AI Proof-of-Concept Portfolio

Turning AI ideas into working, deployed tools for real business audiences

Nathan Schoenkin, FBCI · nathan@schoenkin.fr

EXECUTIVE SUMMARY

The hardest part of an AI strategy is rarely the idea. It is turning ideas into something people can actually use, see and judge. This brief shows that I do exactly that. Over a short period I have conceived and built more than eight distinct, working AI applications — each for a different professional audience and a different business problem, most deployed live on a server rather than left as slides or notebooks, and four now in active in-house development.

I am candid about the method: these tools were produced using AI-assisted development environments, not hand-written as a traditional software engineer — and that is the point. The valuable skill is taking a vague business need, framing it as a concrete use case, orchestrating modern AI tooling into a functioning proof-of-concept, and putting it in front of stakeholders quickly. That is the experimentation-and-scaling loop that moves an AI roadmap from ambition to results: collecting needs from business teams, prioritising them, prototyping fast, and creating tangible artefacts that drive decisions and adoption.

PORTFOLIO AT A GLANCE

Most of these tools share a common, reusable foundation: a clean web interface, document handling, and a multi-agent AI engine (Google Gemini, Mistral AI, Anthropic Claude) that orchestrates specialised analysis steps, with results exportable into everyday business formats such as PDF and PowerPoint. That shared pattern is itself the asset — it is what makes spinning up the next use case fast and low-cost. Tools built within AstraZeneca are described below by capability; independently built products are named.

BUILT WITHIN ASTRAZENECA — ENTERPRISE TOOLS

1. Enterprise AI governance & assessment platform

Built with the AstraZeneca Enterprise AI team · proof-of-concept, in assessment

Audience: risk, governance and technology leaders.

Problem: every AI use case needs proportionate scrutiny against the EU AI Act and internal policy — but assessing each one consistently, and proving you did, is slow and uneven.

What it does: a deterministic intake classifies each AI submission and routes it to a proportionate assessment tier; a multi-agent pipeline then scores it against the regulation and internal frameworks, with human reviewers keeping final authority and every agent call and decision logged and exportable for audit. A companion drift-analysis pipeline monitors live multi-agent systems after deployment.

Why it matters: this is responsible AI as an operating capability, not a slide — the governance plumbing a regulated group needs to scale AI safely. It is a proof-of-concept, currently in assessment and iteration with the relevant teams — not yet approved for use.

2. AI grant & incentive search

Co-built in 24 hours with the Enterprise AI team · demoed to CFO level

Audience: finance, tax and AI-programme leaders.

Problem: public grants and tax credits that could fund AI work are scattered across jurisdictions and rarely surfaced in time.

What it does: a multi-agent web-search system autonomously finds grants and tax credits matching internal AI project profiles. Its standout feature is a no-code admin layer — reviewers submit natural-language feedback that becomes rules stored in a RAG database and consulted on every subsequent run, refining behaviour without code releases.

Why it matters: built in a day and demoed to CFO level — proof that a tightly scoped need can become a working, governed tool almost immediately, with a feedback pattern portable to any review-driven process.

3. Crisis & continuity exercise builder

Approved internal proof-of-concept

Audience: risk, business-continuity and training teams.

Problem: designing a credible crisis or continuity exercise — scenario, phases, injects, facilitation materials — takes significant expert time.

What it does: a multi-LLM pipeline generates an exercise phase by phase and exports ready-to-run packs including PowerPoint, with a custom agent monitoring reputable news to ground scenarios in real-world events. Its version history reflects repeated, hands-on refinement.

Why it matters: demonstrates iteration and a process-improvement mindset, and the discipline to take a tool from idea through formal approval toward internal production.

INDEPENDENTLY BUILT — NAMED PRODUCTS

4. OpRes Horizon — operational-resilience & crisis-management regulatory intelligence

Independently built · live at opres.structuredresilience.app

Audience: operational-resilience, business-continuity, crisis-management and risk leaders operating across borders.

Problem: the operational-resilience and crisis-management obligations that apply to a business shift constantly across dozens of jurisdictions and sectors — knowing what the rules are, who issues them and how heavily a country regulates relative to its peers is a permanent, manual challenge.

What it does: a public regulatory-intelligence platform that maps operational-resilience and crisis-management regulation across 90+ countries and 8 industry sectors. A 3-agent, 3-stage pipeline (Mistral search agents triage every sector, then deep-dive the most regulated; Claude synthesises) produces board-level country profiles — a 0–4 regulatory-intensity score, an executive landscape narrative, and separate per-sector OpRes and crisis-management briefings — plus a live "Most Regulated" league table ranking all scanned jurisdictions. Every profile is fully bilingual (English / French) and exportable to Word. A read-only public portal is free with no login; an authenticated admin portal runs the scans.

Why it matters: this is the tool that sits squarely on my own domain — proof that the same rapid PoC method produces something genuinely useful to the operational-resilience and crisis-management community, deployed live and given away free.

5. ReguCheck Nexus — AI governance & compliance

Independently built

Audience: risk, compliance and governance leads up to CCO level.

Problem: organisations must show their AI and processes meet EU AI Act, GDPR and DORA expectations — laborious and inconsistent to assess.

What it does: guides a structured assessment, scores an organisation or solution against the requirements, surfaces gaps, and produces a dashboard and exportable executive report. RAG grounds reasoning in the authoritative regulatory text, so new guidance is operationalised by document upload rather than retraining.

Why it matters: shows I treat AI governance as a deliverable — and that I can ship and host a compliant product end-to-end.

6. Legal Horizon — legal & regulatory horizon-scanning

Audience: legal and compliance teams.

Problem: keeping up with changing legislation and case law across multiple jurisdictions is slow, manual and easy to fall behind on.

What it does: specialised AI analysis steps monitor new legislation and legal developments and turn them into structured, readable summaries — from a quick topline to a deeper dive, organised by country, processing many items in batches.

Why it matters: a clear analogue of the multi-jurisdiction risk-and-compliance exposure any group operating across borders lives with.

7. Horizon ClinOps — clinical-trial operations support

Audience: clinical operations teams.

Problem: running clinical trials means triaging information, tracking many criteria and synthesising scattered research — detailed, high-stakes, time-consuming work.

What it does: a multi-agent, audit-first platform scans trials, assesses impact and supports the operational work around a study; no AI output triggers automated action — every decision routes through human review.

Why it matters: proof I can take a complex, criteria-heavy problem in a highly regulated domain and make it usable and safe.

8. Crisis Compass — AI crisis management

Audience: crisis managers and senior leadership.

Problem: in a live crisis, decision-makers are overwhelmed with information under acute time pressure.

What it does: structures the response — a verified horizon feed, bespoke AI-generated response plans, jurisdiction-specific reporting obligations injected as plan steps, and situation-report and after-action generation.

Why it matters: designing AI for senior, high-pressure audiences where clarity and trust are everything.

Beyond these, the same approach has produced consumer-grade products such as Monde Français, a CEFR-tracked French-learning ecosystem — evidence the method delivers polished tools well beyond the enterprise.

MY APPROACH TO POC-TO-PRODUCTION

What this portfolio illustrates is a way of working, not a one-off result. I start from a specific audience and a real problem, frame it as a tightly scoped use case, and use AI-assisted tooling to produce a functioning proof-of-concept quickly and cheaply. Because the tools share a common foundation, each new idea costs less to test than the last.

Promising PoCs get deployed live so stakeholders can actually use them; the rest are cheap, fast lessons. This is the experimentation engine an organisation needs to convert AI ambition into adoption: many shots on goal, quick feedback, and a clear path for the winners toward production — with success measured by real adoption rather than by a slide.

IN SUMMARY

Multiple audiences. Working AI tools. Live deployments. One repeatable method for turning business needs into usable, governed AI. I would welcome the chance to walk a team through any of these tools live and to discuss how the same approach could be applied to its highest-value use cases.