

ENTERPRISE KNOWLEDGE INTEGRATION & AI KNOWLEDGE ASSISTANT

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 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

Public sector organisations hold vast amounts of knowledge across document repositories, collaboration tools, architecture models and operational systems. Policies, procedures, guidance and technical information are often spread across platforms such as SharePoint, Confluence, document management systems and enterprise architecture tools. As a result, staff spend significant time searching for information, interpreting conflicting guidance or relying on informal knowledge held by individuals. This slows decision-making, creates inconsistency and introduces operational risk.

Reply's Enterprise Knowledge Integration & AI Knowledge Assistants service for G-Cloud 15 addresses this challenge by securely integrating AI-powered knowledge assistants with enterprise content, systems and workflows. We enable staff to access trusted, up-to-date information through AI-assisted interfaces embedded into their daily work, while maintaining strong governance, access control and auditability.

Our approach focuses on making existing knowledge usable rather than creating new content silos. AI assistants retrieve and synthesise information from approved sources, applying role-based context and policy controls to ensure responses are relevant, accurate and appropriate. This supports consistent interpretation of policies and procedures, reduces reliance on manual search and improves organisational knowledge retention.

Crucially, the service is designed for public sector environments where trust, accountability and compliance are essential. AI-assisted knowledge access is governed, auditable and aligned with existing security and information management practices, enabling productivity gains without compromising control or assurance.

FEATURES

Our service combines secure integration, AI enablement and organisational adoption to deliver trusted knowledge assistance at enterprise scale.

- **Integration of AI Assistants With Enterprise Knowledge Sources**

We integrate AI knowledge assistants with approved enterprise knowledge repositories, including document management systems, collaboration platforms and architecture tools. This allows AI assistants to retrieve information directly from authoritative sources rather than relying on duplicated or manually curated content.

Integration is designed to respect existing structures, permissions and ownership models.

- **Secure Retrieval of Structured and Unstructured Content**

AI assistants retrieve and reason over both structured data (e.g. catalogues, registers, architecture models) and unstructured content (e.g. policies, procedures, guidance documents). Retrieval is governed by security controls to ensure only authorised content is accessed and surfaced.

This supports trusted, consistent responses grounded in approved information.

- **Context-Aware Responses Based on User Role**

Responses are tailored based on user identity, role and permissions. This ensures staff receive information relevant to their responsibilities while preventing inappropriate disclosure of sensitive or restricted content.

Context awareness improves relevance and reduces cognitive overload.



- **Integration With Operational Systems and Workflows**

AI knowledge assistants can be embedded into existing tools and workflows, enabling staff to access information where work is actually performed. This includes integration with service management, delivery and operational platforms.

Knowledge becomes embedded into daily activity rather than accessed through separate portals.

- **Policy-Controlled Knowledge Access**

Access to knowledge sources and response behaviour is governed by configurable policies. These controls define which sources can be queried, how information can be combined and how responses are framed.

Policy controls support compliance with organisational and regulatory requirements.

- **Audit Logging of AI Interactions**

All AI interactions are logged, creating a clear audit trail of queries, sources accessed and responses provided. This supports accountability, governance and assurance, particularly for decision-support scenarios.

- **Continuous Content Updates and Indexing**

Knowledge sources are continuously indexed and refreshed to ensure AI assistants reflect the latest approved content. This reduces the risk of outdated guidance being reused or propagated.

- **Human Validation and Review Workflows**

Where appropriate, human-in-the-loop workflows are introduced to review or validate AI responses, particularly for sensitive or high-impact use cases. This ensures quality and trust while AI capability matures.

- **Cloud-Agnostic Deployment Models**

AI knowledge assistants are designed to operate across cloud, hybrid or on-premise environments. This provides flexibility and avoids unnecessary dependency on specific vendors or platforms.

- **Training and Adoption Support**

We provide training and adoption support to help staff use AI knowledge assistants effectively and confidently, ensuring capability is embedded rather than underused.

BENEFITS

Our Enterprise Knowledge Integration & AI Knowledge Assistants service transforms fragmented organisational knowledge into a **trusted, governed and continuously available capability**. The benefits focus on improving productivity **without losing control**, which is critical in public sector environments.

- **Faster, More Reliable Access to Organisational Knowledge**

Staff no longer need to search across multiple systems, repositories and documents to find relevant guidance. AI knowledge assistants surface the most relevant information quickly, grounded in approved enterprise sources.



This reduces delays in decision-making and enables staff to act confidently with the right information at the right time.

- **Improved Trust Through Authoritative, Governed Responses**

By retrieving information only from approved enterprise knowledge sources, organisations reduce the risk of inaccurate, outdated or unofficial guidance being used. AI responses are explainable and traceable back to source content.

This builds trust among users and stakeholders, particularly where policies and procedures must be followed precisely.

- **Better Knowledge Retention and Reduced Dependence on Individuals**

Organisations often rely on informal, undocumented knowledge held by experienced individuals. AI-assisted access captures and operationalises this knowledge at scale, reducing risk when staff leave or roles change.

This improves organisational resilience and continuity.

- **Consistent Interpretation of Policies and Guidance**

AI assistants help ensure staff interpret policies, procedures and standards consistently across teams and locations. This reduces variability caused by different interpretations of the same documentation.

Consistency supports fairer decision-making and reduces operational and compliance risk.

- **Embedded Knowledge Support Within Daily Workflows**

By integrating AI assistants into operational systems and tools already used by staff, knowledge becomes available at the point of need. Users are not required to context-switch to separate systems to find answers.

This improves adoption and embeds knowledge support into everyday operations rather than treating it as a separate function.

- **Improved Compliance With Policies and Controls**

Policy-controlled access ensures staff see only the information they are authorised to access. Context-aware responses reduce the risk of sensitive information being shared inappropriately.

This supports compliance with information governance, data protection and internal control requirements.

- **Stronger Accountability and Oversight**

Audit logging of AI interactions provides visibility into how organisational knowledge is accessed and used. This supports audit, assurance and investigation activities where decisions rely on AI-assisted information.

Accountability is preserved even as access becomes more automated.

- **Up-to-Date Knowledge Without Manual Rework**



Continuous indexing and refresh of content ensures AI assistants reflect the most current approved guidance. Organisations avoid the risk of outdated documents being relied upon due to manual updates being missed.

- **Maintained Quality Through Human Oversight**

Human-in-the-loop workflows allow responses to be reviewed or validated where accuracy is critical. This balances automation with appropriate human judgement and maintains confidence in AI-assisted outputs.

- **Reduced Reliance on Specialist Teams for Routine Queries**

Staff can self-serve routine knowledge needs, reducing the volume of queries escalated to specialist or central teams. This frees experts to focus on higher-value work while improving responsiveness across the organisation.

- **Significant Reduction in Time Spent Searching for Information**

By eliminating manual search and interpretation effort, staff reclaim time currently lost navigating fragmented knowledge estates. This productivity gain compounds over time across large workforces.

WHAT THE SERVICE INCLUDES

Reply delivers a **comprehensive, end-to-end service** that integrates AI knowledge assistants securely into enterprise environments, with governance and assurance embedded throughout. The service is designed as a **living capability**, not a one-off deployment.

1. Knowledge Landscape Discovery and Assessment

We begin by understanding where knowledge currently resides, how it is used, and where fragmentation causes friction or risk.

This includes:

- Inventory of knowledge repositories and tools
- Assessment of content ownership and governance
- Review of access controls and sensitivity
- Identification of high-value knowledge use cases

This ensures AI assistance focuses on real organisational needs rather than generic search.

2. Knowledge Source Selection and Integration

We integrate AI assistants with approved enterprise knowledge sources such as document repositories, collaboration platforms and architecture tools.

Activities include:

- Secure connector configuration
- Alignment with existing access permissions
- Metadata and structure mapping
- Validation of content suitability

Integration respects existing governance rather than bypassing it.

3. Security, Access Control and Policy Configuration

We design and implement role-based access controls and policy rules governing how AI retrieves and presents information.

This includes:

- User identity and role mapping



- Source-level access enforcement
- Policy rules controlling response behaviour
- Safeguards for sensitive content

These controls ensure AI access aligns with organisational and regulatory expectations.

4. Context Awareness and Response Design

We configure AI assistants to tailor responses based on user context, role and task.

This improves:

- Relevance of responses
- Accuracy of interpretation
- Reduction of unnecessary or inappropriate information

Context awareness increases usefulness without increasing risk.

5. Integration With Operational Systems and Tools

We embed AI knowledge assistants into systems and workflows used in daily operations.

This includes:

- Integration with service and delivery tools
- Contextual knowledge prompts
- Workflow-aware responses

Knowledge becomes part of the operational fabric rather than a separate channel.

6. Human-in-the-Loop Validation and Quality Controls

For sensitive or high-impact knowledge, we design validation workflows where responses can be reviewed, approved or refined by subject matter experts.

This ensures quality is maintained as AI usage grows.

7. Audit Logging, Monitoring and Assurance Enablement

We implement comprehensive logging of AI interactions to provide transparency and traceability.

This supports:

- Audit and assurance activities
- Investigation and review
- Continuous improvement based on usage insight

Logs are designed to be understandable and usable by non-technical stakeholders.

8. Continuous Content Indexing and Knowledge Refresh

We configure ongoing indexing and refresh processes to ensure AI assistants reflect current approved knowledge.

This includes:

- Automated refresh schedules
- Handling of content updates and retirements
- Validation of indexing integrity

9. Training, Adoption and Behavioural Enablement

We support staff adoption through training, guidance and communication, helping users understand how to use AI knowledge assistants appropriately and effectively.

Adoption support focuses on trust, correct use and confidence rather than technology alone.

10. Continuous Improvement and Capability Evolution

We use usage patterns, feedback and operational insight to refine assistant behaviour, improve relevance and identify opportunities for expansion.



The capability evolves as organisational needs change.

WHY THIS MATTERS

Together, these service components replace fragmented, informal knowledge access with a secure, governed and scalable capability. Organisations gain faster access to trusted information while retaining control, accountability and confidence in how AI is used. This enables staff to make better decisions, apply policies consistently and focus on delivering high-quality public services rather than searching for information.



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