



Senior engineers who build like **owners**.

We build and run our own SaaS — then put the same AI-native team on your product, as a dedicated team or a full build.

100%

SENIOR-ONLY STAFFING

1 wk

WORKING SOFTWARE CADENCE

US • UK

PRIMARY MARKETS

Who We Are

Nugen is a product engineering studio. We build and run our own SaaS products — then put that same senior, AI-native team on client products we believe in.

That makes us different from typical outsourcing: our engineers have shipped and operated real software, so they bring product judgment — activation, reliability, cost per user — not just clean code.

Work with us as a full project build, an embedded dedicated team, or targeted staff augmentation. Either way you get product-grade software, shipped weekly, owned end to end.

CATEGORY

Product engineering studio — not a generic IT outsourcer

PRIMARY MARKETS

United States & United Kingdom

BUYERS

Founders, CTOs, VPs of Engineering, product owners

ENGAGEMENT MODES

Product build · Dedicated team · Staff augmentation

WHY WE EXIST

The Gap in Outsourcing

Most software is outsourced to teams that have never shipped and run a product of their own. The result is code that technically works — and a product that commercially doesn't.

WE ARE NOT

- × A generic IT services / staffing company
- × A pool of junior contractors billed by the hour
- × A firm that added an "AI" page to look current
- × A vendor that has never shipped its own product
- × The cheapest bid in the room

WE ARE

- A product engineering studio
- Senior-only, outcome-owned teams
- AI-native as a default
- Operators of our own SaaS
- Premium relative to offshore — and worth it

THE ONE-LINE DIFFERENCE

"We can't hide behind the ticket — because we run the same kind of product we're building for you."

REASONS TO BELIEVE

Why Nugen

01

We run our own SaaS

Product judgment from operating real software — not just code delivered.

02

Senior-only teams

Engineers who have shipped and run real products. No junior bait-and-switch.

03

AI-native by default

AI features that reach production and stay reliable — not demos that die in a pilot.

04

Outcome-owned

Measured on the product shipped and its impact — not hours logged.

05

Working software weekly

Small, visible increments — and an honest status you can trust.

Product-grade engineering,
owned end to end.

— THE NUGEN PROMISE

IDEAL CUSTOMERS

Who We Serve

We don't serve everyone. We serve two profiles precisely.

PRIMARY ICP

US · UK

Venture-backed SaaS teams

Seed – Series B · 10–150 people · B2B SaaS, fintech, healthtech, AI/ML, dev tools

- Can't hire senior engineers fast enough
- Roadmap slipping against a launch or renewal deadline
- Needs AI features but lacks in-house AI depth
- Burning runway on slow delivery

Buyers: Founder / CEO · CTO · VP Engineering

SECONDARY ICP

US · UK

Mid-market product organizations

Post-Series B – enterprise · 150–2,000 people · software, finserv, healthcare, logistics

- Legacy platform slowing the business down
- Internal teams over capacity
- AI and automation initiatives stuck in pilot
- Failed prior vendor; new CTO cleaning house

Buyers: VP / Director Engineering · Head of Data / AI

— WHAT WE BUILD

Five services. **One senior team.**

We keep the surface small on purpose. Every service is delivered by the same senior pool — and defined as a business outcome, never a technology.

01	AI Engineering	Production RAG, agents, evals — AI that stays reliable past the demo.
02	Product Engineering	SaaS platforms, web apps, and MVPs — built to production quality.
03	Dedicated Engineering Teams	An embedded senior squad that owns a slice of your roadmap.
04	Platform Modernization	Slow, legacy systems rebuilt fast and maintainable — no big-bang rewrite.
05	Cloud Engineering	Reliable, observable, cost-efficient software in production.

AI Engineering

PROBLEM

Leadership wants AI in the product, but internal teams lack the depth to get it past a demo — or to keep it reliable and control its cost.

SOLUTION

We design and ship production AI — RAG systems, LLM-powered features, AI agents, and document processing — built to stay accurate, safe, and cost-efficient.

Proven in production: LangChain-gated LLM feedback at Mastery Coding, incident summarization at Intterra, RAG accuracy at Mosaicx.

What you get

- Working AI features in production
- Evaluation harness & guardrails
- Monitoring, cost controls & documentation

TIMELINE	First production feature in 4–12 weeks
PRICING	Project from \$30k · AI squad from \$25k/mo
WHO BUYS	CTOs, Heads of AI / Data, product leaders

AI That Survives Production

Most "AI features" die between the demo and the second week of usage. We build the system underneath — retrieval, evaluation, guardrails, cost — so the feature stays useful when your CFO reads the bill.

01 • INGRESS

Retrieval & Grounding

Hybrid search, chunking strategy, freshness signals, tenancy-aware embeddings. Answers cite sources — always.

02 • CORE

Agents & Orchestration

Task decomposition with typed tools, deterministic fallbacks, and clean state. Not a jumble of chained prompts.

03 • EGRESS

Guardrails & UX

Structured outputs, refusal handling, latency budgets, and graceful degrade paths written into the interface.

04 • FEEDBACK

Evals & Regression

Golden sets, offline evals in CI, prompt versioning, and human review loops so model drift can't ship silently.

05 • COST

Cost & Routing

Model tiers, caching layers, and per-tenant budgets. We ship with a per-request unit-cost dashboard from week one.

06 • SAFETY

Trust & Compliance

PII redaction, audit trails, SOC 2-aligned logging, and red-team playbooks — not a security add-on later.

Product Engineering

PROBLEM

A SaaS platform, web app, or MVP needs to be built to production quality — faster than in-house hiring allows.

SOLUTION

A senior squad builds and ships your product in small increments — from architecture through launch, with product judgment baked in.

Proven in production: Optum Store's prescription commerce, Qudos' real-time mobility platform, Guide to Iceland's booking surfaces.

What you get

- Shipped SaaS / web / mobile product
- Clean architecture, CI/CD, tests
- Docs & a handover-ready codebase

TIMELINE **MVP in 6–12 weeks · platforms 3–6+ months**

PRICING **Fixed build from \$30k · team from \$25k/mo**

WHO BUYS **Founders & CTOs, seed – Series B**

Dedicated Engineering Teams

PROBLEM

You need senior engineering capacity now — and can't wait out a 3–6 month hiring cycle.

SOLUTION

A cohesive senior team embeds with you and owns a slice of your roadmap month over month — with continuity and shared context, not a rotating cast of contractors.

Proven in production: multi-year embedded teams at Apploi (2023–present) and Intterra (2024–present).

What you get

- Embedded senior team, sprint cadence
- Weekly working software
- Shared tooling & clear ownership

TIMELINE

Kickoff within 1–2 weeks • ongoing

PRICING

\$25k – \$60k / month

WHO BUYS

VPs of Engineering, CTOs scaling delivery

Platform Modernization

PROBLEM

A legacy or slow platform is holding the business back — hard to change, expensive to run, risky to touch.

SOLUTION

We rebuild slow systems into fast, maintainable, well-tested platforms — incrementally, without a risky big-bang rewrite.

Proven in production: Apploi's class-to-hooks frontend migration (~40% boilerplate cut), Dapt Tech's payroll re-architecture after a silent-drop incident.

What you get

- Modernized architecture & migration plan
- Improved performance & reliability
- Test coverage & documentation

TIMELINE **Assessment 2–4 weeks · program 3–9+ months**

PRICING **Program-based · dedicated team from \$25k/mo**

WHO BUYS **Mid-market eng leaders, new CTOs**

Cloud Engineering

PROBLEM

Software is live but fragile, opaque, or expensive to run in production.

SOLUTION

We make systems reliable, observable, and cost-efficient — CI/CD, infrastructure-as-code, monitoring, and cloud cost optimization.

Proven in production: Kubernetes-orchestrated sandboxes at Mastery Coding, Kafka burst-load pipelines at Intterra, AWS scale at Qudos.

What you get

- Automated pipelines & infrastructure-as-code
- Observability & alerting
- Documented, cost-optimized cloud setup

TIMELINE Foundation in 4–10 weeks · ongoing as needed

PRICING Project from \$20k · or within a dedicated team

WHO BUYS CTOs, platform / DevOps leads

Engagement Models

We price on seniority and outcomes — never on being the cheapest bid.

MODEL 01

Product Build

We take full ownership of a defined product outcome and timeline — architecture through launch.

TEAM3–6 · full-stack

DURATION6–16 weeks

\$30k – \$150k

Best when scope is definable and you want a fixed outcome.

MOST COMMON

MODEL 02

Dedicated Team

Senior engineers embed with you and own a slice of the roadmap — month over month, with continuity.

TEAM4–10 · cross-functional

DURATION6–24 months

\$25k – \$60k / mo

Best when you need ongoing senior capacity and continuity.

MODEL 03

Staff Augmentation

Specific senior roles added to your existing team — flexible scope, defined period.

TEAM1–3 · specialists

DURATION3+ months

\$7.5k – \$9.5k / eng / mo

Best when you just need particular skills, fast.

Indicative 2026 bands – validated per engagement. Kickoff within 1-2 weeks of scope agreement.

— CASE STUDIES

Proof, **not promises.**

Claims are backed by proof or not made. Five featured engagements — real systems at real scale, owned end to end by senior Nugen teams.

01	Intterra	Emergency management SaaS — real-time geospatial alerting	<500MS • 10K+ RESPONDERS
02	Apploi	Healthcare workforce platform — Inc. 5000, 2024	9,000+ ORGS SERVED
03	Qudos	Mobility & dispatch platform — real-time matching at scale	200 → 1,000 CONCURRENT
04	Optum Store	Prescription subscription commerce — HIPAA-aware at scale	SSR + ISR • GRAPHQL
05	Mastery Coding	K-12 EdTech — sandboxed code execution with AI feedback	LLM-GATED • K8S SANDBOX

CLOUD SAAS • EMERGENCY MANAGEMENT • 2024-PRESENT

Intterra

Real-time geospatial alerting for fire, EMS & law enforcement. Two-tier architecture — Redis GEO hot layer over PostGIS — with Kafka burst absorption and LLM incident summaries gated by LangSmith evaluation.

Nest.js • Redis GEO • PostGIS • Kafka • LangChain • Azure

<500ms

PROXIMITY
MATCH LATENCY

100k+

GPS EVENTS /
DAY

10k+

CONCURRENT
RESPONDERS

HEALTHCARE • WORKFORCE SAAS • 2023-PRESENT

Apploi

All-in-one workforce platform for healthcare — Inc. 5000 (#822 overall, #90 in software, 2024). We led frontend architecture: a typed, hooks-driven design system consumed by multiple product teams.

React • TypeScript • Redux Toolkit • Next.js • GitHub Actions

9,000+

HEALTHCARE
ORGS SERVED

~35%

FASTER TIME-
TO-
INTERACTIVE

~40%

BOILERPLATE
CUT IN
MIGRATION

MOBILITY • DISPATCH • 2018-2022

Qudos

200 → 1,000

CONCURRENT RIDE REQUESTS

Expanding-radius driver matching, live ETA over Socket.io, and a multi-country admin panel — co-architected end to end.

Node.js • PostgreSQL • Socket.io
• React Native • AWS

HEALTHCARE • COMMERCE • 2020-2023

Optum Store

SSR + ISR

HIPAA-AWARE RX COMMERCE

Prescription subscription platform — Next.js front-end, a single typed GraphQL contract, and UI primitives for refill and account flows at scale.

Next.js • TypeScript • GraphQL • Apollo

EDTECH • K-12 • 2022-2024

Mastery Coding

20 → 100

CONCURRENT CODE EXECUTIONS

Sandboxed Python execution on Kubernetes with AI feedback on every student submission — LLM outputs confidence-gated before reaching classrooms.

Python • Docker • Kubernetes • LangChain • Node.js

MORE SHIPPED WORK

Selected Engagements

9 OF 40+ SHIPPED PRODUCTS

FINTECH • WALLET-AS-A-SERVICE

Universal Ledger

Wallet experiences & ops dashboards with module-federated frontends.

FINTECH • MERCHANT PAYMENTS

Blink Payment

Merchant admin dashboard — virtualized data grids, Storybook components.

TELECOM • AT&T

Cricket Wireless

User-management dashboard — state-machine flows, year-long ownership.

SPORTS • INJURY ANALYTICS

Inside Injuries

Player risk scoring with Airflow ETL for high-frequency sports data.

HR TECH • PAYROLL

Dapt Tech

Kafka-driven payroll sync for 10k+ employees across 300+ companies.

TRAVEL • BOOKING

Guide to Iceland

SSR + ISR discovery & booking, LCP tuned under 1.8s on marketing pages.

RETAIL • CORPORATE REWARDS

Lucky Rhino

Shopify & BigCommerce reconciliation with idempotent webhooks.

FINTECH • TAX & BOOKKEEPING

Protax Canada

Unified tax processing — Stripe billing, e-sign, multi-step filing flows.

AI • CONVERSATIONAL SEARCH

Mosaicx

RAG pipeline improving contextual accuracy in conversational AI.

Rebuilt onboarding, doubled activation — without touching pricing.

A workflow SaaS with strong retention but a leaky top-of-funnel. A prior agency had rebuilt onboarding twice and made it worse. We shipped a new activation flow, instrumented it end-to-end, and ran three experiments in eleven weeks.

PRODUCT BUILD

11 WEEKS

SQUAD OF 4

NEXT.JS · POSTGRES · PGVECTOR

“The first agency we’ve worked with that treated activation like a first-class metric — not a UI concern.”

VP PRODUCT · SERIES A B2B SAAS · US

ACTIVATION RATE

+112%

Day-7 activation, cohort over cohort, in 11 weeks.

TIME TO FIRST VALUE

-63%

Median TTFV from 14 min → 5.2 min.

SUPPORT TICKETS

-41%

Onboarding-related tickets per 100 signups.

WEEKS TO PROD

4

First shipped variant of the new flow, live to real users.

Shipped an AI reconciliation agent that actually stayed in production.

The internal team had a working prototype and no path to production — evals were manual, cost was unbounded, and the finance team wouldn't sign off. We rebuilt the system around a tenant-safe RAG, an eval harness in CI, and a per-request cost budget. It went to production in eight weeks.

DEDICATED TEAM

8 WEEKS TO PROD

SQUAD OF 5

LANGGRAPH • BEDROCK • POSTGRES

“They didn’t sell us more AI. They made the AI we already had actually deployable.”

HEAD OF ENGINEERING • SERIES B FINTECH • UK

MANUAL REVIEW TIME

-78%

Hours of ops review per 10k transactions.

COST PER REQUEST

-64%

After routing, caching, and eval-guided prompt cuts.

EVAL COVERAGE

92%

Regression cases in CI before every model change.

HALLUCINATION RATE

<0.4%

On golden set, at ship. Down from ~7% at handover.

PROCESS

How We Work

Working software every week beats a perfect plan. You always know where things stand.

01 • WEEK 0

Kickoff

A one-week intensive: outcomes, constraints, architecture spike, and a written delivery plan you sign off on.

02 • 1-WEEK LOOP

Weekly Ship

Every Friday: a working increment on a live URL, a 10-minute demo, and an honest status — including what we didn't finish.

03 • FINAL 2 WEEKS

QA & Hardening

Regression suite, load runs, security review, and an ops runbook — before we call anything "done."

04 • POST-LAUNCH

Handover

Documented architecture, on-call playbook, a 30-day warranty window, and a knowledge-transfer sprint. All code and IP is yours.

THE WEEKLY ARTIFACT

Live URL • Demo video • Written status • Risk log — not a burn-down chart, not a Jira export

OUR TEAM SHAPE


No Pyramid. No Shadow Team.

A small, flat squad — mostly staff-level, all senior. The people you meet in the pitch are the ones committing code, on the standup, and on the on-call rotation.

- Every engineer **7+ years** in production
- Zero rebadging · zero subcontracting
- Named team from day one — **CVs before contract**
- Overlap with your timezone **4+ hours** daily

A TYPICAL SQUAD

One tech lead. Two-to-four senior engineers. One designer. One product lead. That’s a Nugen team — and it’s the whole team.



GET STARTED

Three Ways to Start

We'd rather earn the next hour than sell the next quarter. Every entry point below ends in a written artifact you keep — whether or not we work together.

STEP 01 • FREE

45-min intro call

Frame the problem, sanity-check the shape of a solution, and get a candid read on whether you actually need us.

BOOK THE CALL →

STEP 02 • FIXED FEE

Two-week scoping sprint

An architecture spike, a delivery plan, a fixed-fee proposal, and — yes — the source of the spike, whether or not you continue.

START SCOPING →

STEP 03 • LIVE TEAM

Four-week pilot sprint

A named senior squad, a Friday demo cadence, and a hard success gate at week four before either side commits further.

RUN A PILOT →

LET'S BUILD

GET STARTED



Your product, owned like **our own.**

If you're a founder or engineering leader who needs to ship something that matters — that's exactly what we do.

[Book a discovery call](#)

TALK TO US

info@nugeninfo.com

ON THE WEB

nugeninfo.com

PHONE

+91 78148 52571

WHERE WE SHIP

US · UK · India