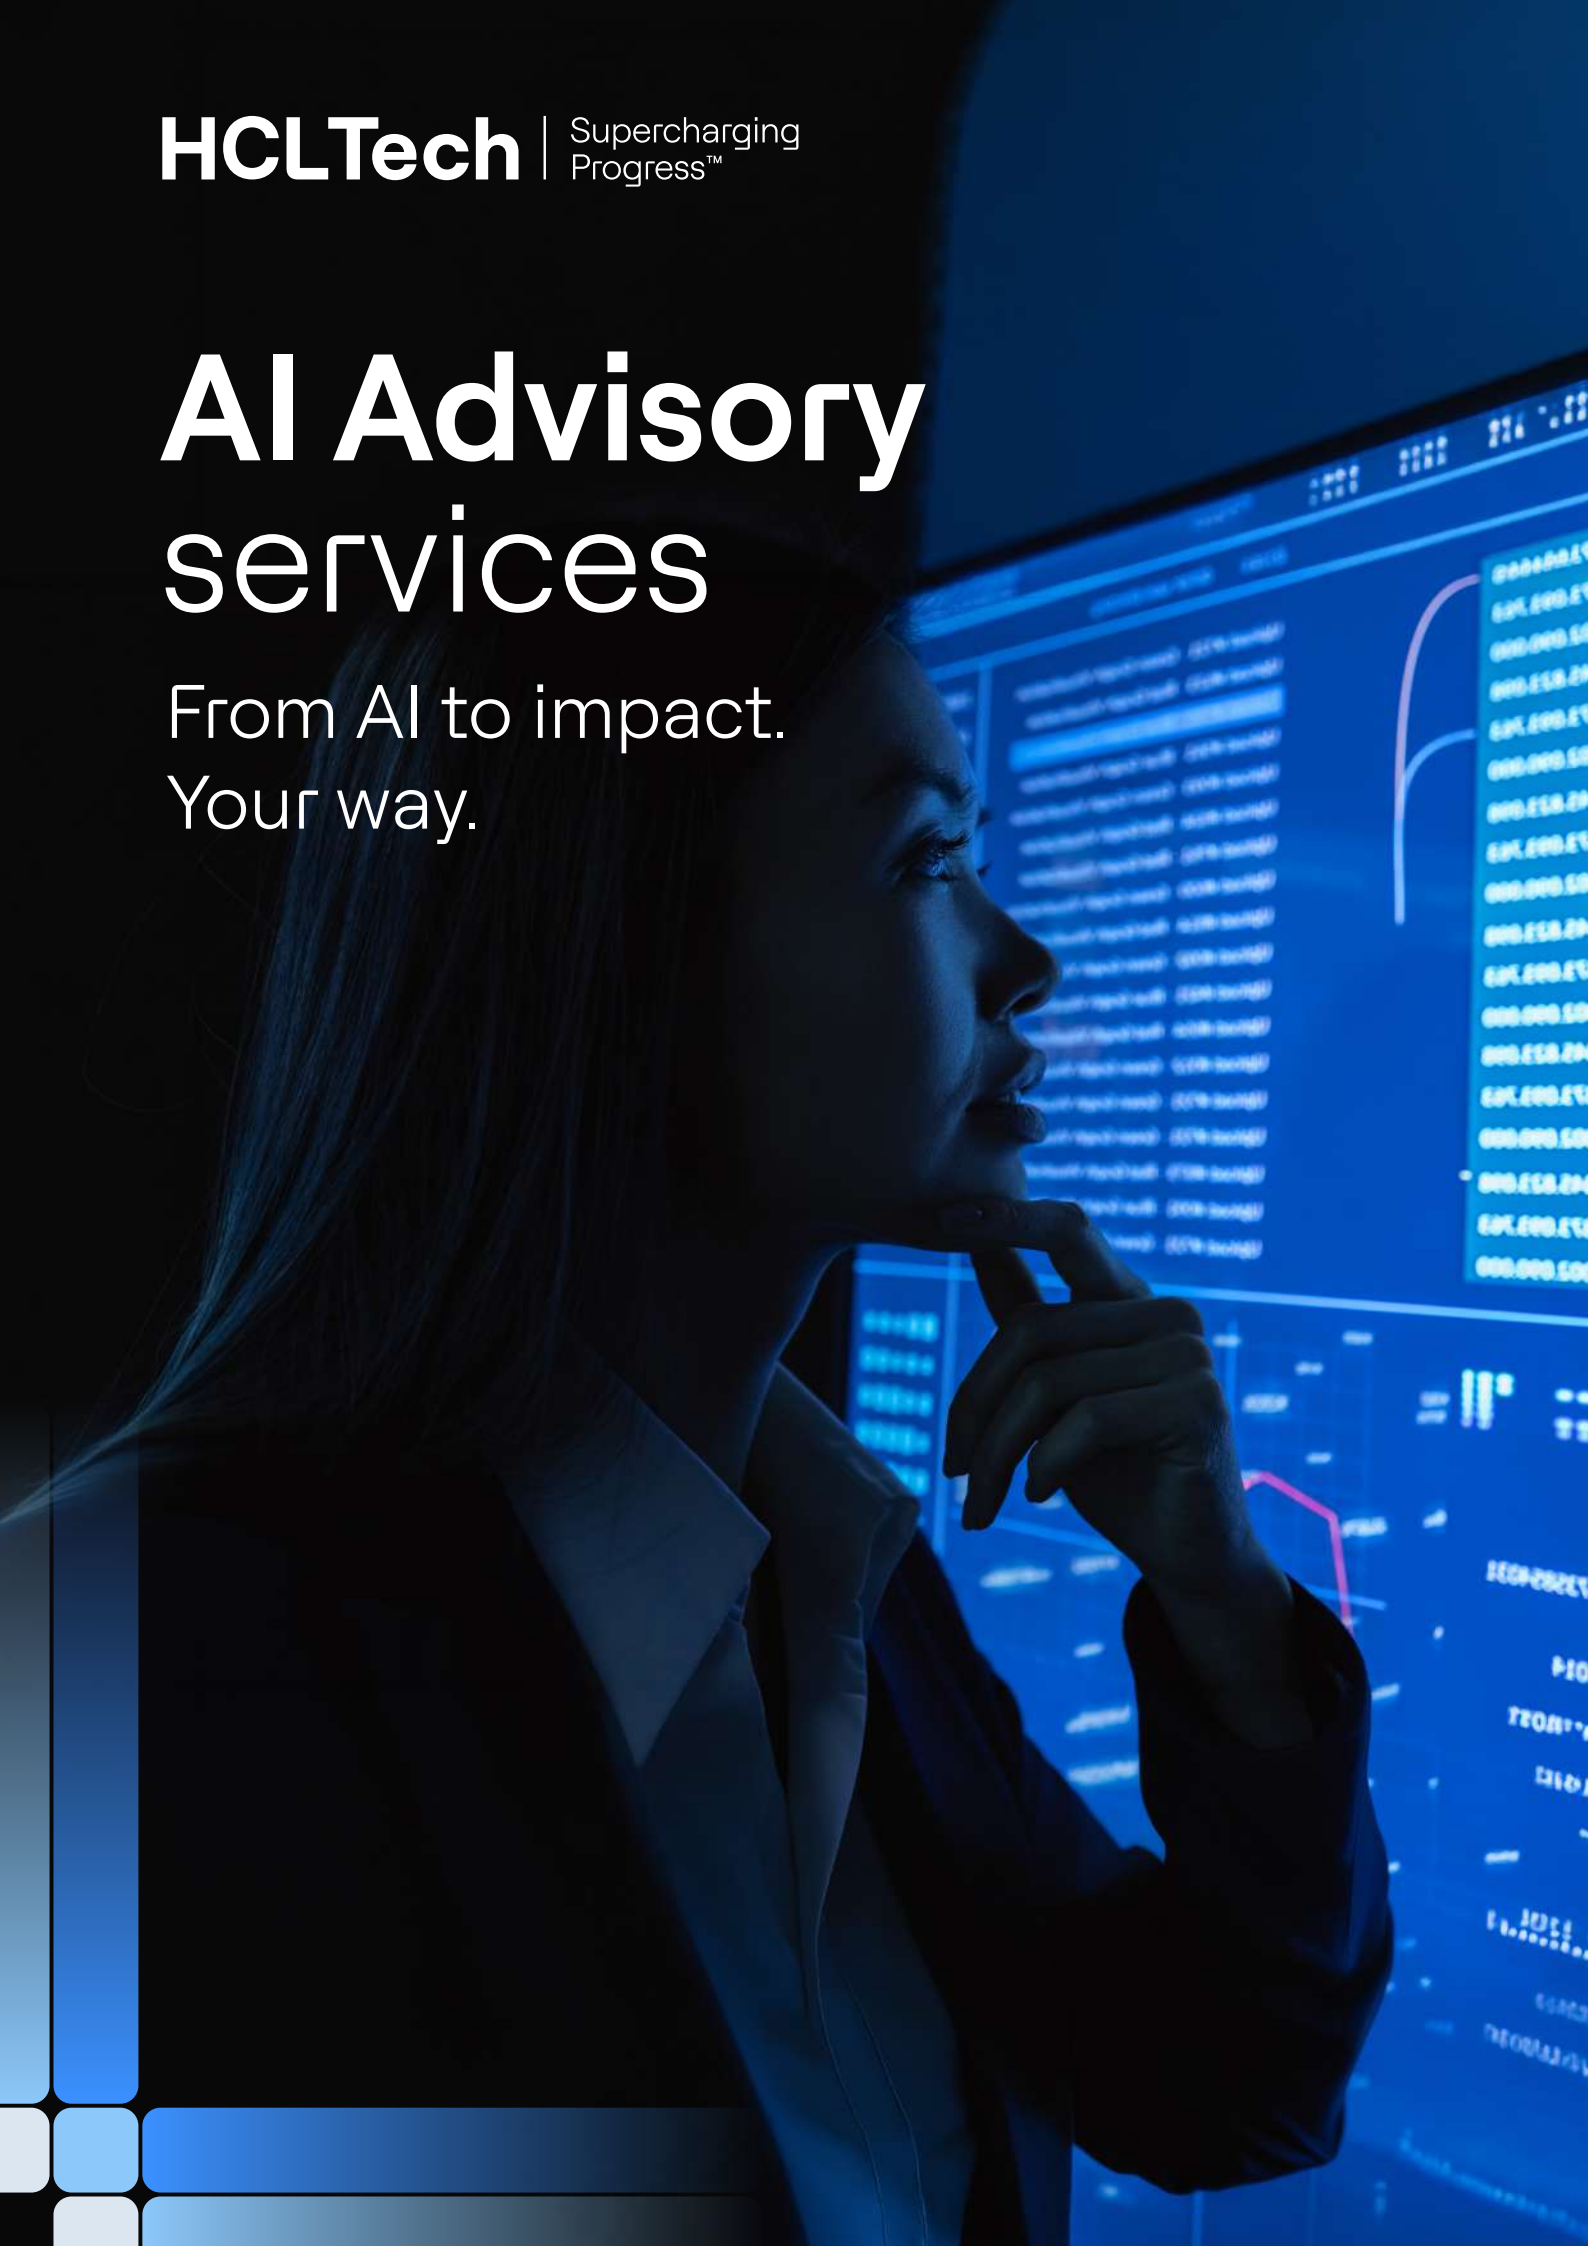
HCLTech | Supercharging Progress™

AI Advisory services

From AI to impact.
Your way.



What AI Advisory is

HCLTech AI Advisory is an enterprise AI consulting service that moves organizations from AI strategy to production using their own data and environment. The offering combines Responsible AI consulting, implementation services and production engineering through a vendor-neutral, fixed-scope and outcome-based model powered by HCLTech AI Labs and the Office of Responsible AI and Governance.

Most enterprises struggle with fragmented AI initiatives, governance frameworks that exist only on paper and innovation labs with no clear path to production. HCLTech AI Advisory solves this through fixed-scope, outcome-driven and AI-accelerated engagements, where teams design the architecture, build the first production system, establish governance controls and transfer full capability to the client. Every engagement is delivered through AI-native squads using agentic tools by default.





Research, analysis, documentation and quality assurance are all agent-assisted, enabling materially faster delivery at a lower cost than traditional advisory models.

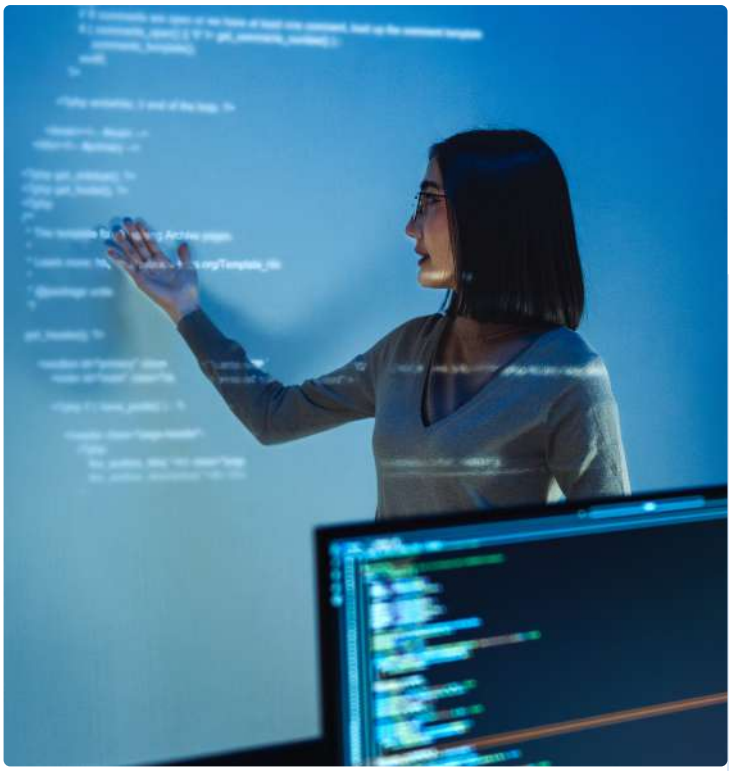
6 Offerings	**6** AI labs	**500+** Clients
1,000+ Engagements	**ISO/IEC 42001** Certified	**275+** Governance use cases

HCLTech has been recertified by the Responsible AI Institute (RAII) for a second consecutive year, earning the Organizational AI Governance Verification Badge for demonstrating alignment with 81 AI governance controls across five governance domains. HCLTech is the first organization globally to achieve this recognition.

Define AI strategy, adoption roadmaps	Build proofs of concept and execute minimum viable product scoping	Build governance, ethics and risk frameworks	Prioritize business use cases and define the investment case for AI	Run AI Red Teaming as a continuous engagement across the AI lifecycle	Design AI operating models and the Office of AI	Provide advisory coverage through sustained operations post-production

Our AI Advisory offerings

Productized. Repeatable. AI-native delivery. Clients can enter at any phase.



Journey to Production

What this is

An advisory-led, time-boxed engagement that demonstrates how AI can work with your real data in your environment. We provide an architecture blueprint with governance embedded, apply Responsible AI controls throughout and close with a working system, documentation and a return on investment baseline. The engagement ends with a clean handover.

What is included

- AI readiness and data architecture review
- Proof of concept and minimum viable product scoping with acceptance criteria agreed upfront
- Architecture blueprint with sign-off
- Advisory-led build with agentic engineering tools and governance embedded
- Return on investment baseline and capability handover pack

Typical scope

Minimum viable product in 1 to 3 months

Pilot in 3 to 12 months

What the client gets

A working production system deployed in your own environment. Full documentation, continuous integration and continuous deployment pipelines, a return on investment measurement dashboard and a capability handover pack. You own what we build. Many clients return for the next use case within three to six months, turning a finite engagement into a strong client partnership.



Agentic Engineering

What this is

Advisory to co-design and support agentic engineering capabilities with your technology teams. Reference architecture, toolchain choices, cross-functional team model design and hands-on knowledge transfer.

What is included

- Reference architecture design and technical review
- Vendor-neutral tool and platform assessment
- Agentic pod model and role design
- Toolchain implementation and adoption plan
- Post engagement coaching and advisory support available

Typical scope	What the client gets
1 to 5 days advisory, extending to on going advisory support	Reference architecture document, toolchain implementation plan, pod model definition and capability roadmap. Your teams can build and run agentic systems independently.



AI Transformation Advisory

What this is

Clients need to know what work agents should do, how agents sit alongside human staff and what governance looks like at scale. AI Transformation Advisory produces an operating model and transition plan to get there, built around your specific workforce and regulatory context.

What is included

- Digital workforce strategy and use case prioritization
- Hybrid workforce operating model design
- Agent persona and workflow design
- Workforce transition and change management plan
- Governance framework for autonomous agent

Typical scope	What the client gets
6 to 16 weeks program	A workforce transition plan with agent designs, a hybrid operating model and a governance framework for autonomous behavior. Board and regulatory ready



Office of AI

What this is

Most organizations have pockets of AI activity with no central coordination. Office of AI designs and implements the governance function that manages priorities, investment decisions and delivery standards across the business, built for the scale you are heading toward, not just where you are today.

What is included

- AI governance structure and operating model
- Center of Excellence charter, scope and team design
- Decision rights framework and escalation model
- AI investment prioritization process
- Metrics, reporting cadence and board-level narrative

Typical scope	What the client gets
8 to 20 weeks program	A fully designed Office of AI including governance structure, operating model, decision rights, investment process and board narrative. Ready to launch and operate.



AI Red Teaming

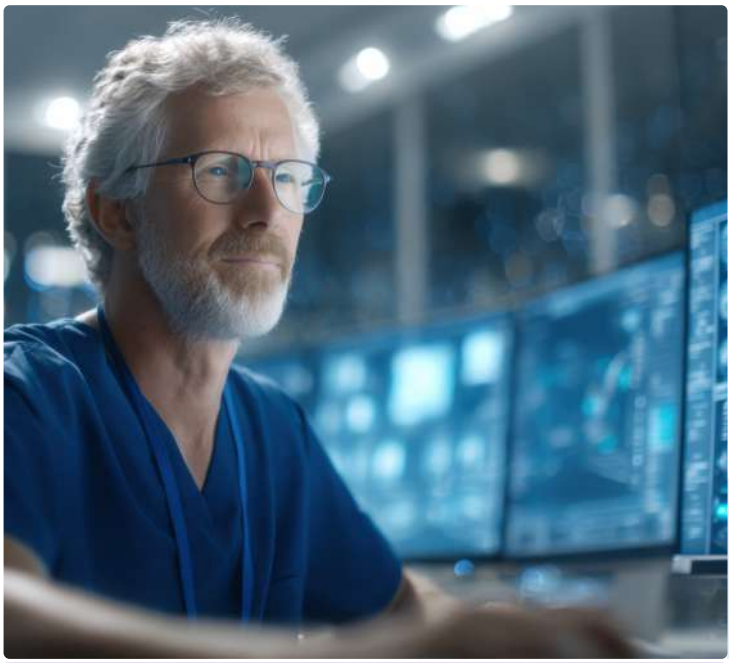
What this is

Structured adversarial and real-world testing of AI systems to find bias, safety gaps, prompt injection risks and unexpected outputs under real conditions. Tests are based on your actual deployment scenarios, not generic checklists.

What is included

- Real-world, adversarial prompt and jailbreak testing
- Bias and fairness assessment across demographic groups
- Data leakage and personally identifiable information exposure checks
- Output consistency and reliability testing
- Findings report with risk ratings and remediation steps

Typical scope	What the client gets
4 to 6 weeks engagement per system	An AI Red Teaming findings report with rated risks, details on how each issue was identified and a prioritized remediation plan. Suitable for board, regulatory or audit committee review.



Responsible AI Advisory

What this is

Responsible AI frameworks matched to your regulatory environment, risk tolerance and AI maturity. Covers the full cycle from policy drafting to the review process embedded in delivery. Practical policy, not just documentation.

What is included

- AI ethics policy and principles design
- Regulatory compliance mapping, for example the EU AI Act and industry-specific rules
- AI risk taxonomy and risk register
- Responsible AI review process integrated into delivery
- Board-level AI accountability narrative

Typical scope	What the client gets
6 to 16 weeks program	A complete Responsible AI framework including ethics policy, risk register, compliance map, embedded review process and board narrative. Audit-ready from day one.

Impact

Every engagement is delivered through AI-native squads using agentic tools by default. Research, analysis, documentation and quality assurance are all agent-assisted, enabling materially faster delivery at a lower cost than traditional advisory.

Up to 2x faster delivery. Up to 50% lower cost.

Metrics	Traditional advisory	HCLTech AI Advisory
Innovation sprint	2 to 4 weeks	1 to 5 days
Minimum viable product delivery	4 to 6 months	1 to 6 weeks
Pod size	5 to 7 people	2 to 3 people



The AI Advisory journey

A productized journey from first conversation to AI at scale. Clients can enter at any phase. Responsible AI Advisory and AI Red Teaming run continuously across every phase.

Assess & Align Day 0 to 30. AI Readiness.	**Govern & Architect** Day 30 to 60. AI Governance.	**Build & Accelerate** Day 60 to 90. AI Innovation.	**Scale & Operate** Day 90 and beyond. AI Scale.
• Baseline AI maturity across 6 dimensions • Map risk and regulatory exposure • Prioritize use cases by impact and feasibility	• Responsible AI framework and policy design • Target-state architecture and platform blueprint • Office of AI operating model	• Fixed-scope minimum viable Product with real data in client environment • Agentic workflows replacing manual processes • Return on investment baseline and handover documentation	• Enterprise agentic workforce design • Continuous performance governance and key performance Indicators • Intellectual property reuse and value measurement at scale

How it works: Agentic delivery powered by HCLTech AI Force

Every AI Advisory engagement is delivered using HCLTech AI Force as the underlying Generative AI and Agentic AI platform. Specialized agents accelerate research, analysis, architecture, governance and documentation, while human experts direct, validate and own the client relationship.

Human access
Users can access AI Advisory services through Microsoft Teams, Slack, email, APIs and a live dashboard.



Orchestration layer

- **Delivery lead agent:** Scope, timeline, client communications
- **Scrum master agent:** Sprint backlog and blockers
- **Knowledge curator:** Frameworks and Intellectual property library

Advisory specialization agents

- **Research agent:** Regulation and market data
- **Analysis agent:** Gap analysis and maturity
- **Architecture agent:** Reference designs
- **Red team agent:** Adversarial testing
- **Governance agent:** Policy and Responsible AI controls
- **Documentation agent:** Reports and frameworks

Knowledge base	Outputs
Framework library. Regulation database. Past engagements. Data vault.	Strategy reports. Governance packs. Risk assessments. Client roadmaps



Why enterprises choose HCLTech AI Advisory

Advice ends with a working system.
Every Journey to Production engagement closes with production-grade code using your data in your environment and a clean handover, not a document about what to build next.

No platform allegiance.
We have no commercial tie to any AI vendor. Our job is to find the best fit for your data, your regulations and your team. A rare position and one that matters most in regulated sectors.

AI-accelerated services.
Every engagement is delivered using optimized squads and extensive AI tools. Research, analysis, documentation and quality assurance are all AI-native and agent-assisted. Up to 2x faster and up to 50% lower cost than traditional advisory. Outcomes, not time and materials.

Outcomes, not time and materials.
Every engagement ends with a concrete output: a roadmap, governance framework, architecture blueprint or working system. Fixed scope. Fixed outcomes.

Governance baked in.
ISO/IEC 42001 certified and the first organization globally to receive the Responsible AI Institute's Organizational AI Governance Verification Badge, verifying 81 governance controls across five domains.

Built for the regulated enterprise.
EU AI Act, NIST AI RMF, industry-specific rules and sovereign AI requirements are built into the advisory blueprint from the start.

Case studies

Case study 1: Driving Responsible AI adoption with 88% faster developer productivity.

Offerings: Agentic Engineering, Responsible AI Advisory

Challenges: Inconsistent Copilot output due to generic or poorly structured prompts. Limited relevance across tech stacks. Lack of standardization. No system for continuous improvement.

Solution: Applied sustainable prompt engineering through a structured, Responsible AI-first methodology, enabling originality, compliance and safety in every prompt. Leveraged the proprietary CRAFT method, which covers Context, Requirements, Architecture, Frameworks and Testing, to guide prompt creation aligned with enterprise standards. Embedded originality safeguards to avoid public code reuse, implementing fairness, transparency and security.

Impact: 88% time savings for developers. $2.8 million annual cost savings. 100% test coverage. 150+ prompts curated.

Case study 2: Establishing Responsible AI governance across 275+ enterprise AI use cases for a leading Middle East banking and financial services enterprise.

Offerings: Responsible AI Advisory, Office of AI

Challenges: Fragmented AI governance. Manual and slow risk controls. Limited explainability and trust. Application programming interface and ecosystem risk. Shariah alignment at scale.

Solution: Human-centric Responsible AI and risk governance framework, purpose-built for the client. Shariah-aligned Responsible AI governance blueprint aligned to personal data protection law, ISO/IEC 42001 and NIST AI RMF. Central AI model registry with mandatory validation and lifecycle visibility across all AI systems. Independent model risk management with bias, explainability, robustness and Shariah impact validation.

Impact: AI portfolio visibility: 275+ enterprise AI use cases consolidated. Enterprise engagement: 82+ leaders aligned on AI strategy and ownership. Strategic engagement: 1,240+ onsite hours with executive and domain teams. Program delivery scale: 5,000+ billable hours delivered across the AI program. AI maturity baseline: 20-dimension AI maturity framework. Governance institutionalization: AI lifecycle governance established from idea to production.

Case study 3: Agentic AI and large language model accelerators for a global pharmaceutical and biotechnology company.

Offerings: Agentic Engineering, Journey to Production

Challenges: The client needed a scalable and governed framework to accelerate development and deployment of Agentic AI and fine-tuned Large Language Model solutions across business units.

Solution: Reusable agent-to-agent and model context protocol accelerator components for Agentic AI solutions. Delivered a low-rank adaptation and Quantized low-rank adaptation based large language model fine-tuning accelerator with MLflow and Unity Catalog governance. Standardized reusable pipelines and orchestration frameworks for AI engineers and data scientists. Enabled rapid validation through REST Application Programming Interfaces, automated Databricks jobs and Gradio User Interface integration. Designed for scalability, governance and enterprise-wide reuse.

Impact: $14.9 million combined annual savings. 3.5+ million annual hours saved through AI automation. Improved productivity for 200+ AI engineers and users. Accelerated deployment

Case study 4: Accelerate ChatGPT enterprise rollout with security operations integration for a global luxury fashion retailer.

Offerings: Journey to Production, Responsible AI Advisory

Challenges: The client required enterprise-grade governance and security visibility before scaling ChatGPT enterprise adoption.

Solution: Integrated ChatGPT enterprise compliance application programming interface logs into existing security operations systems. Enabled centralized audit, monitoring and compliance visibility. Leveraged Amazon Web Services security and ChatGPT Enterprise expertise for rapid deployment.

Impact: Unblocked enterprise-wide ChatGPT rollout. Improved AI governance and security visibility. Reduced deployment timelines from weeks to days. Strengthened compliance readiness for enterprise AI adoption.

HCLTech AI Advisory services

From AI to impact. Your way.

AI Advisory services · Six productized services · AI-native delivery

Visit www.hcltech.com/ai-advisory

HCLTech | Supercharging Progress™

hcltech.com