REPLY

MCP-BASED AGENTIC AI INTEGRATION SERVICE

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
infouk@reply.com

www.reply.com

TABLE of CONTENTS

SERVICE DESCRIPTION

Public sector organisations are increasingly exploring agentic AI to support complex decision-making, automate actions and improve service delivery. However, without strong integration patterns, governance and security controls, AI agents can quickly become disconnected from enterprise systems, operate with unclear permissions and create unacceptable operational, security and reputational risk. Many early AI initiatives remain experimental because organisations lack a safe, auditable way to connect AI agents to live data, systems and workflows.

Reply's MCP-Based Agentic AI Integration Services for G-Cloud 15 address this challenge by providing a secure, standardised and governed approach to integrating AI agents with enterprise environments using Model Context Protocol (MCP). We enable AI agents to interact with systems, integrations and data in a controlled, policy-driven and auditable way across cloud, hybrid and legacy estates.

Our service embeds agentic AI into existing integration, identity and security architectures, rather than bypassing them. MCP provides a consistent mechanism for managing AI context, permissions and actions, ensuring agents operate within clearly defined boundaries. This allows organisations to move confidently from AI experimentation to production use cases that are transparent, trusted and compliant.

Crucially, we focus on long-term sustainability. Solutions are designed to be vendor-agnostic, extensible and aligned with GDS and public sector security standards, enabling organisations to scale agentic AI adoption while maintaining accountability, explainability and public trust.

FEATURES

Our service combines agentic AI, secure integration patterns, identity-aware access control and governance-by-design to create AI solutions that are safe, scalable and suitable for real-world public sector delivery.

- **MCP-Based AI Agent Integration with Enterprise Systems**
  We design and implement AI agents that integrate with enterprise systems using Model Context Protocol rather than bespoke or hard-coded connections. MCP provides a standard, structured way for agents to request context, invoke actions and receive responses.
  This approach:
  - Reduces fragile, point-to-point AI integrations
  - Enables consistent integration across platforms
  - Simplifies governance and auditability
  - Allows agent behaviour to evolve without re-engineering integrations

- **Secure Context Sharing Between AI Agents and Services**
  We carefully design how context is shared with AI agents, ensuring they receive only the minimum information required to perform a task. Context is scoped, filtered and governed, reducing the risk of data leakage or inappropriate AI inference.

  This prevents scenarios where AI systems operate with unrestricted or poorly understood access to sensitive information.

- **Policy-Driven Agent Permissions and Actions**
  AI agent capabilities are controlled through explicit policies defining what an agent can see, what actions it may take and under what conditions. Policies can vary by environment, use case or risk profile and are enforced consistently through integration and identity controls.

  This enables organisations to maintain operational control over AI behaviour, even as use cases expand.

- **Integration of AI Agents with APIs and Workflows**
  We integrate AI agents directly into existing APIs, process workflows and orchestration platforms rather than allowing them to act independently. This ensures AI-driven actions align with existing business logic, approvals and exception handling.

  AI becomes part of controlled enterprise processes rather than an unmanaged decision-maker.

- **Human-in-the-Loop Approval Mechanisms**
  For sensitive or high-impact activities, we design human-in-the-loop controls that allow AI agents to recommend or prepare actions rather than execute them autonomously. Human approval points provide oversight, accountability and reassurance.

  This supports trusted AI adoption while maintaining human responsibility.

- **AI Agent Orchestration Across Cloud Environments**
  We design architectures that support multiple AI agents working together across cloud and hybrid environments. Orchestration controls how agents collaborate, share context and delegate tasks while remaining observable and governed.

  This enables scalable AI adoption without losing visibility or control.

- **Full Audit Logging of Agent Decisions and Actions**
  All agent interactions, decisions and actions are logged with sufficient detail to support audit, investigation and assurance. Logs capture what context was used, what decision was made and what action was taken.

  This ensures traceability, explainability and public-sector-grade accountability.

- **Secure Access to Enterprise Data and Integrations**
  AI agents access enterprise data through governed APIs and integration layers rather than direct database or file access. This ensures consistent enforcement of data protection, access controls and security monitoring.

  AI access follows the same rules as other enterprise systems.

- **Vendor-Agnostic and Extensible AI Architecture**
  Architectures are designed to be independent of specific AI model vendors or platforms. This allows organisations to adopt new models, tools or capabilities as the AI landscape evolves without redesigning governance and integration foundations.

- **Alignment with GDS and Security Standards**
  All designs align with GDS guidance, public sector security standards and responsible AI principles. This ensures AI solutions are suitable for production use in regulated, high-scrutiny environments.

BENEFITS

Our MCP-Based Agentic AI Integration Services enable organisations to adopt agentic AI as a **trusted operational capability**, not an experimental tool. The benefits extend beyond efficiency gains to improved governance, reduced risk and long-term readiness for AI-enabled public services.

- **Standardised AI Interaction Across Platforms**

Without a common integration standard, AI agents often interact with systems in inconsistent and hard-to-manage ways. By using Model Context Protocol, organisations establish a single, standardised approach for how AI agents request context, invoke actions and return outcomes across cloud, hybrid and legacy systems.

This consistency simplifies architecture, reduces duplication and enables AI capabilities to be reused across departments and services rather than rebuilt each time.

- **Reduced Risk of Uncontrolled or Unintended AI Access**
  Unconstrained AI access to enterprise systems and data presents significant security, privacy and reputational risk. Through scoped context sharing and policy-driven permissions, organisations retain control over exactly what information an agent can see and what actions it can take.

  This prevents accidental overreach, limits exposure of sensitive data and ensures AI behaviour remains predictable and controlled, even as use cases evolve.

- **Maintained Governance and Operational Control**
  Agentic AI introduces new operational risks if decisions and actions cannot be governed or explained. Our approach embeds governance directly into AI integration, ensuring all agent actions are subject to defined policies, approval mechanisms and audit requirements.

  Organisations retain control over AI-driven activity, supporting assurance, accountability and compliance with internal governance and external scrutiny.

- **AI Aligned to Existing Integration and Process Architecture**
  Rather than creating parallel AI systems that bypass existing controls, AI agents are integrated into established APIs, workflows and orchestration platforms. This ensures AI-driven actions follow the same business rules, exception handling and approval paths as existing processes.

  AI becomes an extension of enterprise architecture rather than a disconnected decision-making layer.

- **Controlled and Trusted AI Operations**
  Human-in-the-loop mechanisms allow AI agents to recommend or prepare actions while preserving human authority over sensitive decisions. This creates trust among users, stakeholders and regulators, enabling AI adoption in high-impact areas without removing human accountability.

  Over time, this trust enables more advanced automation while maintaining appropriate oversight.

- **Scalable Multi-Agent Architectures Without Loss of Control**
  As AI adoption grows, organisations often need multiple agents performing specialised roles. MCP-based orchestration enables agents to collaborate while remaining governed, observable and auditable.

  This allows AI capability to scale across services, teams and environments without becoming opaque or unmanageable.

- **Clear Traceability, Explainability and Accountability**
  Every AI decision and action is logged with associated context, permissions and outcomes. This provides clear traceability for audits, investigations and assurance reviews.

  Organisations can explain how and why actions were taken, supporting public accountability and building confidence among oversight bodies.

- **Secure Foundation for Future AI Services**
  By embedding security, identity and governance at the integration layer, organisations create a foundation that supports future AI use cases without repeated risk assessment or redesign.

  This enables incremental expansion of AI capabilities with confidence.

- **Faster Adoption of New AI Capabilities**
  Vendor-agnostic, MCP-based architecture reduces dependency on specific models or platforms. Organisations can adopt newer models or tools as they emerge without reworking integration, governance or security foundations.

  This keeps AI strategy flexible in a fast-moving market.

- **Public-Sector-Ready AI Delivery**
  Alignment with GDS guidance, security standards and responsible AI principles ensures solutions are suitable for real-world government deployment, scrutiny and assurance.

  AI becomes a credible delivery capability rather than a research exercise.

WHAT THE SERVICE INCLUDES

Reply delivers a **comprehensive, end-to-end service** that covers strategy, architecture, implementation, operation and continuous improvement of agentic AI integration. The service is designed to embed agentic AI safely into enterprise environments over time.

1. Discovery and Agentic AI Readiness Assessment

We begin by understanding organisational readiness for agentic AI, focusing on where AI-driven actions deliver value without introducing unacceptable risk.
Activities include:

- Identification of candidate processes and decisions for agentic AI
- Assessment of data sensitivity, risk and regulatory constraints
- Evaluation of required oversight and approval mechanisms
- Alignment with organisational AI strategy and delivery priorities

This ensures agentic AI is targeted at suitable use cases rather than applied indiscriminately.

2. Agent Roles, Responsibilities and Operating Model Design

We define clear roles for AI agents, including what they are responsible for, how they interact with systems and where human oversight is required.
This includes:

- Definition of agent purpose and boundaries
- Human-AI interaction models
- Escalation and exception handling design
- Ownership and accountability models

This prevents ambiguous responsibility and uncontrolled behaviour.

3. MCP-Based Integration and Context Architecture

We design MCP-based architectures that define how agents receive context, access systems and take action.
Design includes:

- Context scoping and data minimisation
- Policy-driven permission models
- Secure interaction with APIs and workflows
- Environment-specific controls (e.g. dev, test, prod)

Architecture prioritises safety, clarity and auditability.

4. Secure Implementation of AI Agent Integrations

We implement secure connectors between agents, integration platforms and enterprise systems using standardised patterns.
Implementation includes:

- MCP connectors and interfaces
- Integration with API and workflow layers
- Identity and access enforcement
- Centralised logging and observability

This ensures AI actions flow through governed enterprise pathways.

5. Human-in-the-Loop and Oversight Controls

We design and implement approval steps, review workflows and escalation mechanisms where AI actions require human validation.
This includes:

- Approval thresholds for sensitive actions
- Separation of recommendation and execution
- Clear user interfaces for oversight
- Audit capture of human decisions

6. Testing, Validation and Assurance

AI agents are tested for correctness, behaviour, security and governance compliance. Testing includes edge cases and failure scenarios to ensure predictable behaviour.
This supports production readiness and assurance.

7. Deployment and Operational Readiness

We deploy solutions with appropriate monitoring, alerting, documentation and operational support processes to ensure safe live operation.

8. Monitoring, Audit and Day-to-Day Operations

We enable continuous monitoring of agent activity, decisions and outcomes.
This includes:

- Behaviour monitoring and alerts
- Audit dashboards and reports
- Operational insight into agent performance

This supports both operations and governance.

9. Governance, Change and Lifecycle Management

We establish governance frameworks covering updates to agents, policy changes, onboarding of new use cases and retirement of agents.
This prevents unmanaged growth and ensures long-term control.

10. Knowledge Transfer and Organisational Enablement

We provide documentation, training and knowledge transfer to ensure internal teams understand how agentic AI operates and how to manage it responsibly.

11. Continuous Optimisation and Improvement

We support ongoing refinement of agent behaviour, policies and integration patterns as use cases mature and confidence grows.

WHY THIS MATTERS

Together, these services transform agentic AI from a high-risk experiment into a governed, explainable and scalable capability. Organisations gain confidence that AI agents act transparently, securely and in alignment with public service values, enabling innovation while protecting trust, accountability and control.

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