understands questions, drafts, summaries or checks, without rebuilding the system. Every AI feature runs through Sentinel's AI gateway, so it is logged, scored for risk and kept within your cost limits.


✦ **Smart search**  
Staff phone list

✦ **Drafted replies**  
Holiday requests

✦ **Risk flags**  
Supplier register

✦ **Summaries**  
Quality reports

# Every old system goes through Sentinel and comes out healthy.



The five systems shown are examples. Each one keeps its own container, backups and version history, gains AI features where they help, and appears on the same Live, Firewall and AI screens as everything else in Sentinel.

# Know exactly how AI is used. Nothing reaches a model unseen.

Every AI request from every attached system passes through Sentinel's gateway. It is logged in full, scored for risk before it is sent, and linked to the person, the page and the model — with its cost counted in real time.

100%

of AI requests logged — prompt, answer, tokens and cost

0–100


security score on every question before it is sent

60+

scores are stopped before a model ever sees them

1 click

to an AI Use Statement for auditors, as a PDF



The prompt, the complete answer, tokens, cost and who asked. Search it, filter it, export it to Excel.

An automatic reviewer rates each question from safe to dangerous — prompt injection, secrets, personal data, off-topic use.

With the pre-send check on, requests that score 60 or more never reach a model. Your own review always has the last word.

See every page that uses AI, which user groups can open it and which models answer there.

Spend per system, per model, per page and per person as it happens, with a maximum cost per use wherever you want one.

Open-weight models run on your own hardware at no usage charge. Cloud models are used only when needed.

AI Use - pages × user groups × models

pages × user groups × models

Rows are pages that use AI. Columns show which user groups can open each page, how much each model is used there, and the totals with the maximum cost per use. Click any cell for details.

Requests

Spent

Users

Enabled for group

No restriction

Hidden

Sentinel admins

Model use (low → high)

LOCAL: own server

CL-PROD: via OpenRouter

AI pages

7 pages - 30 days

Compliance working...  
2 members

Drug & Alcohol Adm...  
2 members

Feature Request  
1 member

IT Security  
1 member

Line Mechanics  
1 member

Rank/Kahuna  
2 members

Sentinel Team  
2 members

TechOps Management  
1 member

No group  
2 members

CL-PROD  
gov-6-astro  
Ops/ml

CL-DEV  
cloud-01  
Astro

Knowledge Library

Semantic search (just embedding)...

2 used

2 used

<

# AI Use · models & security review

**MODEL REGISTER**

## Models in use

| Model Name               | Provider | Usage Details                                                          | Status  |
|--------------------------|----------|------------------------------------------------------------------------|---------|
| gpt-4o-extra             | Cloud    | OpenAI - 50k requests - 2 people - \$10.00 / \$30.00 per M tokens      | \$42.77 |
| claude-opus-4.5          | Cloud    | Anthropic - 138 requests - 1 people - \$4.00 / \$30.00 per M tokens    | \$32.40 |
| glm-3.3-flash            | Cloud    | Z.ai(Zhipu AI) - 89 requests - 2 people - \$0.00 / \$0.14 per M tokens | \$0.44  |
| gpt-4o-turbo             | Cloud    | OpenAI - 5 requests - 1 people - \$0.10 / \$0.50 per M tokens          | \$0.01  |
| gemma-4-31b-it           | Cloud    | Google - 8 requests - 2 people - \$0.09 / \$0.34 per M tokens          | \$0.01  |
| llm-m3-mix-8bit          | Local    | BAAI (open weight) - 140 requests - 2 people - no usage charge         | \$0.00  |
| llm-expert               | Local    | Open-weight model - 15 requests - 2 people - no usage charge           | \$0.00  |
| Owen3.8-27B-8bit         | Local    | KoboldCPP (open weight) - 11 requests - 1 people - no usage charge     | \$0.00  |
| gemma-4-31B-E-MIX-8bit   | Local    | Google (open weight) - 5 requests - 2 people - no usage charge         | \$0.00  |
| GEMMA-3.3-Flash-MIX-8bit |          |                                                                        | \$0.00  |

**OVERSIGHT**

## Security review & controls

5 Stars 68
Low risk 9
Medium-risk 1
High risk 5
Dangerous 33

110 of 250 application requests scored automatically - 25 reviewed by an administrator - 32 stopped by Sentinel - 1 errors.

- ✓ **Control AI gateway**  
Every request tagged with user, page, model and cost
- ✓ **Automatic security review**  
Scored by a z-writer-5.5 flash
- ✓ **Pre-send check**  
Stops requests scoring 60+ below a model seen threshold
- ✓ **Held delivery**  
Answers delivered directly
- ✓ **Maximum cost per use**  
0 of 1 cloud-using pages limited
- ✓ **Access by user group**  
Single sign-on and group page access

## — COST CONTROL

# See the cost *as it happens*. Set the limits yourself.

Sentinel prices every AI request the moment it is made, so you always know what AI costs you: today, per system, per person, per page and per model. Not as a surprise at the end of the month.

- ✓ Live spend per system, per model, per page and per person
- ✓ A maximum cost per use on any AI page, set in dollars
- ✓ Answers are shortened to stay under the limit, or the request is stopped before it is sent
- ✓ Local models run at no usage charge and are never limited

● LIVE AI SPEND · TODAY

\$149.61

Maximum cost per use

\$1.00

Sales · Offer writer

Jennifer W. · Cloud model

\$0.1842

HR · Policy search

Michael T. · Local model

Free

Finance · Invoice check

Robert J. · Cloud model

\$0.9934

Answer shortened to stay under \$1.00

Legal · Contract review

James M. · Cloud model

\$0.0000

Stopped before sending: above the \$1.00 limit

Support · Reply draft


Ashley R. · Cloud model

\$0.0716

Illustration · names and values are examples

# Your own AI, plugged into Sentinel. Still under your rules.

Developers keep Claude Code or Codex on their own computer and connect it with one pasted text. Agents that run inside your network can be locked into an NVIDIA OpenShell sandbox. Every AI call — hosted or on-prem — passes through Sentinel, and Sentinel itself stays out of reach.



## Your own AI, connected in a minute

ENTERPRISE

Claude Code or Codex talks to Sentinel over MCP with a personal, revocable token. It reads the code of the systems the person can access, asks Sentinel, pulls container logs (secrets masked) and files. Read-only, every call logged — Sentinel itself is off-limits.


## NVIDIA OpenShell, an agent in a cage

ENTERPRISE · OPTION

NVIDIA's open-source sandbox runtime for AI agents. Each sandbox runs in its own micro-VM; a policy decides which networks and files the agent can reach, and every policy is proven against your Sentinel boundary before it is saved.

# See everything *as it happens.*

Sentinel's Live page follows every request through the gateway in real time: where it came from, which system it went to, how long it took, and how risky the AI part was. Dot speed follows response time — you can literally watch your systems breathe.



# Systems that keep getting better – and fix themselves when they break.



## ONE-CLICK DEVELOPMENT

From idea to update, again and again.

- 1 A user requests a feature**  
From inside the system they use
- 2 An administrator starts AutoBuild**  
One click, after a quick read
- 3 Sentinel builds the feature**  
Code, checks and deployment
- 4 The application is updated**  
With a snapshot to roll back to



## REAL-TIME ERROR CORRECTION

From bug to fix, automatically.

- 1 A bug is detected**  
In any attached system
- 2 Sentinel detects and diagnoses it**  
Logs, code and recent changes
- 3 Sentinel runs AutoFix**  
A focused, reviewable repair
- 4 The application works again**  
Recorded in the change history

# Your systems. Your source code. Your hardware, or your cloud.

RUNS WHEREVER YOU CHOOSE

● Your own servers

● AWS

● Microsoft Azure

● Any other hosted platform



## Docker deployment

Every system runs in standard Docker containers, on your own servers or in AWS, Azure or any other hosted platform. Backed up, restorable and recoverable on its own.



## Your plugins

You own every plugin you create, including its source code, and can reuse it across systems.



## Your systems

You own every system you develop. Nothing is locked inside someone else's platform.



## No licensing fees

No per-seat or per-system licences for the plugins and systems you create.

## AI security by design

Every AI request from every system passes through Sentinel, is logged in full, risk-scored and reaches a model only when it is safe.

## User-aware, context-aware AI

Sentinel knows who is asking, from which system and module — and with what data — so answers and access follow the person.

## One-click development

Describe a change. Sentinel's AI builds it in the right module, with cost, history and rollback on every step.

## One-click deployment

Accepted changes go live in the system's container automatically, with health checks and instant recovery.

## AI value alignment

Your company's values, context and rules shape how every request is judged — from safe to dangerous.

## EDITIONS

# Two editions. Start with Essentials, grow into Enterprise.

Both editions share the same AI gateway, Live monitoring and firewall. Enterprise adds scale, self-healing, compliance tools and high availability.

| FEATURE                                                                                                   | Sentinel Essentials<br>For a team or a department | Sentinel Enterprise<br>For the whole organisation |
|-----------------------------------------------------------------------------------------------------------|---------------------------------------------------|---------------------------------------------------|
| <b>SYSTEMS</b>                                                                                            |                                                   |                                                   |
| Systems (Docker containers)                                                                               | Up to 3                                           | Unlimited                                         |
| Existing systems imported and converted                                                                   | 1                                                 | Unlimited                                         |
| Administrators                                                                                            | 2                                                 | Unlimited                                         |
| <b>BUILDING</b>                                                                                           |                                                   |                                                   |
| Builder and one-click development                                                                         | ✓                                                 | ✓                                                 |
| Builds                                                                                                    | One at a time, monthly quota of change packages   | Several at once, no quota                         |
| AutoFix — self-healing error correction                                                                   | —                                                 | ✓                                                 |
| Bring your own AI — your local Claude Code or Codex, plugged into Sentinel                                | —                                                 | ✓                                                 |
| <b>AI AND SECURITY</b>                                                                                    |                                                   |                                                   |
| AI gateway: every request logged, risk-scored and cost-limited                                            | ✓                                                 | ✓                                                 |
| Live monitoring and firewall                                                                              | ✓                                                 | ✓                                                 |
| AI value alignment with your own company policies                                                         | —                                                 | ✓                                                 |
| Single sign-on                                                                                            | —                                                 | ✓                                                 |
| OpenShell integrated — NVIDIA's sandbox runtime for AI agents                                             | —                                                 | ✓                                                 |
| Daily brief — your code checked against vulnerability databases, plus the day's news, as a PDF newsletter | —                                                 | ✓                                                 |
| <b>HISTORY AND COMPLIANCE</b>                                                                             |                                                   |                                                   |
| AI request logs, Live history, snapshots and backups                                                      | 30 days                                           | Unlimited                                         |
| One-click AI Use Statement PDF for auditors                                                               | —                                                 | ✓                                                 |
| Lockdown with database dumps and a full incident record                                                   | —                                                 | ✓                                                 |
| <b>HOSTING AND SUPPORT</b>                                                                                |                                                   |                                                   |
| Hosting                                                                                                   | A single host                                     | Several hosts or clouds, high availability        |
| Support                                                                                                   | Standard support                                  | SLA                                               |
| Licensing fees for the systems and plugins you build                                                      | None                                              | None                                              |

# Think *It.* Build *It.* Protect *It.*

One control centre for your business applications and the AI inside them — built with you, running on your hardware or in your cloud, watched over every second.

[sentinel.loglevel.com](https://sentinel.loglevel.com)

## LOGLEVEL USA

4770 Duke Drive  
Mason, OH 45040  
USA

## LOGLEVEL EUROPE

Robert Jacobsens Vej 68  
2300 Copenhagen  
Denmark

Screens in this document come from the running system; names and addresses in the screens have been replaced.

© 2026 Loglevel