

Abhishek Mishra

VP Product & Growth · AI Products for Insurance, Reinsurance & Risk

Hartford, CT | abhishek.mishra1218@gmail.com | (860) 904-3391 | abhishekmishra.ai | [LinkedIn](#) | [Substack](#) | [Medium](#)

PROFESSIONAL SUMMARY

VP-level AI product leader with nearly two decades building and commercializing enterprise products across **insurance, reinsurance, and catastrophe intelligence**. Owns product strategy, roadmap, P&L, and go-to-market for AI platforms across claims, underwriting, and risk, with delivery experience in nine markets on four continents.

Pairs deep insurance domain expertise with **hands-on AI engineering** — designing agentic architectures, evaluation pipelines, and production guardrails, and building working prototypes personally. Equally comfortable presenting an executive business case and reviewing an architecture with engineering teams.

Operating philosophy: *AI amplifies human judgment rather than replacing it. In regulated domains the discipline that matters most is knowing where AI belongs — and where deterministic controls must stay in place.*

HOW I BUILD AI PRODUCTS

- **Domain first, model second.** Most enterprise AI failures are workflow and domain failures, not model failures. Starts with the decision being made and the person making it, then works back to the architecture.
- **Deterministic controls around probabilistic systems.** Rule-based triage, database-as-source-of-truth, and human approval gates on money and denials — AI proposes, people decide.
- **Observability ships with the product.** Evaluation suites, red-teaming, and drift monitoring built in from the start rather than added after launch — a prerequisite for regulated deployment.
- **Prototype to de-risk.** Builds working prototypes to test product hypotheses and expose integration gaps before committing significant engineering investment.

IMPACT HIGHLIGHTS

- Delivered AI-driven claims and catastrophe products for global carriers — reducing client **claims expense 20%, loss leakage 4–6%, and processing time 30%**
- Built and commercialized AI products generating **\$1M+ ARR within 18 months**, and grew a \$30M service-line P&L by 40%
- Led AI and cloud transformation for **top 25 global insurers** across P&C, Commercial, and Workers Compensation in **nine markets**: the US, UK, Denmark, Singapore, Hong Kong, India, Australia, Saudi Arabia, and the UAE
- Incubated a venture-funded claims start-up inside Cognizant — **\$6M raised, 25-person founding team, 9 enterprise customers**
- Built and led **global product, data science, and engineering teams of 30+** across three continents
- Co-inventor, **issued European patent** for Rapid Damage Assessment Risk Score; pending in the US and Australia
- Designs and builds full-stack AI applications independently — **four deployed prototypes** across claims, catastrophe analytics, FNOL intake, and voice ordering

CORE COMPETENCIES

AI Product Strategy & Leadership AI-native product strategy · product operating models · build-vs-buy and model selection · AI pricing and monetization · responsible AI and model governance · executive stakeholder management	AI / ML Foundations Machine learning · deep learning and CNNs · computer vision · predictive analytics · model evaluation and production deployment · MLOps foundations
Insurance & Regulated Domain Claims lifecycle & FNOL · underwriting intelligence · reinsurance · catastrophe response and exposure analytics · fraud/SIU · core platform modernization · NAIC-aligned governance	Growth, Commercial & P&L P&L ownership · consulting-led GTM · pricing and packaging · platform monetization · partner and alliance strategy
Agentic AI & Applied Engineering Multi-agent orchestration · RAG · MCP tool integration · AI evaluation, red-teaming, drift monitoring and observability · deterministic guardrails · schema-first architecture · voice agents	Leadership & Delivery Global team building and coaching · program governance · Agile/SAFe · cross-cultural delivery · executive communication · vendor and partner strategy

PROFESSIONAL EXPERIENCE

Product Lead & Head of Growth (Vice President) — Claims AI & Data Products, Risk Data & Solutions

Swiss Re · New York / Connecticut · July 2022 – Present

- Own global product vision and roadmap for an **AI-driven claims platform addressing climate and catastrophe risk**, now serving carrier clients across multiple regions
- Reduced client claims expenses by **over 20%** and improved claims processing time by **30%** through deep learning applied to imagery, weather, and event data
- Established the architecture and governance model for AI/ML pipelines, agentic and RAG workflows, and evaluation/observability practice — integrated with Guidewire ClaimCenter
- Scaled a global team of 30+ engineers, designers, and data scientists across the US and Europe, and defined the product operating model connecting research, engineering, and commercial teams
- Designed pricing models and regional GTM playbooks across markets, generating **\$1M ARR within 18 months** while owning global product P&L, budgets, and cost controls

P&C Products & Services Market Leader (Director) — P&L Owner, Core Platform Transformation & Innovation

Cognizant Technology Solutions · Teaneck, NJ · May 2020 – July 2022

- Repositioned a \$30M service line around platform-led transformation, growing it 40% in two years by securing 5 new clients including 4 multi-year deals
- Led Insurtech and platform solution strategy, sales, and delivery (Guidewire, Duck Creek, Unqork) across US P&C insurers
- Directed program governance for policy administration, claims, and distribution implementations across personal, commercial, and Workers Compensation carriers
- Designed joint GTM offerings with AWS and Insurtech partners, building a \$0.5M net-new revenue pipeline within one year

Senior Product Manager (Associate Director) — Cognizant Accelerator, AI-Based NATCAT Claims Solution

Cognizant Technology Solutions · Teaneck, NJ · Oct 2017 – Apr 2020

- Incubated a property claims venture inside the Cognizant Launchpad Accelerator, securing \$6M in funding on the thesis that catastrophe claims would shift from manual inspection to AI-assisted assessment
- Built the 25-person founding team — full-stack engineers, UI/UX designers, data scientists, ML engineers, and product managers
- Owned product strategy, growth, and market entry; onboarded 6 beta and 3 commercial customers in 3 years on an AWS-based web and mobile claims platform

Consulting Leader (Senior Manager, Consulting) — US, UK & Emerging Markets

Cognizant Technology Solutions · Jul 2015 – Sep 2017

- Drove business and technology consulting sales and execution across the US, UK, Denmark, Singapore, Hong Kong, India, Australia, Saudi Arabia, and the UAE — delivering for insurers operating under distinct regulatory regimes

Key engagements: Post-merger claims org consolidation (UK insurer) · Lean implementation across 23 underwriting offices (US insurer) · Workers Compensation claims transformation (US insurer) · Technical accounting optimization (global insurer) · Guidewire PolicyCenter implementation (US insurer)

AI-NATIVE PRODUCT PORTFOLIO · SELF-BUILT

Four full-stack AI applications architected and built independently, deployed on Vercel with source on GitHub — built to test how AI-native products should work in regulated domains. Live demos and case studies at [abhishekmishra.ai](#).

FACIA — Fast Cat Impact Analytics

Identifies impacted insured locations within 30 minutes of a weather event, replacing hours of manual spreadsheet triage with an auditable, GIS-backed pipeline.

- Architected six independently testable agents (geocoding, weather, footprint, impact, reporting, supervisor) evaluating full NOAA weather grids to synthesize event footprints and classify policies as directly impacted or adjacent
- **Design decision:** made the database — not agent-to-agent messaging — the orchestration bus, so every run is restartable, auditable, and independently testable
- **Judgment call:** the LLM writes narrative prose only; every number in every report comes from the database, never the model

Stack: React · TypeScript · Mapbox GL · FastAPI · CrewAI · n8n · Supabase PostgreSQL + PostGIS · NOAA GRIB2 · Claude · Vercel

AI-Native Property Claims Platform

Runs the full claims lifecycle — intake through payment — on one AI-native platform, with automation confined to where it is provably safe.

- Full lifecycle spanning multi-channel FNOL (form, voice, phone), a role- and region-aware Claims Desktop, fraud analytics, payments, and AI quality operations on a single shared claim model
- **Design decision:** agents propose, humans approve — the desktop assistant can never execute; fraud holds hard-block payout regardless of model confidence
- **Judgment call:** shipped observability as product, not afterthought — golden and red-team eval suites, drift monitoring, and promotion gated on thresholds (red-team 100%, fraud accuracy ≥ 85%)

Stack: Next.js on Vercel · MCP server on Render · n8n Cloud · ElevenLabs · Langfuse · Supabase RLS · Stripe Connect · Twilio

WC FNOL Copilot

Compresses workers' compensation intake from fragmented phone, paper, and email threads into one verified, triaged, submittable incident record.

- Mobile-first PWA guiding injured workers through deterministic intake graphs, with optional AI chat extracting multiple fields per message and supervisors contributing via secure invite links
- **Design decision:** triage scoring, validation, and completeness gates remain rule-based and explainable — AI phrases the questions, but never scores the injury
- **Engineering rigor:** multi-tenant row-level security treated as a first-migration design choice rather than pre-launch hardening; 80+ automated tests across domain, contracts, RLS, and E2E

Stack: Next.js PWA · Supabase (Auth, RLS, Edge Functions, Storage) · TypeScript domain layer · Zod · rule-based triage

Restaurant Voice Ordering & Operations Copilot

Conversational ordering in a safety-critical domain: guests order naturally by voice while allergen decisions remain deterministic and server-enforced.

- Voice- and chat-first copilot handling multi-turn ordering, disambiguation, and full read-back confirmation, with a staff operations portal for kitchen queues and escalations
- **Design decision:** the agent never writes to the database — every menu fact and cart change flows through validated server tools, with explicit guest approval before a kitchen ticket exists
- **Judgment call:** when allergen data is incomplete the system warns and escalates to staff rather than guessing; safety copy is templated from backend classification, never model-generated

Stack: ElevenLabs Conversational AI (voice + text) · Supabase · Edge Functions · n8n · React monorepo · Vercel · 114 automated tests

WRITING, PATENTS & SPEAKING

- Writes on AI-native product management, agentic AI in insurance, and the evolving product-manager role — published on [Substack](#), [Medium](#), and [abhishekmishra.ai](#)
- Co-inventor, **European Patent No. 4409484** — Rapid Damage Assessment Risk Score (issued 2025); pending in the US and Australia
- Swiss Re publications: *Using data to get ahead of the storm* (2022); Swiss Re Annual Report 2022
- Cognizant publication: *Right Touch Commercial Underwriting Playbook* (2021)
- Speaker, Duck Creek Formation (Orlando, 2022): *How Commercial Insurers Are Transforming Underwriting with SaaS Platforms*

CERTIFICATIONS & EXECUTIVE EDUCATION

Product School	AI Product Management · Product Growth · Advanced AI Agents · AI Evals · Product Leader · Product Marketing (2025)
Product Faculty	AI Product Management · Product Sense (2025)
Wharton	Product Management & Strategy, Wharton Executive Education (2025)
The Institutes	Associate in Insurance (AINS), 2025 · Associate in Claims (AIC), in progress — expected 2026
Cloud & AI	AWS Certified Cloud Practitioner & AI Practitioner (2025) · Microsoft Certified Azure Fundamentals & Azure AI Fundamentals (2025)
Accelerator	Cognizant Startup Accelerator Graduate (2018)

EDUCATION

Xavier Institute of Management	College of Engineering & Technology
MBA, Marketing · 2007	B.Tech, Instrumentation & Electronics · 2005

TECHNICAL FLUENCY

AI & ML	Deep learning & CNNs · computer vision · LLMs · RAG · agentic & multi-agent orchestration · AI evaluation, red-teaming & drift monitoring
Agentic Tooling	LangChain / LangGraph · CrewAI · n8n · Model Context Protocol (MCP) · Langfuse · ElevenLabs
Build & Languages	Cursor · Claude Code · Python · TypeScript · SQL
Cloud & Data	AWS (Bedrock, SageMaker) · Azure (Azure ML) · Supabase · Vercel · PostgreSQL / PostGIS
Enterprise Platforms	Guidewire (ClaimCenter, PolicyCenter) · Duck Creek · Unqork