

AI & AGENTIC AI PILOT DESIGN AND DELIVERY SERVICES

Service Definition
G-Cloud 15

REPLY

38 Grosvenor Gardens - London SW1W 0EB - United Kingdom - tel +44 (0) 207 730 6000 - fax +44 (0) 207 259 8600
info@reply.com

www.reply.com



TABLE of CONTENTS

SERVICE DEFINITION	1
TABLE OF CONTENTS	2
SERVICE DESCRIPTION	3
ONBOARDING	9
KICK-OFF MEETING	9
ACCOUNT MANAGEMENT	9
<i>DESIGNATED ACCOUNT MANAGER</i>	9
<i>FORTNIGHTLY REVIEW MEETING</i>	9
OFFBOARDING	9
ACCESS TO DATA	9
SECURITY	10
PERSONNEL SECURITY	10
INFORMATION SECURITY	10
PHYSICAL SECURITY	10
TRAINING & KNOWLEDGE TRANSFER	11
ORDERING AND INVOICING PROCESS	12
CUSTOMER RESPONSIBILITIES	13
TECHNICAL AND CLIENT-SIDE REQUIREMENTS	13
OUTCOMES AND DELIVERABLES	13
ABOUT REPLY	14



SERVICE DESCRIPTION

Many public sector organisations recognise the potential of AI and agentic AI but struggle to move beyond experimentation in a safe, credible and evidence-based way. Pilots are often either too abstract to prove real value, or too tightly coupled to live services, introducing unacceptable operational, security or reputational risk. As a result, organisations are left with inconclusive proofs of concept and limited confidence in whether, or how, AI should be scaled.

Reply's AI & Agentic AI Pilot Design and Delivery Services for G-Cloud 15 address this challenge by designing and delivering controlled, well-governed AI pilots that balance rapid learning with strong operational safeguards. We enable organisations to explore AI and agentic AI capabilities in realistic conditions, while maintaining clear boundaries around risk, data access and decision-making authority.

Our service focuses on learning and evidence rather than technology demonstration alone. Pilots are scoped around clear objectives, success criteria and evaluation methods, ensuring outcomes directly inform strategic decisions about scaling, further investment or refocusing effort. Where agentic AI is involved, pilots are designed with appropriate human oversight, auditability and integration controls from the outset.

Crucially, we treat pilots as stepping stones, not dead ends. Each pilot produces actionable insight, documented lessons learned and clear recommendations that enable organisations to progress confidently to production-ready AI capabilities or to stop safely where value is not demonstrated.

FEATURES

Our service combines rapid delivery with disciplined governance and public sector delivery experience to ensure AI pilots are safe, credible and decision-focused.

- **AI and Agentic AI Pilot Scoping**

We work with stakeholders to define clear pilot objectives, hypotheses and scope. This ensures pilots answer specific questions, such as feasibility, value, risk or operational impact, rather than acting as unfocused technology experiments.

Scoping includes defining boundaries, data sources, success measures and acceptable risk levels.

- **Rapid Prototype and Proof-of-Concept Delivery**

We design and deliver prototypes and proofs of concept quickly, enabling rapid feedback and learning cycles. This allows organisations to test assumptions, explore alternative approaches and adapt direction early.

Rapid delivery avoids long-running pilots that lose momentum or relevance.

- **Controlled Environments for AI Testing**

Pilots are delivered in controlled environments that isolate experimental AI capabilities from live services. This reduces operational risk while still allowing meaningful testing of behaviour, performance and integration.

Controls are tailored to the sensitivity and criticality of the use case.

- **Integration With Real Enterprise Systems**



Where appropriate, pilots are integrated with real enterprise systems, APIs and workflows rather than synthetic or disconnected data. This ensures outcomes reflect operational reality and avoids misleading results based on artificial test conditions.

Integration is designed to be non-disruptive and reversible.

- **Data Preparation and Integration Support**

We support data identification, preparation and integration required for pilots, ensuring data is suitable, governed and accessible. This reduces the risk of pilots failing due to data quality or availability issues rather than AI capability.

- **Governance and Assurance for Pilots**

We embed governance and assurance controls appropriate to pilot scope and risk. This includes defining decision rights, review points and escalation paths to ensure pilots remain controlled and defensible.

Governance supports learning rather than constraining it.

- **Human-in-the-Loop Controls**

For agentic or decision-support pilots, we design human-in-the-loop mechanisms that ensure AI outputs are reviewed, approved or challenged as appropriate. This supports safe experimentation while maintaining accountability.

- **Evaluation Metrics and Success Criteria**

We define clear evaluation metrics and success criteria aligned to pilot objectives. These may include accuracy, efficiency, usability, risk reduction or operational impact, depending on the use case.

Clear metrics enable objective assessment rather than subjective judgement.

- **Lessons Learned and Scale Recommendations**

Each pilot concludes with a structured review capturing outcomes, limitations, risks and dependencies. We provide clear recommendations on whether and how to scale, adjust or stop the initiative.

This ensures pilots directly inform future decisions.

- **Knowledge Transfer and Documentation**

We document pilot design, outcomes and technical approaches, transferring knowledge to internal teams. This builds organisational understanding and capability beyond the life of the pilot itself.

BENEFITS

Our AI & Agentic AI Pilot Design and Delivery Services enable organisations to explore AI safely, learn quickly and make evidence-based decisions about adoption and scale. The benefits focus on **reducing**



uncertainty, managing risk and turning experimentation into actionable outcomes, rather than simply delivering technical prototypes.

- **Clear, Focused Objectives for AI Experimentation**

Many AI pilots fail because they are unfocused demonstrations rather than structured experiments. By clearly defining pilot objectives, hypotheses and success criteria upfront, organisations ensure that each pilot answers specific questions about feasibility, value, risk or impact.

This clarity prevents wasted effort and ensures pilots directly inform strategic decisions.

- **Faster Learning Without Committing to Large-Scale Change**

Rapid prototyping and proof-of-concept delivery allow organisations to test ideas quickly, challenge assumptions and learn from real behaviour. Short learning cycles enable teams to explore multiple options without committing to long, expensive programmes.

This accelerates understanding while keeping investment proportionate.

- **Reduced Operational and Reputational Risk**

Pilots are delivered in controlled environments with clear boundaries, ensuring experimental AI does not disrupt live services, expose sensitive data or create unmanaged dependencies. This protects critical public services while still allowing meaningful experimentation.

Risk is deliberately managed rather than accepted by default.

- **Realistic Outcomes Through Enterprise Integration**

By integrating pilots with real enterprise systems, workflows and data where appropriate, organisations gain insight into how AI behaves in realistic operational conditions. This avoids misleading results that often arise from pilots based solely on synthetic data or isolated environments.

Decisions about scale are grounded in reality rather than optimism.

- **Avoidance of Failed AI Initiatives Caused by Poor Data**

Data quality and integration issues are a common cause of failed AI initiatives. By addressing data preparation and integration as part of the pilot, organisations identify data readiness issues early, before significant investment is made.

This prevents technology being blamed for underlying data problems.

- **Faster Understanding of AI Value and Limitations**

Well-designed pilots reveal not only where AI performs well, but also where it does not. Organisations gain a balanced understanding of strengths, limitations, risks and dependencies, enabling more realistic expectations and planning.

This supports mature, informed adoption rather than hype-driven decisions.

- **Controlled, Trusted AI and Agentic AI Behaviour**



Human-in-the-loop controls ensure AI and agentic AI outputs are reviewed, approved or challenged where appropriate. This builds trust with users, leaders and oversight bodies, particularly for sensitive or high-impact use cases.

Trust established during pilots supports wider adoption later.

- **Evidence-Based Decisions About Scaling or Stopping**

Clear evaluation metrics and success criteria provide objective evidence of pilot outcomes.

Organisations can confidently decide whether to scale, refine or stop an initiative based on data rather than anecdote or pressure to proceed.

This reduces the risk of sunk-cost decision-making.

- **Reduced Risk Before Major Investment**

Lessons learned and scale recommendations provide a clear view of what needs to change — technically, operationally or organisationally — before AI is deployed at scale. Risks are identified early, when they are cheaper to address.

Large investments are made with eyes open rather than hope.

- **Growth of Internal AI Capability and Confidence**

Knowledge transfer and documentation ensure learning from pilots remains within the organisation. Teams build confidence in AI concepts, limitations and governance requirements, reducing long-term dependence on external support.

WHAT THE SERVICE INCLUDES

Reply delivers a **complete, end-to-end pilot service** covering design, delivery, evaluation and enablement. The service is designed to turn controlled experimentation into clear direction for future AI adoption.

1. Discovery and Pilot Context Definition

We begin by understanding organisational goals, pressures and constraints driving the interest in AI or agentic AI.

This includes:

- Clarifying what decisions the pilot needs to inform
- Understanding service, policy and operational context
- Identifying acceptable levels of risk and autonomy
- Aligning pilot scope with organisational priorities

This ensures pilots are relevant, purposeful and appropriately bounded.

2. AI and Agentic AI Use-Case Selection and Scoping

We work with stakeholders to identify and prioritise pilot use cases, focusing on those that are suitable for experimentation and learning.

Scoping includes:

- Defining pilot objectives and hypotheses
- Identifying required data, integrations and controls
- Setting boundaries for AI behaviour and access
- Agreeing success criteria and evaluation measures

This avoids pilots becoming uncontrolled or unfocused.



3. Pilot Architecture and Control Design

We design pilot architectures that support rapid iteration while maintaining governance and separation from live services.

This includes:

- Selection of deployment environments
- Integration design with enterprise systems
- Security, access and data-handling controls
- Observability and audit considerations

Architecture is designed to support learning safely.

4. Rapid Build and Prototype Delivery

We design and build prototypes and proofs of concept using appropriate AI, agentic AI and integration technologies. Delivery prioritises speed to insight over production-grade optimisation.

This allows assumptions to be tested early and adjusted quickly.

5. Data Preparation and Integration Enablement

We support identification, preparation and integration of data required for pilots.

Activities include:

- Data sourcing and validation
- Integration with enterprise data and systems
- Governance and access control alignment
- Identification of data quality constraints

This ensures pilots fail or succeed for the right reasons.

6. Controlled Testing and Experimentation

Pilots are run in controlled conditions with monitoring and safeguards in place. Behaviour, performance and user interaction are observed to generate meaningful insight.

Testing focuses on learning rather than perfection.

7. Human-in-the-Loop and Oversight Mechanisms

Where pilots involve decision support or agentic behaviour, we implement human oversight controls.

This includes:

- Approval workflows
- Review and escalation paths
- Separation of recommendation and action
- Capture of human decisions

Oversight mechanisms support trust and accountability.

8. Evaluation Framework and Outcome Assessment

We apply agreed evaluation metrics to assess pilot outcomes objectively.

Evaluation may include:

- Effectiveness and accuracy
- Operational impact
- User trust and usability
- Risk and control effectiveness

Findings are clearly documented and evidence-based.

9. Lessons Learned and Scaling Recommendations

Each pilot concludes with a structured review capturing:

- What worked and what did not



- Constraints and dependencies identified
- Risks uncovered during testing
- Recommendations for scaling, refinement or stopping

This turns experimentation into actionable insight.

10. Documentation and Knowledge Transfer

We provide clear documentation and knowledge transfer to internal teams, ensuring learning persists beyond the pilot.

This includes:

- Pilot design and decisions
- Technical and operational findings
- Governance and control considerations
- Suggested next steps

WHY THIS MATTERS

Together, these services replace unfocused experimentation with structured, low-risk learning. Organisations gain clear evidence about where AI and agentic AI deliver value, where risks lie and what must be in place to scale safely. This enables confident progression from curiosity to capability while protecting public services, data and trust.



ONBOARDING

Kick-Off Meeting

When we start an engagement with a new client, we begin with a workshop introducing the team (from both Reply and the client), the project objectives, and our anticipated timelines and approach. This workshop is most impactful when the majority of client stakeholders involved in the project are present and involved (either in person or remotely). Where this is impossible, we can hold separate sessions with teams that cannot participate in the kick-off workshop to ensure all stakeholders are engaged and bought into the project's objectives.

Account Management

Designated Account Manager

Each client has a designated Account Manager, who will be the main point of contact throughout the engagement. This individual is responsible for ensuring that delivery progresses and is executed to the standards that Reply and our clients expect.

Fortnightly review meeting

Alongside regular ceremonies and informal daily interaction between the Reply project team and the client, we hold formal review meetings with key stakeholders, usually fortnightly. This includes critical client stakeholders, senior Reply project team members, and the Reply Account Manager. Where appropriate, these meetings will consist of a review of work completed and demonstrations or walk-throughs of deliverables.

Offboarding

Where an engagement continues to completion, unless the client requests an extension or new engagement, all client IT, documents, artefacts and deliverables will be returned to the client. Our team will ensure a complete handover of data and knowledge to client staff. An ongoing support or retained advisory service may be implemented where appropriate to provide additional continuity.

Access to Data

Ensuring secure access to client data upon exit is paramount to maintaining confidentiality and integrity. We begin by conducting a comprehensive inventory of all client data stored within our systems, encrypting it both in transit and at rest using industry-standard protocols. Strict access controls and authentication mechanisms are implemented to restrict access to authorized personnel only, with permissions regularly reviewed and updated. Client data is segregated from other data within our systems, ensuring isolation and protection. Upon exit, all access rights to client data are promptly revoked, including user account deactivation and permission removal. If requested, we securely delete or anonymise client data in accordance with retention policies and regulatory requirements. Detailed documentation and audit trails of data access are maintained, providing transparency and accountability. Before exiting, we verify with the client that all necessary data access has been revoked. Finally, we adhere to legal and regulatory obligations regarding data protection and privacy, ensuring compliance with contractual agreements and applicable laws such as GDPR or HIPAA. Through these measures, we ensure that client data remains secure throughout the exit process, mitigating the risk of unauthorized access or data breaches.



SECURITY

Personnel Security

As a Specialist Cloud Service, the capability being offered is not limited to specific Impact levels (as it is not infrastructure, software or a platform) and can be used, subject to personal Security Clearance levels. Reply consultants are mostly Security Cleared (SC), some have higher-level Developed Vetting (DV) clearances. We also have a pool of NPPV3 Consultants for Policing work. The majority of our work for both public and private sector clients is at IL2 but we work in the Official, Secret and Top-Secret domains. Our hard and soft information security processes have been designed and approved by independent CESG CLAS accredited consultants.

Information Security

We are ISO27001, Cyber Essentials certified and our Quality Assurance processes are based upon and compliant with our ISO 9001 accreditation. Our Information Security processes are directly guided by our ISO27001 accreditation. We shall adhere to local information and other security policies and will apply local Security Operating Procedures (SyOPs) as may exist. If such do not exist we shall apply our own SyOPs.

Physical Security

The Main Reply Office (London, UK) is a Police Accredited Secure Facility (PASF) with access controls for each point of entry. All Laptops/Phones must be stored securely overnight. When working on a client-site our consultants adhere to client security policies.



TRAINING & KNOWLEDGE TRANSFER

We enable customer teams throughout delivery by embedding knowledge transfer alongside implementation activities. Training is aligned to organisational roles, ensuring relevance for finance users, operational staff, and system administrators.

Knowledge transfer is delivered through:

- A train-the-trainer approach to establish internal subject matter experts.
- Technical training for administrators covering configuration, security roles, and integrations.
- Walkthroughs of system design, data structures, and operational processes.
- Provision of comprehensive documentation, including user guides and runbooks.

Formal training is supported by ongoing access to expert guidance during delivery and early life support. Where required, refresher sessions or additional enablement can be provided to support organisational change or onboarding of new users.



ORDERING AND INVOICING PROCESS

To begin the process, please contact glue.frameworks@reply.com with details of your organisation, role and high-level requirements.

Our Framework Manager will connect you to our experts and from here we will organise an introductory discussion to go into more detail and shape a proposed engagement that suits your needs.

We invoice for both fixed price and time and materials engagements one month in arrears. For fixed price engagements, this takes place upon completion of the deliverable(s), or in stages, if the size of the engagement is greater than £40,000 excluding VAT.

Payment terms are 30 days net of receipt of invoice.



CUSTOMER RESPONSIBILITIES

To deliver our responsibilities as a partner, we require the following:

- Full access to systems and data that are relevant to the specific project
- Reasonable and timely access to key stakeholders and internal SMEs for information gathering, requirements definition and other critical tasks associated with the delivery of our outputs

Technical and Client-Side Requirements

Other than the system access specified above, there are no specific technical or client-side requirements that would limit our capacity to begin work.

Outcomes and Deliverables

Our engagement's specific outcomes and deliverables will be determined during the scoping phase before beginning the project based on your goals and objectives and the level of depth you require. These will include:

- Workshop(s) to discuss and assess the existing architectural design
- Detailed assessment of design with identification of risks and opportunities
- Report with recommendations and proposed action plan



ABOUT REPLY

Reply is a company that specialises in Consulting, Systems Integration and Digital Services with a focus on the conception, design and implementation of solutions based on the new communication channels and digital media.

Reply partners with key industrial groups in defining and developing business models made possible by the new technological and communication paradigms such as Artificial Intelligence, Big Data, Cloud Computing, Digital Communication, the Internet of Things and Mobile and Social Networking. In so doing, it aims to optimise and integrate processes, applications and devices.

Reply's offer is aimed at fostering the success of its customers through the introduction of innovation along the whole economic digital chain. Given its knowledge of specific solutions and due to a consolidated experience, Reply addresses the main core issues of the various industrial sectors.

Through its network of specialist companies, Reply supports some of Europe's leading industrial groups in Telco & Media, Industry & Services, Banks & Insurance, and Public Administration to define and develop business models, suited to the new paradigms of Artificial Intelligence, Big Data, Cloud Computing, Digital Media and the Internet of Things.

We make innovation happen.

We started with a small team and a purpose: to help the digital revolution happen. Today, we are a team of more than 14,600 people in 16 different countries but we still have the same DNA, made up of agile, vertical task forces and an inner passion for innovation.

We are a decentralised network of specialised companies.

Among Replyers, you will find passionate geeks, visionary strategists, and creative minds, each with sharp skills in a specific business, who cooperate together and never stop learning from each other.

Make forward, act sustainably.

We know a sustainable future is possible and technology can be a strong asset to reach it. As leaders in digital transformation, we push for change and operate in full accordance with the highest ethical standards and with respect for the rights of future generations