

Nitish Gupta

Gurugram, India · nitish.gupta101@outlook.com · linkedin.com/in/nitish11 · github.com/nitishiot

AI Product Manager — Agentic & LLM Systems · 15 years in Product/SaaS, 8 years as Product Owner / Product Manager · Sovereign & Regulated AI

SUMMARY

Product manager for LLM and agentic systems in a regulated enterprise, moving into India's defence and security sector. Self-funded study since January 2025 (Basel Institute on Governance, Rashtriya Raksha University, BITS Pilani CRENS under CERT-In aegis, CNAWS) and a sovereign, air-gapped OSINT platform designed and delivered end to end for a defence innovation challenge. Sole product owner of a conversational AI assistant at Fidelity International taken from Sprint 0 to production go-live; one of three product managers on an enterprise AI platform, owning model selection, guardrails and observability. Fifteen of those years sit in product and SaaS, with the last eight owning products end to end.

PRODUCT MANAGEMENT PROFILE

- **Zero-to-live AI product ownership.** Sprint 0 to production go-live on a "high material risk" AI system: discovery, perimeter definition, guardrail design, governance approval.
- **Product definition for probabilistic systems.** PRDs, epics, stories and acceptance criteria where the criterion is a threshold, a coverage target and an explanation requirement: a credibility score from 18 signals, under 200 ms, ≥95 % coverage, unscorableables flagged.
- **Enterprise AI platform ownership.** One of three product managers on the AI Centre of Excellence programme: model selection, guardrails, observability, from vision to execution.
- **Agentic and LLM systems.** LLM copilots and RAG at enterprise scale; agent gateway, agent registry and MCP gateway/registry from design into build. Orchestration and planning patterns (ReAct, task decomposition, multi-agent coordination), tool-calling design, four memory types.
- **Evaluation, trust and AI security.** Evaluation as executable policy: task success, hallucination and abstention rates, latency and cost per task. Authored the AI-security roadmap (prompt injection, data leakage, policy violation, unsafe tool use), routed into SOC workflows with human checkpoints named and incident response aligned to NIST.
- **Spec-driven delivery.** Drove adoption of spec-driven development in the platform team's workflow; ran the same discipline independently: 45 phase specifications written before code.
- **Regulated delivery.** Single point of contact for senior leadership, regulators, engineering and vendors; led cross-functional teams across product, engineering, AI/ML, security, data and compliance.

EXPERIENCE

Fidelity International — Product Owner / Product Manager since 2018 11/2016 – Present · Gurugram, India

Conversational AI assistant — Sole Product Owner · Q4 2024 – Present - Wrote the PRD and shipped what it specified: epics, user stories and acceptance criteria carried through governance approval to production go-live. - Cleared a "high material risk" AI system through full governance: a three-layer guardrail design (prompt-injection detection, an anti-hallucination check, a filter blocking investment-style advice) through PIA/DPIA and architecture-review approval. - Own a public customer-facing assistant and an internal ethics-advisory bot on one conversational UI stack.

Enterprise AI Platform (AI Centre of Excellence) — Co-owner: model selection, guardrails, observability - One of three product managers on the programme; co-authored the platform's Level 1 Conceptual Design with the architecture team. - Benchmarked the enterprise agentic-platform category (Moveworks, Aisera, ServiceNow AI): reference architectures and demo environments; market analysis fed build-versus-buy and roadmap decisions, with ROI and time-to-value assessed per initiative through POCs and POVs. - Contributing to the platform's agent gateway, agent registry and MCP gateway/registry, specified and in build through 2026. - Took the platform from an internal LLM assistant (API surface grown to ~25 endpoints) through a Copilot proof-of-concept to a governed platform plus marketplace; the marketplace has been live since January 2026. - Ran discovery workshops and workflow deep-dives with operations, reconciliation, security and compliance teams; defined personas and pain-point models. - Enabled enterprise LLM adoption across several assistants and coding tools under responsible-use governance.

Earlier in the role: enterprise digital investment products and native mobile applications, prioritised by product analytics, NPS and customer feedback; an operational data store integrating 30+ systems of record; regulatory-technology

delivery (AML, KYC, MiFID II, UK GDPR); led a 10–15-person business-analyst team through time-critical EU regulatory change in a SAFe environment.

Earlier experience - Senior Business Analyst, payment security and fraud detection, Collabera (payments client) · 08/2015 – 09/2016 · Pune — requirements, epics, stories and acceptance criteria for card-services payment-security and fraud-detection platforms - Business Analyst, Capita · 06/2014 – 07/2015 · Pune - Software Engineer, Hexaware Technologies and Binary Softzone · 07/2009 – 11/2013 · Pune

INDEPENDENT PROJECTS

A sovereign OSINT platform built for a defence innovation challenge · 2026 - Designed and delivered the architecture end to end: a 195,657-item live corpus; an 18-signal credibility engine (Bayesian log-odds + XGBoost) at 100 % signal coverage and 0.64 ms P95 latency; hybrid dense + full-text retrieval fused by RRF, 0.182 ms vector-leg P95. - Wrote the product definition before the code: 45 phase specifications, 29 phase sign-off tags, 1,711 tests passing (2026-08-28). - An evaluation gate that rejected a production candidate: citation validity 100 % (43/43), abstention 12/12. The same gate blocked a weighting proposal that would have cost 6.6 % against an 85.9 % floor. - Agent operating discipline, measured: profiled 41 of the build's own agent working sessions for context-cost amplification (133× mean, 364× peak per tool result), then enforced policy: a model-routing rule that moved one model's share of spend from 21.6 % to 1.9 %, nine automated lifecycle hooks checking cost and process drift, and four role-scoped review agents with separated write permissions. - Build-versus-buy decisions recorded with reasons: pgvector over FAISS · RRF over a weighted blend · LangChain assessed and not adopted · hosted evaluation tooling excluded on air-gap grounds · YOLOv8 rejected because AGPL licensing reaches the delivered platform.

FrameProof — screen-recorded lectures to verified markdown · 2026 · github.com/nitishiot/FrameProof - Pipeline from long-form video to structured, citable documents: audio transcription plus on-screen visual extraction. 19 of 19 lectures delivered; 34,097 frames sampled, 13,538 kept. - Seven verification layers (V0 to V6), most of them model-free: 8,492 grid cells re-read by OCR; 79 real errors found and fixed, each counted from the verification record that found it. - Delivery status generated from disk state, so the tracker cannot drift from what exists.

CERTIFICATIONS

Defence, security and intelligence - BITS Pilani CRENS — Professional Development Course in Cybersecurity, under the aegis of CERT-In, MeitY (Jul–Sep 2025) - Rashtriya Raksha University — Information Warfare (May 2025); Cybercrime, Cyber Laws and Technology (May 2025) - Atma Nirbhar Soch — PLA certificate series (2025); Tactical Drones and Counter-UAS - Basel Institute on Governance — Open-source Intelligence (Jan 2025) - CNAWS — Certificate Course in OSINT: Collection, Analysis and Applications (Feb 2026)

AI, data and product - IISc Bengaluru, Centre for Continuing Education — Large Language Models: A Hands-on Approach (in progress) - IIT Madras — Advanced Certification, Data Science and Artificial Intelligence - Scaled Agile — SAFe 6.0 Agilist; SAFe 5.0 Product Owner / Product Manager

TECHNOLOGY AND PRODUCT TOOLING

Agentic and LLM: orchestration and planning, tool-calling design, agent memory, MCP, LLM APIs, frameworks assessed (LangChain, LangGraph, AutoGen, CrewAI patterns) · Retrieval and data: RAG, pgvector/HNSW, hybrid retrieval with RRF, embeddings, Kafka, Snowflake, SQL · Evaluation: gates as code, hallucination and abstention rates, cost per task, guardrails, failure-case review · Product craft: PRDs, KPI trees, anti-metrics, roadmaps, discovery, POC/POV, build-versus-buy, RICE, Kano · Governance: PIA/DPIA, UK GDPR, MiFID II/MiFIR, AML, KYC, SAFe/Scrum/Kanban, JIRA, Confluence, Azure, AWS

EDUCATION

B.Tech, Information Technology — Bharati Vidyapeeth University, College of Engineering, Pune (2009), first class with distinction