



AI Transformation Leader

Governed, working AI — grounded in ~15 years of operational resilience across regulated, global organisations

AI transformation leader who ships **governed, working AI** in regulated, global organisations. Has conceived and built **more than eight AI applications**, four now in active in-house development, and independently shipped an AI compliance platform navigating the EU AI Act, GDPR and DORA. At AstraZeneca, leads AI transformation within Oncology R&D and partners with the Enterprise AI team, Technology Governance, Legal and Information Security to operationalise AI governance across a 10,000-employee organisation — building, in daily customer-driven sprints, a proof-of-concept **enterprise AI governance platform** that classifies AI submissions against the EU AI Act and internal frameworks via a deterministic routing layer and a multi-agent assessment pipeline with full audit trail. Designs multi-agent LLM architectures across Gemini, Mistral and Claude, drives rapid **proof-of-concept** → **production** delivery, and leads adoption, acculturation and training from prompting fundamentals through Copilot Studio agents. Public work includes **OpRes Horizon** (live at opres.structuredresilience.app), a regulatory-intelligence platform spanning 90+ countries. This AI delivery is grounded in ~15 years of operational resilience and crisis leadership — including Incident Command on national activations such as the Colonial Pipeline cyber breach — giving deep fluency in risk, governance and regulated environments. Fellow of the Business Continuity Institute and committee member of the BCI AI SIG; dual U.S./French citizen, fluent English, intermediate French.

CORE COMPETENCIES

AI Strategy, Roadmap & Use-Case Portfolio	AI Governance & Responsible-Use Frameworks
Multi-Agent LLM Architecture (Gemini · Mistral · Claude)	Rapid Proof-of-Concept → Production Delivery
EU AI Act, GDPR, DORA & NIS2 Operationalisation	RAG Pipelines & AI-Assisted Development
AI Adoption, Acculturation & Training	Model-Drift, Bias & Fairness Monitoring
Power Platform & Microsoft Copilot Studio	Operational Resilience & Crisis Domain Grounding

CAREER EXPERIENCE

Associate Director — AI Transformation (Business Continuity & Crisis Management) 2023 – Present

AstraZeneca, Oncology R&D, Gaithersburg, MD (Global Remit) · North America · Europe · Asia

Lead AI transformation within the resilience function for a 10,000-employee R&D organisation. Partner with the Enterprise AI team, Technology Governance, IT, Legal, Information Security and Data Engineering to operationalise AI governance; act as de facto product owner for internally developed AI applications across their full lifecycle.

AI Governance, Risk Classification & Approval Workflows

- Lead daily customer-driven sprints building a **proof-of-concept enterprise AI governance platform** with the Enterprise AI team — designed to evaluate AI submissions against the EU AI Act and internal frameworks. A deterministic classification layer routes each submission to a fast-track or full-assessment lane; a multi-agent pipeline then scores it, with human reviewers retaining final authority on higher-risk tiers and every agent call and decision logged and exportable for audit.
- Designed the platform's approval workflow, documentation requirements and lifecycle checkpoints, anticipating **EU AI Act high-risk classification** guidance (Article 6) — realistic failure mode, not commercial label, determines classification.
- Conceived and built a **model-drift analysis pipeline** that measures divergence across runs of internal multi-agent pipelines and generates targeted refinement recommendations — already used to improve several pipelines — operationalising the EU AI Act's post-deployment monitoring obligation.
- Currently workshoping enterprise-wide **AI bias and fairness monitoring** with the Enterprise AI team — critical in pharma, where biased outputs can affect patient selection and clinical decision support.

- Led the first enterprise-wide **GenAI emerging-risks scan** across the 10,000-employee organisation, establishing proactive data-governance and responsible-use protocols and translating findings into internal policy for Legal, IT and business leadership.

AI Use-Case Portfolio, PoC-to-Production & Delivery

- Conceived and built **more than eight AI applications** across enterprise, governance and productivity domains; four are now in active in-house development, demonstrating the path from rapid proof-of-concept toward production.
- Co-developed **in 24 hours** with the Enterprise AI team a multi-agent AI grant & tax-incentive search platform that autonomously identifies grants and tax credits applicable to internal AI projects worldwide, with a no-code admin layer where natural-language feedback becomes RAG-stored rules. Demoed up to CFO level.
- Built and deployed independently an **AI compliance platform** navigating the EU AI Act, GDPR and DORA — a centralised AI project register, automated risk classification and a multi-agent LLM pipeline grounded by RAG in the authoritative regulatory text, so new guidance is operationalised by document upload rather than retraining.
- Built and deployed independently **OpRes Horizon** (live at opres.structuredresilience.app) — a public regulatory-intelligence platform across 90+ countries and 8 sectors; a 3-agent, 3-stage pipeline (Mistral search → Claude synthesis) produces board-level country profiles with a 0–4 regulatory-intensity score, a live league table and bilingual EN/FR output with Word export. Mistral routing chosen for European data sovereignty.
- Built independently a **six-agent clinical-trial supply-intelligence** proof-of-concept (React/TypeScript, Express/Node, Google Cloud Run) with an audit-first architecture — no AI output triggers automated action; every decision routes through human review.

Adoption, Acculturation & Enablement

- Designed and delivered structured internal training from prompting fundamentals through building **Microsoft Copilot Studio** agents; deployed a Copilot Studio crisis agent globally across AstraZeneca.
- Mentored a non-technical colleague to independently build a globally deployed Power Application; built a Power Platform application at **>90% first-deployment adoption across three continents** as the standardised gateway for critical-process data.

Resilience Domain Grounding

- All AI work is grounded in leading an ISO 22301 / 22361-aligned resilience and crisis-management programme across global laboratories and clinical manufacturing — the domain expertise and governance fluency that make the AI deliverables credible in a regulated environment.

Business Continuity Program Manager

2022 – 2023

Financial Industry Regulatory Authority (FINRA), Corporate Risk Management · Rockville, MD

Led the overhaul of business-continuity guidance and crisis-response protocols for U.S. financial-market operations housing data for 100M+ accounts — building the regulated-environment and risk fluency the later AI-governance work draws on.

- Stood up an **OSINT threat-monitoring feed** and an ArcGIS people-accountability dashboard; ran vendor scans and API-fit evaluations for third-party risk SaaS — early data- and tooling-driven decision support.
- Directed improvements to the crisis-management framework and designed executive exercises (active-shooter, cyber-ransom, prolonged-power-failure), translating outcomes into board-level proposals; authored the enterprise risk library and ransomware response plan.
- Conducted black-swan market-risk assessments before the Russian invasion of Ukraine, including NATO-territory escalation and cyber-risk exposure analysis.

Senior Resilience Engineer · General Engineer · Emergency Support Specialist

2012 – 2022

U.S. DOT — Pipeline & Hazardous Materials Safety Administration (PHMSA) · Washington, DC

A decade leading critical-infrastructure crisis and resilience work while turning manual federal processes into working software — the technical and crisis foundation the AI work builds on.

- Designed and built a **SharePoint workflow** that digitised the federal oil-spill-response-plan review process (49 CFR Part 194): a three-stage routing engine adopted for 99% of national reviews and used for five years — a direct precursor to today's low-code and agentic-workflow work.
- Pioneered **predictive geospatial (ArcGIS) modelling**, fusing National Weather Service forecasts with pipeline-infrastructure data to pre-position federal resources; built a real-time COVID-19 operational dashboard and coordinated FEMA PPE distribution nationwide.

- Incident Commander on national activations — **Colonial Pipeline** cyber breach, **Merrimack Valley** gas explosions, **Aliso Canyon** and **Hurricane Maria** — briefing up to White House level; Cyber-OT/SCADA subject-matter expert and TSA liaison.
- Co-authored legislation with digital-operational-resilience provisions **mirroring DORA**, and participated in **DARPA RADICS** Exercise 6 (black-start grid recovery after a simulated cyber-attack).

AI APPLICATION PORTFOLIO

Built with multi-agent pipelines (Google Gemini, Mistral AI, Anthropic Claude) and AI-assisted development environments. Tools built within AstraZeneca are described by capability; independently built products are named.

- **Enterprise AI governance & risk-assessment platform** (AstraZeneca, PoC in assessment) — tiered EU AI Act classification, multi-agent assessment, human-in-the-loop, full audit trail, drift monitoring.
- **AI grant & incentive search platform** (AstraZeneca, with the Enterprise AI team) — multi-agent web search with a no-code feedback layer; built in 24h, demoed to CFO.
- **Crisis & continuity exercise builder** (AstraZeneca, approved internal PoC) — multi-LLM scenario, document and PowerPoint generation.
- **OpRes Horizon** — independently built public regulatory-intelligence platform across 90+ countries and 8 sectors; 3-agent / 3-stage pipeline, 0–4 intensity ranking, bilingual EN/FR. Live at opres.structuredresilience.app.
- **ReguCheck Nexus** — independently built AI-compliance platform (EU AI Act, GDPR, DORA).
- **Horizon ClinOps** — independently built clinical-trial supply-intelligence platform, audit-first multi-agent architecture.
- **Legal Horizon** — multi-jurisdiction legal & regulatory horizon-scanning. **Crisis Compass** — AI crisis-management hub with jurisdictional reporting logic. **Monde Français** — CEFR-tracked French-learning ecosystem.

EDUCATION & CREDENTIALS

- M.S., Engineering Management (Crisis, Emergency & Risk Management) — The George Washington University
- B.S., Mechanical Engineering — University of Maryland, College Park
- Fellow of the Business Continuity Institute (FBCI) — highest grade of BCI membership

TECHNICAL PROFICIENCIES

AI & Development: Multi-agent LLM orchestration (Gemini, Mistral, Claude), Microsoft Copilot Studio, AI-assisted development (Claude Code, Google AI Studio, Mistral Vibe CLI), prompt engineering, React/TypeScript, Node/Express, RAG pipelines, Power Platform (PowerApps, Power Automate, Power BI, PowerFx)

AI Governance & Regulation: EU AI Act, GDPR, DORA, NIS2, NIST AI RMF, ISO 42001 (awareness); risk classification, model-drift & bias monitoring, audit-first / human-in-the-loop architectures

Data & Domain: ArcGIS (geospatial dashboards, predictive models), Power BI, Excel (advanced); ISO 22301 / 22361 operational resilience & crisis management

Languages: English (native), French (intermediate, actively developing). Dual U.S. / French citizen.