



AI-ENABLED TEAMS

Three AI engagement models

This document presents three purpose-built team configurations from the Devshore Partners talent pool. Each team addresses a distinct AI engagement model — from engineers who use AI to move faster, to specialists who ship AI-native products, to practitioners who automate entire workflows using agentic tooling. All teams are available for remote engagement.

12 Engineers	10+ Avg. years experience	3 Team configurations	100% Available for remote
------------------------	-------------------------------------	---------------------------------	-------------------------------------

TEAM 01 | AI-Augmented Engineers

Proven engineers who use AI as a force multiplier to ship faster and smarter

This team is made up of senior software engineers with strong production track records who have actively integrated AI copilots, code-generation tools, and AI-powered testing into their day-to-day delivery. They are not AI researchers — they are engineers who have learned to wield AI to compress timelines, reduce errors, and tackle more ambitious architecture. Ideal for clients who want experienced developers that arrive with an AI-accelerated workflow already in place.

BEST FOR

Product builds and feature delivery where speed matters, fintech or enterprise backend work, teams that need to scale throughput without scaling headcount.

TEAM MEMBERS

Arif — Senior Backend & Fullstack Engineer · 15+ yrs

- Senior fintech and banking specialist — led engineering teams building Loan Origination Systems, fraud-detection pipelines, and microservices handling millions of daily transactions.
- Actively uses OpenCode + DeepSeek V4 Pro / Gemini Flash for architecture design, scaffolding, test generation, and technical docs — cutting dev time meaningfully without sacrificing standards.
- Founder-engineer: built GoShow (WhatsApp-first queue management SaaS) and Kopintar (cooperative management platform) solo using AI-accelerated development, including 200+ E2E tests and full security audits.
- Stack: Java / Spring Boot, NestJS, ReactJS, PostgreSQL, Kafka, Docker, MCP / Agentic AI workflows.

Eko — Senior Fullstack Engineer · 9+ yrs

- Built and shipped an AI-powered Ingredients Scanner using Ollama (Qwen) and Gemini for a 12,500+ user global platform — covering prompt engineering, parameter tuning, and latency optimisation.
- Led a dev team of 5, managing end-to-end delivery of 6 major enterprise web applications; rigorous daily code reviews.

- Strong backend fundamentals: NestJS, Node.js, PostgreSQL, Redis, GraphQL, AWS. Also built Python/Selenium web-scraping pipelines and performance-optimised 4M+ product datasets.
- Known for translating complex business requirements into clean, maintainable architecture.

Dedy — Fullstack Engineer · 6+ yrs

- Built an AI-powered semantic search feature using OpenAI embeddings and Pinecone, letting employees find training materials by context rather than keywords.
- Implemented an AI recipe recommendation system and has worked on RAG-based chatbot frontends for international clients.
- Solid CI/CD automation: GitHub Actions, Docker; cut deployment time by 66% on one engagement.
- Comfortable across the stack: Next.js, React Native, Flutter, Node.js, Golang, Rust, PostgreSQL.

Budi — Senior Backend Engineer · 8+ yrs · FHIR R4 & Healthcare Data

- Architected Python services syncing 2,000+ clinical resources daily across a federated national healthcare platform — reduced sync time 6x and cut API costs by 90% through batching, conflict resolution, and delta-based fetch strategies.
- Uses Claude Code as an agentic development tool to autonomously generate FHIR R4 converter structures across 15+ health information system integrations — AI-driven schema analysis, mapping scaffolding, and iterative code refinement in production.
- Designed a terminology mapping system covering 35+ government health systems with 475+ LOINC, 200+ SNOMED CT, and 180+ ICD-10 mappings; operates 58 Cloud Run microservices and a BigQuery medallion warehouse (50+ datasets) serving 42+ community health centres.
- IEEE-published researcher in big data analytics (200M+ records, PySpark); NLP/sentiment analysis background. Stack: Python/FastAPI, Java/Spring Boot, TypeScript, GCP, FHIR R4, Claude Code.

TEAM 02 | AI & LLM Product Builders

Engineers and data scientists who build AI-native products end to end

This team specialises in building products where AI is the core, not a feature. They have shipped agentic pipelines, voice and vision AI systems, LLM-powered recruitment engines, GenAI document processing, and ML recommendation systems — all in production. Combined, they cover the full stack from data science and model integration through to backend services and cloud infrastructure. Ideal for clients building AI-first products or embedding LLMs deeply into existing platforms.

BEST FOR

Greenfield AI products, LLM integration into SaaS, agentic automation pipelines, ML recommendation or matching systems, voice/vision AI features.

TEAM MEMBERS

Rafinno — Senior Fullstack Engineer · Engineering Manager · 13+ yrs

- Sole engineer on two full AI product builds: an AI wellbeing platform (OpenAI Whisper voice check-in → GPT-4o extraction → real-time WebSocket) and an AI dating companion app (multi-provider LLM: OpenAI + LLaMA via AWS Bedrock + Together AI, plus voice chat via VAPI and DALL-E image gen).
- AWS-native: provisions full production stacks from scratch with Terraform (~70 resources per engagement) and GitHub Actions CI/CD. Comfortable owning infra, backend, and mobile simultaneously.
- Former Engineering Manager at GovTech Edu — cut a plagiarism-validation workflow by 70% and tripled approved submissions across 163,000+ records.

- Stack: TypeScript, Node.js, Go, React/React Native, PostgreSQL, Redis, AWS (ECS, Bedrock, RDS, S3), Terraform.

Denny — AI / LLM Engineer + Senior Fullstack · 13+ yrs

- Purpose-built agentic pipelines in production: an autonomous recruitment engine (Claude API → resume parsing → semantic candidate-job matching → Go tool functions → Redis-queued agent loops with human-in-the-loop approval gates) and a self-hosted Claude Code DevOps automation system monitoring ~148 domains via MCP.
- Built an in-house AI code-review system integrated into Bitbucket pull-request workflows using the Anthropic Claude API with structured-JSON output prompts, multi-tenant RBAC, and per-client AI cost tracking.
- Delivered an AI business intelligence suite over Xero and HubSpot data: predictive financial dashboards, budget-vs-actual trackers, and a conversational financial-maturity advisory chatbot.
- Stack: Go, TypeScript, NestJS, React/Vue, PostgreSQL, Redis, Elasticsearch, Docker, Jenkins, Claude API, MCP.

Kurnia — Senior AI & Data Scientist · 5+ yrs · AWS Certified ML Engineer

- Built end-to-end ML pipelines for a 50M+ record job-platform dataset: ETL/ELT, NLP-based CV-to-job matching using LLM embeddings, automated biweekly retraining via AWS CDK + Step Functions — 4x business metric improvement, 45% infra cost reduction.
- Deployed production GenAI solutions: RAG chatbots, document QA, OCR-to-LLM pipelines, AI transcription and diarisation (pyannote.audio + OpenAI Realtime API), and real-time Kubernetes monitoring with LLM summarisation.
- Active mentor: 500+ students trained in data science, ML, and AI engineering across Dibimbing.id and Kelas.com; consistent 4.8-4.9 / 5 rating.
- Stack: Python, FastAPI, LangChain, AWS SageMaker / Lambda / Bedrock, PostgreSQL, PySpark, Kubernetes.

TEAM 03 | Agentic Workflow & AI Operations

Specialists in Hermes-style automation, AI-enabled QA, and AI-native product oversight

This team works at the intersection of AI and operational leverage. They build and operate the orchestration layer — agentic tools that replace multi-step manual workflows, QA systems that evaluate non-deterministic AI outputs, and product leadership that drives AI initiatives from strategy to delivery. Warda literally architected and shipped Hermes, a production agentic NestJS tool that routes tasks across specialised LLM agents with human-in-the-loop controls. Najmi runs AI-enabled QA with LLM-as-a-Judge evaluation frameworks. Derryl deploys AI agents via MCP to automate research, content, and reporting pipelines. Lian bridges AI product management with data-driven health informatics. Together they represent the governance, quality, and operations layer that AI-first organisations need.

BEST FOR

Organisations adopting AI at scale who need workflow automation, QA coverage over AI outputs, and product management that can own the full AI product lifecycle.

TEAM MEMBERS

Warda — Senior Fullstack Engineer (Mobile + Backend + AI) · 12+ yrs

- Built Hermes — a production agentic tool with a Flutter frontend and a NestJS backend that connects an LLM to internal tools, routing tasks across specialised agents with human-in-the-loop approval and full action logging.
- Developed AI vision (VLM) pipelines for smart site monitoring: face detection, PPE compliance, crowd monitoring, and licence-plate recognition across 20+ concurrent live camera streams — fine-tuned PyTorch VLM models for ~15% accuracy gain.

- Integrated OpenAI to generate personalised health archetypes for a multi-platform health app, reducing content curation time by 50%; built the full dbt transformation layer for health data idempotency.
- Stack: Flutter, React Native, NestJS, Python/FastAPI, PostgreSQL, Docker, Azure DevOps, PyTorch, LLMs/VLMs.

Derryl — Hybrid PM / CEO Office Lead · 6+ yrs

- Bridges C-suite strategy and hands-on product delivery: managed \$3M+ government tech contracts across 8 provinces, aligning Product, IT, and Finance; currently building a core AI product after leading GTM and commercial operations.
- Originated and shipped a RAG knowledge base + WhatsApp AI customer-service agent — engineered unit economics to ~11x lower cost-per-session than market and carried it from concept to a signed enterprise engagement.
- Deploys AI agents via MCP across Google Drive, Notion, and Slack, collapsing multi-hour research/QA/reporting cycles into on-demand automated runs; built reusable custom Claude skills and an automated content pipeline (work artefacts → publish-ready posts + image prompts).
- Trained in Scrum by a certified Scrum instructor; authored a Scrum Playbook adopted across an engineering organisation. Now coaches PM and CEO Office candidates as a hiring-manager mentor — strong read on what makes product leaders effective.

Lian — AI Product Manager · Digital Health & Data Informatics · 13+ yrs

- Led AI-enabled digital health initiatives funded by the Gates Foundation SPHERES programme: OCR automation, clinical decision support prototypes, and predictive analytics improving maternal and child health outcomes.
- Built an AI healthcare chatbot from scratch — NLP workflows, offline on-device LLM models, three-year AI roadmap from model improvement through to clinical decision support deployment.
- Led development of a national COVID-19 contact-tracing app — now used by 172,000 field workers with the web version serving 100,000 daily users at 80,000 accesses per minute.
- Bridges technical delivery and stakeholder management across the Ministry of Health, WHO, and USAID; publishes at international digital health conferences.

Chiqo — DevOps / Platform Engineer · 4+ yrs · CNCF Contributor

- Manages end-to-end cloud infrastructure and platform strategy for a Gates Foundation-funded national healthcare system across multiple district deployments — balancing scalability, security, and strict budget constraints using AI-assisted architecture evaluation.
- Built and deployed Model Context Protocol (MCP) server integrations on a federated healthcare gateway, exposing standards-based healthcare data as agent-accessible tools for secure, governed AI agent integration across national and district systems.
- Pioneered agentic SDLC tooling: built reusable Claude Code Skills (SKILLS.md) and agentic prompting workflows to standardise CI/CD pipeline authoring, GitOps configuration, and infrastructure-as-code generation. Active Cloud Native Computing Foundation (CNCF) open-source contributor.
- Multi-award winner: ASEAN AICTA 2021 3rd place and national Kominfo IdenTIK winner. Stack: Go, TypeScript, Kubernetes, Terraform, ArgoCD, GCP, Docker, Claude Code, MCP, pgvector.

Rahmat — AI/ML Security Engineer · Application Security Specialist · 8+ yrs

- OSCP and CEH certified; led penetration testing, red-team engagements, and DevSecOps integration across fintech, telco, and quick-commerce — including ISO 27001 compliance programmes and incident response.
- Building LLM-assisted security tooling in production: fuzzing payload generators using Nuclei templates, ML-based log-anomaly detection pipelines, and AI-assisted secure code review integrated into CI/CD workflows.
- Defined internal governance for safe LLM adoption at an engineering organisation: prompt-injection controls, data-leakage prevention, and secure AI-assistant usage policies — covering the emerging attack surface of AI systems.

- Stack: Python, Go, BurpSuite, Nuclei, Ghidra, scikit-learn, PyTorch (applied), CI/CD security automation, Linux infrastructure hardening.

Next Steps

Teams can be engaged as full units or as individual contributors depending on project scope. All members are available for remote collaboration and are accustomed to asynchronous, cross-timezone delivery.

01

Choose your model

AI-augmented engineering, AI product build, or agentic workflow — or a mix.

02

Intro call

Meet the relevant team members and align on scope, timeline, and engagement model.

03

Proposal

Receive a tailored commercial proposal within 48 hours of the intro call.



DEVSHORE PARTNERS
TALENT BEYOND BORDERS