ai-ta — Autonomous Infrastructure Triage Agent

AI investigates. Humans decide. Infrastructure learns.

THE PROBLEM

Platform ops runs 8-5. Infrastructure runs 24/7. The 15-hour gap means overnight incidents go undetected until morning — often after they've cascaded into application outages that wake up the wrong team. **NIS2 Article 20 makes the board personally liable** for incident response capability you don't have outside business hours.

HOW IT WORKS

Alert (Grafana / Any webhook) → **Investigate** (SSH, metrics, logs, Docker) → **Resolve or Escalate** (destructive = human approval) → **Learn** (knowledge accumulates) → **Repeat** (24/7)

	Without ai-ta	With ai-ta
3am disk alert	Engineer wakes up, SSHes in, investigates	Agent investigates, requests approval, cleans up. Morning summary.
Recurring issue	Reinvestigated from scratch every time	Auto-resolved in milliseconds from learned knowledge — no LLM cost
NIS2 audit	Scramble to reconstruct timeline	Timestamped trail: detection, investigation, action, approval
New team member	Weeks of shadowing for tribal knowledge	Queries accumulated knowledge from any AI tool

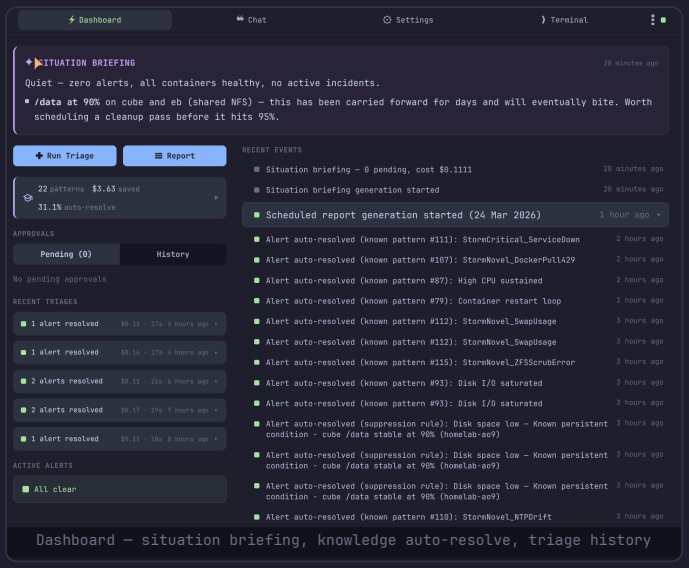

Dashboard — situation briefing, knowledge auto-resolve, triage history

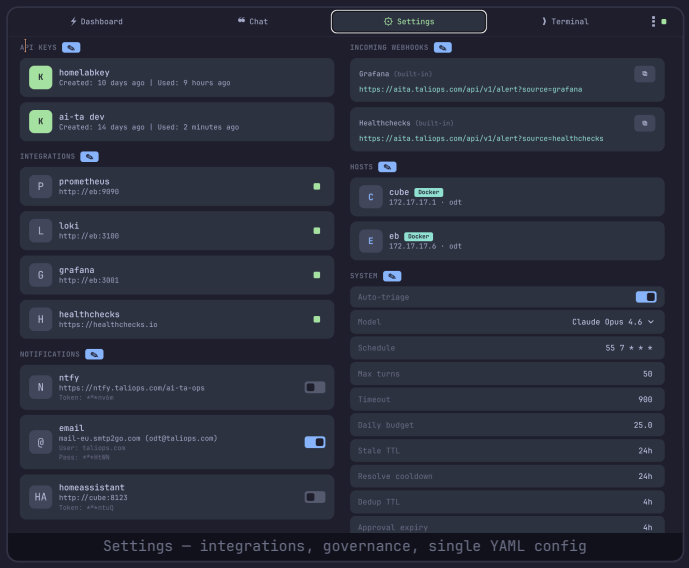

Settings — integrations, governance, single YAML config

BY THE NUMBERS

200+	75	~60s	$40	99.9%
TRIAGES LAST 30D fully autonomous	AUTO-RESOLVED zero human effort	AVG INVESTIGATION vs 30-45 min manual	MONTHLY LLM COST all-in, 200+ triages	UPTIME (YTD) production SLA

NIS2 & GDPR COMPLIANCE

Incident handling (Art. 21.2b)	Full lifecycle: detect, investigate, respond, document
Continuous monitoring (Art. 21.2a)	24/7 autonomous triage with trend detection
Incident reporting (Art. 23)	24h early warning, 72h detail, monthly summary
Supply chain oversight (Art. 21.2d)	Independent monitoring of MSP-managed infra
Board accountability (Art. 20)	AI policy version-controlled, commit-stamped
GDPR data minimization	Configurable retention + automated purge per data type

INTEGRATION

Works with what you have.

MONITOR: Prometheus · Grafana · Loki · Healthchecks.io

INFRA: Docker · Kubernetes · OpenShift

NOTIFY: Slack · Teams · ServiceNow · ntfy · Email

AI: MCP protocol · Code mode · MCP elicitation

AUTH: OIDC · Bearer tokens · Approval gates

Oliver Tali — oliver@taliops.com

AI investigates. Humans decide. Infrastructure learns.

ai-ta — The Future of Operations

Software stops being the constraint. Humans steer. AI executes. Contracts enforce.

WHERE OPERATIONS IS GOING

The operator's job becomes **steering processes, owning accountability, and managing risk** — not navigating dashboards. Three shifts are converging:

Machine-to-machine by default
Monitoring talks to triage, triage talks to remediation, remediation talks to approval gates. Humans intervene at decision points, not execution points. Contracts and schemas enforce determinism between autonomous systems.

Frontends become contextual projections
What matters isn't how the UI looks — it's that the right information reaches the right person at the right moment. A morning email, a mobile approval, an MCP query from another agent — these are all valid interfaces.

Governance becomes the product
When AI acts autonomously, the value is the audit trail, the approval gates, the policy enforcement, the explainability. The strictest compliance environments aren't obstacles — they're the reason it must be built this way.

BUILT FOR THIS FUTURE

Contract-driven	MCP protocol: any AI agent — your team's, your MSP's, your toolchain's — queries and acts on infrastructure knowledge through structured contracts. Not screen scrapes. Not API wrappers.
Process-native	The same triage result renders as an HTML email, an approval webhook, an MCP response, or a terminal session — whatever the process demands.
Self-learning	Every triage cycle feeds the next. Known patterns auto-resolved in milliseconds — no LLM call, no cost. Stale context drops out automatically. Situation briefing proactively tells the operator what matters. Knowledge survives team turnover.
Compliance-first	Every action timestamped. Every decision traceable. Every destructive command gated. Every AI execution governed by version-controlled organizational policy.
Cross-industry	Declarative, pluggable governance. A maritime operator, a hospital, and a fintech run the same agent with different policy files. NIS2, DORA, sector-specific — same architecture.

SCOPE & BOUNDARIES

ai-ta **observes across every layer** and **contains with human approval**. It does not make irreversible decisions.

What ai-ta does	
Observe	L7 edge (Cloudflare WAF), L3/L4 perimeter (firewall), L2 network, host & container — all layers, no limits
Contain	Block at edge, isolate at network, restart services — every action human-approved, auto-expiring
Learn	Accumulate knowledge, detect patterns, skip known-good investigations, flag trends
Document	Full audit trail from detection to resolution — timestamped, traceable, exportable

What ai-ta never touches	
Data	No database operations, backup restores, or storage modifications
Identity	No credential rotation, access policy, or permission changes
DNS	Slow to reverse, high blast radius — stays with the human
Hypervisor	Bare metal and VM lifecycle are last-resort human decisions

INSTITUTIONAL MEMORY

- Every resolved alert becomes a learned pattern
- Known patterns auto-resolve in milliseconds — zero LLM cost
- Learning curve visualizes novel vs known ratio over time
- Knowledge survives team turnover — it's in the system, not in people's heads

Oliver Tali oliver@taliops.com

AI investigates. Humans decide. Infrastructure learns.