



Artificial Intelligence

AI Development Service

Overview

We design, build, and integrate AI-enabled services that are safe, effective, and ready to operate in the real world.

CIVIC's AI Development Service supports organisations to move from validated ideas, strategies, and discovery insights into delivery. We build AI components and end-to-end services that are robust, secure, and grounded in real user needs, organisational priorities, and public-sector accountability.

Our work spans a broad range of AI approaches — including Generative AI (GenAI), Large Language Models (LLMs), predictive and statistical models, optimisation and decision engines, and automation. We do not promote a single technology or model. Instead, we select and combine techniques that are appropriate for the problem being addressed, the quality and availability of data, and the environment in which the service must operate.

We focus on production-ready, cloud-based services that integrate cleanly with existing systems and platforms, fit established operating models, and can be run, supported, assured, and evolved with confidence over time. Our emphasis is on building services that work in practice, not just in controlled or experimental settings.

Why AI Development

Many organisations now have AI strategies, early pilots, or proofs of concept, often developed in response to rapid advances in tooling and capability. The challenge is turning this early momentum into services that are reliable, secure, and genuinely useful in day-to-day operations.

Common barriers include moving beyond experimentation, integrating AI safely with existing systems and data, managing quality and security at scale, and ensuring that AI-driven outputs are understandable, trusted, and appropriate for users. In public-sector and regulated environments, these challenges are compounded by the need for transparency, assurance, audit ability, and clear accountability.

AI Development is where many initiatives succeed or fail. Decisions made at this stage — about architecture, data flows, integration, governance, and operating models — determine whether AI becomes a sustainable capability or a fragile dependency.

Our role is to help organisations bridge the gap between intent and operation. We focus on delivery that works not only technically, but organisationally and operationally, so AI-enabled services can be introduced with confidence and continue to perform under real-world conditions.

What the Service Does

Our AI Development Service supports organisations in designing, building, integrating, and deploying AI-enabled services ready for live use.

In practice, this means we:

- Design and build AI-enabled services based on validated requirements, user needs, and discovery outcomes.

- Select and implement appropriate AI techniques, from automation and predictive models through to GenAI and LLM-based systems, based on fitness for purpose.
- Integrate AI components with existing cloud platforms, data sources, and operational systems, ensuring they function as part of a wider service.
- Put in place safe operating practices, including monitoring, escalation, human oversight, and fallback mechanisms.
- Deliver production-ready services, supported by clear documentation, governance arrangements, and handover to operational teams.

The emphasis throughout is on building services that are usable, maintainable, and resilient — not one-off builds, isolated models, or experimental solutions that cannot be supported over time.

Key Features

- Delivery of AI-enabled services using a range of techniques, including GenAI, LLMs, predictive modelling, optimisation, and automation, applied where they add genuine value.
- Cloud-native implementation across platforms such as AWS, Azure, and GCP, aligned with existing organisational standards and controls.
- Integration with existing systems, APIs, and data pipelines to support complete, end-to-end services rather than stand alone components.
- Secure handling of data, models, prompts, and outputs in line with organisational policies and regulatory requirements.
- Deliberate design of AI behaviour, including tone, constraints, and guardrails for user-facing services.
- Validation, testing, and quality assurance of AI outputs, system performance, and known failure modes.
- Support for deployment, monitoring, logging, and operational readiness.
- Clear documentation and knowledge transfer to support long-term ownership and continuous improvement.

Benefits

Working with CIVIC on AI Development provides:

- AI-enabled services that are fit for purpose and ready for live operation within real organisational environments.
- Reduced implementation risk through proven delivery patterns and early consideration of security, governance, and operational impact.
- Strong alignment between user needs, service design, and technical implementation, reducing rework and unintended consequences.
- Confidence that AI components behave appropriately, transparently, and predictably in real-world use.
- Services that can be supported, monitored, and improved over time without reliance on

fragile, opaque, or overly bespoke solutions.

Our Approach

We deliver AI development using an iterative, service-focused approach that combines software engineering, data science, and service design. Work is structured to provide regular checkpoints, transparency, and opportunities for assurance, ensuring progress is visible and risks are identified and addressed early.

1. Service and Technical Design

We begin by translating discovery outcomes, requirements, and constraints into a clear service and technical design.

This phase typically includes:

- Refining service scope, success measures, and non-functional requirements such as security, performance, resilience, and accessibility.
- Selecting appropriate AI techniques and architectural patterns that align with organisational standards.
- Designing data flows, integrations, access controls, and dependencies.
- Defining AI behaviour, constraints, and guardrails, particularly for user-facing or decision-support services.
- Agreeing on delivery milestones, dependencies, assurance activities, and decision points.

The Output a clear, agreed service and technical design that provides a sound and auditable foundation for implementation.

2. Build and Integration

We build and integrate AI components into a wider digital service, working incrementally to manage risk and maintain transparency.

Activities in this phase may include:

- Developing AI models, prompts, workflows, or automation components.
- Implementing APIs, services, and integrations with existing platforms.
- Connecting to data sources and pipelines, including data validation and quality checks.
- Applying security, privacy, and access controls throughout the service.
- Iterative testing with representative data and users to validate behaviour and performance.

Delivery is incremental, allowing assumptions, risks, and trade-offs to be reviewed and adjusted as work progresses.

3. Validation, Deployment, and Handover

We ensure services are ready to operate safely and effectively in live environments.

This phase focuses on:

- Validating AI outputs, behaviour, and performance against agreed criteria and user needs.
- Testing for robustness, bias, edge cases, and known failure modes.

- Deploying services to cloud environments using agreed release and change processes.
- Establishing monitoring, logging, alerting, and escalation processes.
- Producing documentation and supporting structured handover to operational teams.

The result is a service that can be run, supported, assured, and improved with confidence.

Team

AI Discovery works best with a small, senior team that combines business insight with technical depth. Depending on the scope, this typically includes:

- A product or business specialist
- A GenAI or machine-learning engineer
- A security or risk specialist (where appropriate)

The exact team structure is tailored to your context and regulatory environment.

Experience

We have delivered AI-enabled services across a range of practical contexts, including:

- Conversational and chat bot services integrated with live operational systems.
- Decision-support and co-pilot tools for frontline and operational teams.
- Predictive and analytical services supporting planning, forecasting, and risk management.
- Document and text analysis services using LLMs and supporting models.

In each case, the focus has been on reliability, transparency, and integration into real services — not experimentation for its own sake.

Public Sector Context

We have extensive experience delivering digital and cloud services in complex public-sector environments. Our work aligns with GDS Service Standards, security accreditation processes, and data-protection obligations.

We design AI-enabled services to work effectively within multi-supplier landscapes, reinforce clear accountability, and support confident decision-making at both delivery and governance levels.

Security, Ethics, and Assurance

Security and responsible use of AI are integral to our development work and are considered from the earliest stages of delivery. They shape not only the technical design of AI-enabled services, but also how those services behave in practice and how users experience them.

We pay close attention to how AI presents itself — including its conversational tone, the clarity of its outputs, and the trust it builds with users — alongside the guardrails and controls that govern its use within live services.

This includes:

- Data protection and information security by design, embedded throughout service architecture and delivery.
- Secure-by-design implementation of AI-enabled services, including protection of data,

prompts, model interactions, and outputs.

- Consideration of model risk, bias, explainability, transparency, and appropriate use, including how decisions and outputs are communicated.
- Definition of appropriate AI behaviours, prompts, and responses, particularly for chatbots and conversational tools.
- Clear guardrails covering escalation, human oversight, fallback, and acceptable use.
- Consideration of operational, service, and reputational impact as part of day-to-day operation.
- Alignment with organisational governance, assurance, and accountability processes.

This approach helps ensure that AI-enabled services are trustworthy, defensible, and suitable for use in regulated and public-facing environments, and that any move into live operation is confident, responsible, and sustainable over time.

Building Capability

Building capability is a core part of our delivery approach, not a by-product.

Throughout development, we work openly with client teams, sharing knowledge, explaining design decisions, and making our thinking visible. This builds confidence not just in the service itself, but in the ability to operate, assure, and evolve it over time.

Our aim is not dependency on CIVIC, but capability within your organisation — enabling teams to manage AI as part of a wider service, data, and technology landscape with confidence.



Contact

CIVIC

12 South Charlotte Street
Edinburgh
EH2 4AX
T: 0131 624 9830
E: info@civicuk.com

www.civicuk.com
Registered in Scotland SC221925
VAT Registration No 774906100
CIVIC Computing Ltd

