



One AI-native platform for every incident.

Monitoring, on-call, logs, synthetic, runbooks, postmortem. Unified, with an AI that works your tools.

PRE-SEED · RAISING \$2M · LEAD INDICATED BY GACSYM VENTURES (MATCHES AT \$1M)

founders@novaaiops.com · novaaiops.com

THE PROBLEM

It's 03:14 AM. The pager fires.

A senior SRE is paged for a latency spike on a tier-1 service. They flip between Datadog (metrics), Splunk (logs), PagerDuty (on-call), Confluence (runbooks), Slack (comms), and Jira (tickets). Six disconnected tools, no shared context. They copy-paste dashboard screenshots, hand-write a postmortem, and by the time they diagnose it, the error budget is half burned.

This exact sequence happens every night, at every enterprise with a production service. There are 700,000+ SREs worldwide doing this by hand. Industry benchmark for a 1% latency regression: \$260K per hour in lost revenue for a Fortune 500 service.

4h

MEDIAN MTTR, 2026

6+

TOOLS PER INCIDENT

\$260K

LOST PER HOUR OF DOWNTIME

53%

SRES REPORT BURNOUT

Three inflections converged in 2026.

01

Tool use got reliable.

Claude Sonnet + GPT-5 tool calling hit production quality in late 2025. Agents can now read logs, run remediations, and rollback deploys as reliably as humans, for 1/50th the cost.

02

SRE labor can't scale.

SRE headcount demand is up 34% YoY, supply is flat. Average loaded cost: \$248K. Every team has the same on-call rota, no team can hire its way out.

03

Dashboards, not workers.

Datadog, PagerDuty, Rootly all surface the problem. None resolve it. The category opened up when incumbents added 'AI Assist' chatbots instead of actual agents.

Six ops tools. One AI-native platform.

BEFORE · TODAY'S SRE STACK

~\$600K / year · six vendors

Datadog	Metrics / APM	\$180K
Splunk / Loki	Logs	\$120K
PagerDuty	On-call + paging	\$60K
Rootly / Incident.io	Runbooks + postmortem	\$80K
Checkly / Pingdom	Synthetic checks	\$30K
Confluence + Jira	Docs + tickets	\$50K

AFTER · NOVA AI OPS

\$14K - \$180K / year · one vendor

- ✓ Real-time metrics + golden signals
- ✓ Log explorer + distributed tracing
- ✓ Active incidents + on-call rotation
- ✓ AI runbooks + postmortem generator
- ✓ Synthetic monitors + service map
- ✓ Ask Nova: AI that works your tools

33

UNIFIED PAGES

100

AI AGENTS

8

SAFETY GATES

1

DATA MODEL

Live at app.novaiops.com · Blue-green zero-downtime deploy · Multi-tenant · SOC-2 ready architecture

Replace the stack, don't fight the leader.

Capability	Datadog + PD + Splunk + Rootly	Nova AI Ops
Unified data model	5 separate stores	One schema, one timeline
AI assistant	Chatbots (summarise only)	Page-aware, takes actions
Incident > runbook > postmortem	3 tools, 3 logins	One workflow
Synthetic + real-user + logs	Separate vendors	One query surface
Safety layer for autonomy	None shipped	8 gates, production-enforced
Median annual spend (mid-market)	\$450K - \$620K	\$14K - \$180K

A VP-Eng doesn't need a better Datadog. They need their six renewals to stop arriving.

Three moats, not one.

01 CONSOLIDATION SWITCHING COST

Once a team unifies monitoring, on-call, logs, runbooks, and postmortem on one data model, there's no unwinding it. Every incident written into our timeline is one a competitor would have to re-import to steal the account. Datadog + PagerDuty + Splunk can't consolidate themselves. They each depend on their separate billing relationships.

02 INCIDENT-RESOLUTION DATA GRAVITY

Every auto-resolved incident becomes a training signal scoped to that tenant's services, runbooks, and people. After 90 days the AI is materially better at that customer's environment than on day one. Competitors who ship an AI chatbot in 2026 start at day zero against a 12-month data advantage.

03 SAFETY LAYER THAT CUSTOMERS TRUST

Eight production-enforced gates (kill-switch, prompt-injection defense, cost breaker, SLO gate, tenant isolation, ground-truth verifier at T+5m/1h/24h, consensus arbiter, digital-twin simulation). This is what lets a VP-Eng flip autonomous mode on for a tier-2 service. The thing none of the 'AI assist' chatbots can earn.

The founders lived this pain for years at JPMorgan, US Navy, and trial-firm operations. That's the zero-th moat.

WHERE WE ARE

Product is live. Real users. 805K actions logged.

\$1M

LEAD INDICATED · GACSYM VENTURES

Verbal, matches at \$1M co-invested

Mar 20

2026 PRODUCT LAUNCH

33 days live at app.novaiops.com

15

ORGANIC SIGNUPS, UNPAID

11 non-founder users, 16 workspaces

805K+

ACTIONS LOGGED IN 30 DAYS

14.5K actions in the last 7 days alone

4

UNSOLICITED FEEDBACK ITEMS

From real users, zero prompting

~0.8

ORGANIC SIGNUPS PER DAY

No paid ads, no outbound yet

ACCEPTED INTO STARTUP PROGRAMS

Confluent

Databricks

Google for Startups

Nvidia Inception

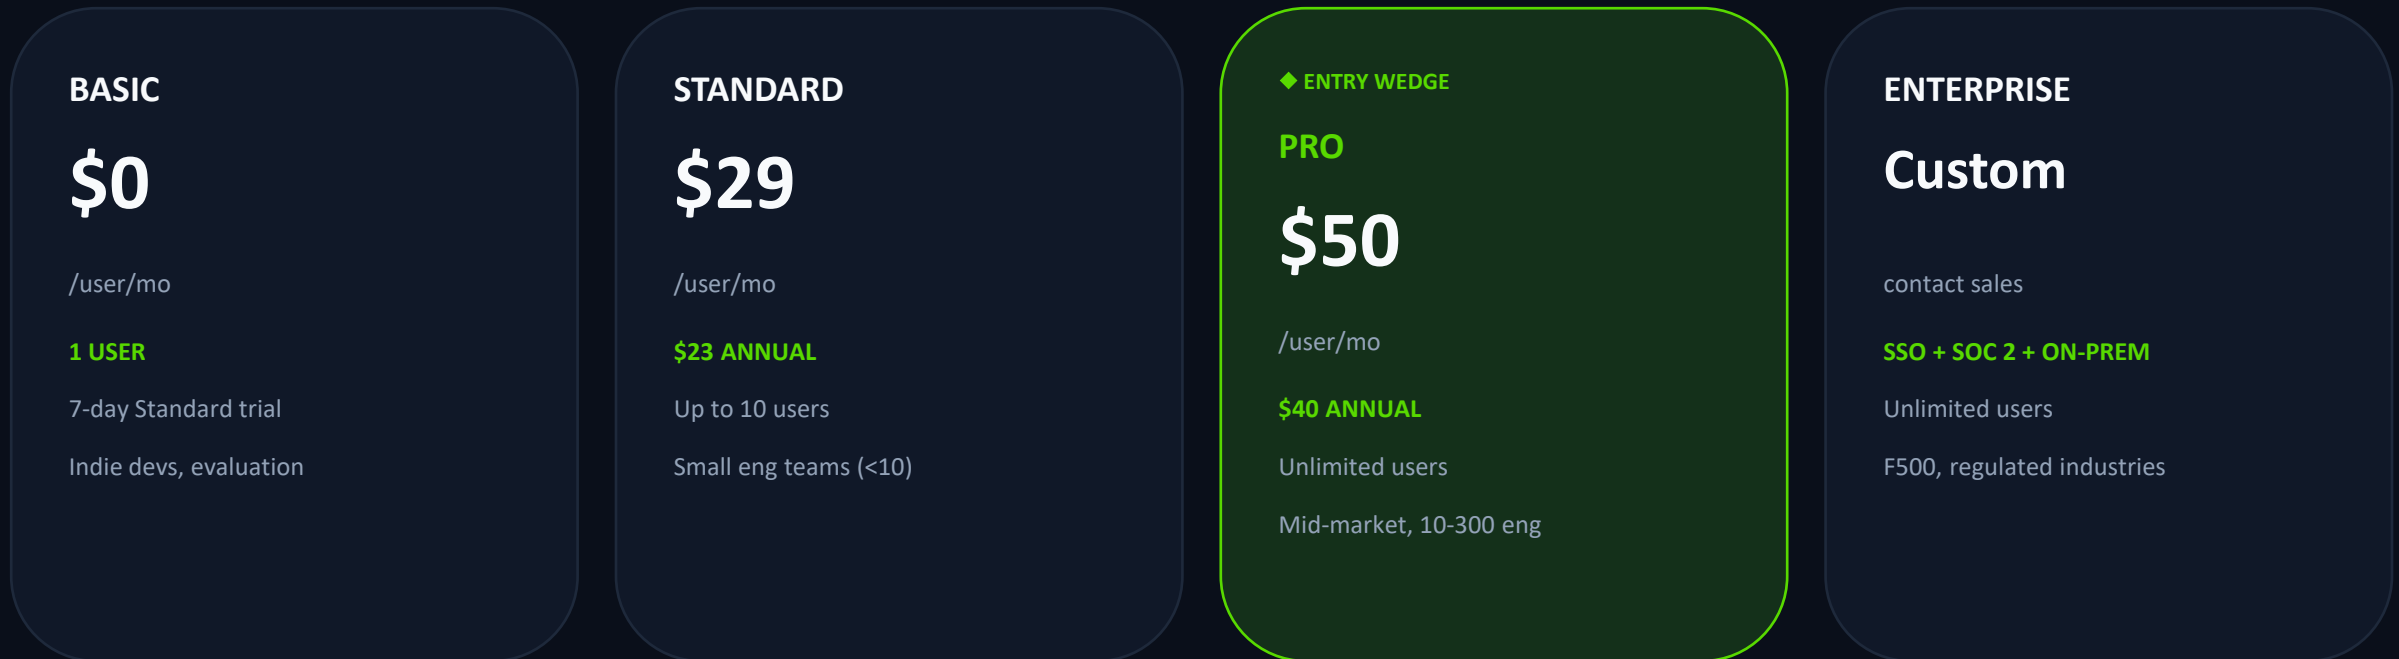
MongoDB for Startups

Redis for Startups

AWS Activate

Pre-revenue but not pre-usage. 11 organic non-founder signups in 14 days with zero paid acquisition. This round funds the CTO + 3 founding engineers + conversion to first 5 paid design partners.

Per engineer. One bill replaces six vendor contracts.



ACV math (real numbers, live pricing):

- Design-partner pilot (30-eng team) on Pro annual: $30 \times \$40 \times 12 = \$14.4K$ ACV. Mid-market scale (80 eng): $\$38.4K$ ACV.
- First Enterprise logo (200+ eng, SSO + SOC 2): $\$100K - \$180K$ ACV + on-prem / DR premium.

Year-1 target: 15 paid teams avg 25 seats → \$180K ARR. Year-2: 2 Enterprise + 40 Pro teams → ~\$780K ARR.

\$6.8B spent on SRE labor we can offset.

700,000

SREs worldwide (StackOverflow survey, LinkedIn)

x \$248,000

Average loaded cost per SRE (levels.fyi)

= \$174B

Global SRE labor spend (annual)

x 4%

Conservative automation displacement (agent-resolved SEV-3/4 toil)

= \$6.8B

Serviceable obtainable market for Nova AI Ops

Bottom-up only. We exclude APM-category spend (\$14B Datadog + \$9B NewRelic + others = \$45B+) to avoid inflating TAM with observability tools we don't replace.

GO-TO-MARKET

Wedge: the mid-market CFO doing vendor math.

WEDGE · MONTHS 1-6

10 design-partner eng teams (100-300 engineers, 50-150 services). Trigger event: they're staring at five ops-vendor renewals totalling \$300-500K and the CFO has asked 'can this be one contract?'. We land on Pro annual with 20-30 seats (\$10K-\$14K ACV), replacing 2-3 of the five. Bridge the rest in phase 2.

EXPAND · MONTHS 7-12

Expand each paid account from 20-30 seats to 60-120 seats as trust grows. Displace remaining vendors (logs vs Splunk, synthetic vs Checkly, postmortem vs Rootly). Land-and-expand compounds because each new surface adds data gravity. Target: 3-4x seat growth per account by month 12.

PLATFORM · MONTHS 13-18

First Enterprise logo (\$100K-\$180K ACV, SSO + SOC 2). Distribution: (1) founder-led LinkedIn + technical blog, (2) open-source the 8 safety modules as a reference impl for Hacker News + agent-safety research community, (3) AWS Marketplace listing (SRE teams have pre-approved spend).

TEAM

The founders who lived this pain. Plus who we're hiring next.



Dr. Samson Tanimawo

FOUNDER & CEO

SRE with 10+ years at JPMorgan Chase and the US Navy. PhD, MSc, MBA. Built the full product himself: 33 pages, 100 agents, 8 safety modules.



Lashae Tanimawo

CO-FOUNDER & CMO

Product marketing leader. Positioning, messaging, and GTM for technical AI products. Creator of Pain to Profit.



Jennifer Broxson

DIRECTOR, AI STRATEGY & BD

Computer scientist (cybersecurity, R1 ABET). Ships end-to-end ML pipelines. Past: 11 yrs senior paralegal. Turns complex systems into revenue-ready positioning.

HIRING WITH THIS ROUND · FIRST 90 DAYS

CTO / Co-Founder

Senior platform / SRE leader. Ex-Datadog, ex-Linear, ex-Stripe preferred.

Founding AI Engineer

Agent-tool-use + safety systems. Ship the next 50 agents.

Founding SRE / DevRel

First 5 design-partner rollouts + the SRE community channel.

Backed by the platforms we're built on.

Accepted into the startup programs of every major data + AI + cloud stack vendor Nova integrates with.

Confluent

For Startups — streaming + Kafka credits

Databricks

Ventures / Startup Program — lakehouse + ML

Google for Startups

Cloud credits + Gemini API credits

Nvidia Inception

GPU credits + Deep Learning Institute access

MongoDB for Startups

Atlas credits + technical architecture review

Redis for Startups

Cloud credits + enterprise features

AWS Activate

Credits + marketplace listing support

THE ASK

\$2M

Pre-Seed round (open)

ROUND SIZE

\$2M pre-seed

LEAD INDICATED

\$1M · Gacsym Ventures (verbal)

TRIGGER

Matches at \$1M raised from co-investors

Close target: Q3 2026

Open to lead-replacement offers and co-lead structures. SAFE or priced round, founder-friendly terms welcome.

USE OF FUNDS

50%

Engineering & AI Talent

4-6 hires: Engineering, SRE, ML/AI

20%

Infrastructure & Compute

Cloud, GPU compute, production reliability

15%

Go-To-Market

Sales team, enterprise pilots, devrel

15%

Operations & Compliance

SOC 2 Type II, ISO 27001, legal

KEY MILESTONES

- ▶ 5 - 10 enterprise design partners
- ▶ Scale team to 10+ engineers
- ▶ 100K+ events/sec in production
- ▶ Launch mobile on-call app
- ▶ SOC 2 Type II certification
- ▶ Custom AI persona training

founders@novaiops.com · novaiops.com