

Data & AI Services- Service Definition

G-Cloud 15 Lot 3 Services

(Cloud Support)



Registered Address

142 Cromwell Road
London
SW7 4EF

Contact information

Phone: +44 20 3151 0307
Email: info@7Dxperts.com

Company Details

Company Reg	07515669
VAT	109841900

Document Control

Purpose

This document provides a single, buyer-facing Core Service Definition for 7Dxperts’ G-Cloud 15 Lot 3 Cloud Support services, together with ten (10) Service Addenda. The Core Service Definition describes how we deliver, govern, assure and support engagements, while each Addendum provides service-specific scope, deliverables and constraints.

How to use this document

Buyers can use this document to:

- Confirm that the service meets their needs and that the scope aligns to their intended outcomes.
- Understand the delivery approach
- Identify assumptions, dependencies and constraints that may need to be clarified in the call-off contract.

Where buyer requirements differ (for example, specific technology platforms, governance standards, or additional workshops), variations can be agreed and documented in the call-off contract and may affect price and timescales.

Use Sections 1–12 as the common description for all Lot 3 services.

Use the relevant Service Addendum (Appendix A1–A10) to describe service-specific activities, deliverables, dependencies and constraints.

Use the separate SFIA-aligned rate card for pricing, and the common terms and conditions for contractual terms (both provided as separate attachments).

Document conventions

- “Buyer” means a public sector customer purchasing via G-Cloud 15.
- “Call-off” means the contract formed under the framework for a specific engagement.
- “Acceptance Authority” is the Buyer’s nominated approver for acceptance decisions and sign-off.

Version

Date	Version	Distribution
12th Jan 2026	0A	Internal
30th Jan 2026	1A	Public, Digital Marketplace

Contents

1 *Service Overview* _____ 5

 1.1 What the service covers _____ 5

 1.2 What the service is not _____ 6

2 *Who this service is for* _____ 6

3 *Platforms and technologies* _____ 6

4 *Service model and supplier type* _____ 7

5 *Delivery lifecycle* _____ 7

 5.1 Typical mobilisation lead time _____ 8

6 *Governance and ways of working* _____ 9

 6.1 Governance approach _____ 9

 6.2 Roles and responsibilities _____ 9

7 *Quality, assurance and acceptance* _____ 9

 7.1 Quality approach _____ 9

 7.2 Default acceptance method _____ 10

 7.3 Service-specific acceptance variants _____ 10

8 *Security, privacy and compliance* _____ 10

9 *Support and operations* _____ 10

 9.1 Engagement model _____ 10

 9.2 Channels _____ 11

 9.3 Operating hours, response targets and service levels _____ 11

 9.4 Support tiers _____ 11

 9.5 Processes and Functions _____ 12

10 *Reporting and service reviews* _____ 14

11 *Buyer responsibilities and dependencies* _____ 14

12 *Assumptions and constraints* _____ 14

13 *Appendix A – Service Addenda* _____ 15

 13.1 Data, AI and Operational DB Transform, Implement, Train and Support _____ 15

 13.2 Data & AI Agentic Platform Architecture and Design Assurance _____ 16

 13.3 Analytics Modernisation, AI, Agentic, Semantic Layer and BI Enablement _____ 17

 13.4 Data Engineering and Integration Services _____ 18

 13.5 Data Migration, Modernisation and Application Retirement _____ 19

13.6 Agent Tooling, Function Calling and Workflow Orchestration _____ 20

13.7 AI Agentic Platform Enablement and ML Model Operations _____ 21


13.8 Operational PostgreSQL Enablement for Data Apps and AI Agents _____ 22

13.9 RAG Enablement and Enterprise Knowledge Integration _____ 23

13.10 Structured Data Text-to-SQL and Semantic Enablement for Agents _____ 24

1 Service Overview

We build targeted data driven solutions, to answer tough questions through analytics and machine learning and enable organisation to leverage AI



Our Footprint

- Offices in 2 countries
- 45+ headcount

Our Expertise

- Data & Analytics
- Machine Learning & Data Science
- Gen AI, Agentic AI
- GeoSpatial Analytics
- Data & Information Governance

7Dxperts provides cloud support services that help Buyer's design, implement, modernise, secure, operate and continuously improve data, analytics and AI platforms and workloads. We support Buyers across the full lifecycle—from discovery and design, through implementation and migration, to operational stabilisation and managed support—tailored per engagement.

We have identified the following ten (10) services that will benefit buyers across the Digital Marketplace

1. Data & AI Agentic Platform Architecture and Design Assurance
2. Analytics Modernisation, AI, Agentic, Semantic Layer and BI Enablement
3. Data Engineering and Integration Services
4. Data Migration, Modernisation and Application Retirement
5. Agent Tooling, Function Calling and Workflow Orchestration
6. AI Agentic Platform Enablement and ML Model Operations
7. Operational PostgreSQL Enablement for Data Apps and AI Agents
8. RAG Enablement and Enterprise Knowledge Integration
9. Structured Data Text-to-SQL and Semantic Enablement for Agents

This granular level of services have been designed to ensure that the buyer has flexibility to select one or more services and has clarity on what outcomes to expect.

1.1 What the service covers

- Design and delivery of cloud-based data and analytics solutions, including modern Lakehouse / warehouse patterns, semantic layers and BI enablement.

- Implementation and operation of AI/ML and agentic capabilities (including RAG, text-to-SQL and workflow orchestration) with governance, security and observability.
- Migration, modernisation and retirement of legacy data platforms and applications, including cutover planning and post-migration validation.
- Operational support, incident management, change management and optimisation, aligned to Buyer operating hours and agreed SLAs.

1.2 What the service is not

- We do not provide unrestricted 24x7 managed service by default; enhanced / 24x7 support is available only if separately contracted.
- We do not act as the Buyer's data controller.
- We do not provide statutory, clinical or legal decision-making; outputs support operational decision-making only.

2 Who this service is for

This service is suitable for Buyers that need to modernise and operate cloud data and AI platforms, improve reliability and security of existing platforms, or introduce governed AI and agentic capabilities across operational teams and digital services.

- Public sector organisations with complex data estates and reporting obligations.
- Teams adopting Databricks, Snowflake, Microsoft Fabric, Azure, AWS, Google Cloud/BigQuery and related analytics and AI tooling.
- Organisations seeking to improve service continuity via runbooks, monitoring, incident management, and operational handover.

3 Platforms and technologies

We are platform-agnostic and deliver on Buyer-selected platforms. Typical platforms and tools include:

- Cloud platforms: Microsoft Azure, Amazon Web Services (AWS), Google Cloud Platform (GCP).
- Data platforms: Databricks, Snowflake, Microsoft Fabric, BigQuery, PostgreSQL and associated cloud-native services.
- Analytics and BI: ThoughtSpot, Power BI, Tableau, Sigma, semantic layer tooling and governed metrics.
- AI/ML and agentic tooling: MLflow / model registry, feature stores, evaluation harnesses, vector search, RAG patterns, tool/function calling, orchestration and observability.
- DevOps and operations: CI/CD tooling, infrastructure-as-code (Buyer-selected), monitoring/alerting stacks, and ticketing/service management tooling (e.g., Jira Service Management or Buyer equivalents)

However, we do have strategic partnerships with the following vendors that we believe are leading edge for data analytics, machine learning and AI



Implementing Lakehouse architecture for consolidating all your data to support advance analytics capability working in multiple verticals and/or building ML learning platforms.

<https://www.databricks.com/company/partners/consulting-and-si>



Forward generate Databricks assets into ThoughtSpot and accelerate enablement of GenAI Natural Language Search on your data. We support customer how to model data for Ai and recommend Human in the loop framework for Governance.

<https://www.thoughtspot.com/partners/directory>


4 Service model and supplier type

7Dxperts provides **implementation** and operational **support** on Buyer-selected platforms. We also provide additional features, **accelerators** and support services that are not available from original software vendors, where appropriate (for example: reusable delivery patterns, migration tools, governance artefacts, operational runbooks, evaluation harnesses, and integration accelerators).

5 Delivery lifecycle

We deliver through standard phases, tailored per engagement and agreed in the call-off. Typical phases are:

1. Mobilise and onboard
2. Discover / assess and plan
3. Design
4. Build / configure
5. Test and validate
6. Deploy and transition
7. Operate and optimise
8. Knowledge transfer and close




For Agile deliveries, stages 3-6 are iterative and handled within a sprint.

5.1 Typical mobilisation lead time

Typical mobilisation is 1–2 weeks from call-off signature, subject to Buyer access approvals, security onboarding and availability of nominated Buyer stakeholders (or as agreed in the call-off).

6 Governance and ways of working

6.1 Governance approach



20

- Delivery leads: a 7Dxperts Delivery Lead (Project Manager) and a Buyer Product Owner / Service Owner.
- Regular cadence: stand-ups / working sessions during delivery, project boards/SteerCo, plus planned service reviews during operate/support where applicable.
- Risk and issue management: RAID log (risks, assumptions, issues and dependencies) maintained and reviewed as part of governance.
- Change control: scope, timeline and cost changes handled through a documented change control process.

6.2 Roles and responsibilities

- Buyer: nominates appropriate contacts to support decisions and acceptance; provides timely access to environments, data sources and stakeholders; and provides licences/tenancies required for the platforms.
- 7Dxperts: delivers agreed scope with reasonable skill and care; maintains delivery artefacts; provides operational runbooks and knowledge transfer; and supports acceptance and transition activities.

7 Quality, assurance and acceptance

7.1 Quality approach

- Definition of Done and acceptance criteria agreed per deliverable / milestone.
- Peer review and technical assurance of key artefacts (architecture, security design, code, configuration).
- Rigorous test approach from Unit Testing, System Integration Testing, Data Acceptance Testing, User Acceptance Testing and Production Acceptance Testing
- Validation activities appropriate to the engagement: reconciliation checks, test evidence, performance checks, and operational readiness checks.

7.2 Default acceptance method

Default acceptance is Buyer acceptance by written sign-off against agreed acceptance criteria. Unless otherwise agreed in the call-off, acceptance decisions are made by the Buyer's nominated Acceptance Authority.

7.3 Service-specific acceptance variants

Where a service requires service-specific acceptance variants (for example, additional operational readiness checks or model evaluation gates), these are captured in the relevant Service Addendum and/or the call-off statement of work.

8 Security, privacy and compliance


Security and privacy are designed into delivery and operations. We implement least-privilege access, role segregation, and auditable controls aligned to Buyer policies and governance requirements.

- Identity & access: SSO integration, RBAC/ABAC patterns, and segregation of duties.
- Data protection: classification-aware access control design and secure connectivity patterns.
- Auditability: lineage, approval evidence, and operational runbooks to support governance.
- Security standards: controls aligned to Cyber Essentials Plus requirements and ISO/IEC 27001-aligned policies and procedures (evidence available on request).

9 Support and operations

9.1 Engagement model

7Dxperts provides a **platform-agnostic**, ITIL-aligned support service that can be delivered as a standalone support engagement or as the operational wrap-around to delivery/migration work.



9.2 Channels

Support is provided through the following channels (as agreed in the Call-Off):

- **Ticketing / email:** Support requests are logged via the Buyer's nominated service desk tool (or a shared mailbox/ticketing queue where agreed). 7Dxperts can operate as a **resolver group** within the Buyer's existing service desk, receiving escalations from first-line support and coordinating investigation and resolution.
- **Phone (optional):** Phone support can be provided where agreed in the Call-Off (typically used for Priority 1/2 incidents and operational coordination).
- **Web chat:** Web chat support is **not provided by default**, but may be agreed where required.

9.3 Operating hours, response targets and service levels

- **Operating hours, severity definitions, response targets and service levels** are agreed per Call-Off to reflect the criticality of the service, environments supported, and Buyer demand patterns.
- Unless otherwise agreed in the Call-Off, a typical baseline support response target is **within 4 hours, Monday–Friday 09:00–17:00 UK time**.
- **24/7 and/or on-call support** can be provided on request at additional cost, subject to the size, complexity and criticality of the Buyer's environment and the commitments required.

9.4 Support tiers

7Dxperts provides three support tiers. The selected tier determines the level of proactivity, the service management processes in scope, and the governance cadence. **Operating hours and SLAs remain Call-Off specific.**

9.4.1 Tier 1 – Standard Support (Reactive + controlled)

Best for Buyers who need reliable “run and fix” support with appropriate operational controls.

- **Service Operation:** incident triage, investigation, service restoration, stakeholder updates; engineering support for fixes and minor corrective actions.
- **Service Transition (lightweight):** change control for fixes/updates required to restore service; release/deployment coordination as needed.
- **Operational enablement:** onboarding to support, escalation paths, operational readiness checklist, and initial runbooks (where available).
- **Reporting:** basic operational reporting (e.g., volumes, response/restoration, top recurring issues) at a cadence agreed in the Call-Off.

9.4.2 Tier 2 – Enhanced Support (Proactive + measurable)

Best for Buyers who want proactive stability, measurable performance and structured improvement.

Includes everything in Standard, plus:

- **Event management / monitoring (where in scope):** agreed monitoring coverage, thresholds and alerting for early warning and exception reporting.
- **Problem management (where in scope):** root cause analysis for material incidents, trend analysis, identification of recurring issues, and prioritised recommendations for long-term fixes.
- **Testing & validation (proportionate):** validation of fixes and changes prior to production where relevant.
- **Knowledge management:** creation/maintenance of knowledge articles, known errors and operational playbooks.
- **Governance:** service level reporting aligned to Call-Off targets and a regular service review (e.g., monthly) including an improvement backlog covering reliability, performance, security posture and cost optimisation.

9.4.3 Tier 3 – Premium Support (Mission-critical + rapid escalation)


Best for Buyers with business-critical services requiring rapid escalation, deeper assurance and stronger operational control.

Includes everything in Enhanced, plus:

- **On-call / extended cover:** on-call capability for Priority 1 incidents (and extended hours where required) as agreed in the Call-Off.
- **Proactive health checks:** planned reviews for capacity, performance hotspots, reliability risks, cost anomalies and security control effectiveness (scope agreed per Call-Off).
- **Release/change assurance:** enhanced release planning, cutover readiness, rollback/contingency planning where applicable, and post-implementation review for high-risk changes.
- **Service ownership:** a named Technical Account Manager (TAM) for escalation management, service alignment, and continuous improvement leadership.
- **Supplier/vendor coordination:** proactive engagement with third-party providers where dependencies impact service outcomes (where applicable).

9.5 Processes and Functions

The various processes and functions are illustrated below.



9.5.1 Incident, event, problem and change management

Support is delivered using structured operational processes, with scope dependent on the selected support tier and the Call-Off:

- **Incident management:** coordinated investigation, classification, prioritisation, service restoration and communications.
- **Event management (where in scope):** monitoring and automated alerting for early warning and exception reporting; thresholds and coverage agreed per Call-Off.
- **Problem management (where in scope):** RCA, recurring issue identification, known error records, and recommendations/implementation of long-term fixes.
- **Change and release management:** governance for planned/unplanned change, approvals, release scheduling, deployment coordination and validation prior to production (proportionate to risk and criticality).
- **Resolver group model:** where the Buyer uses an existing service desk, 7Dxperts can operate as a resolver group, receiving escalations and working within Buyer workflows and policies. Where monitoring is in scope, 7Dxperts may also raise incidents proactively when issues are detected.

9.5.2 Operational handover and transition to support

Where 7Dxperts transitions a platform, product or capability into support, we provide (as applicable):

- **Operational readiness checklist and handover pack** (including escalation paths, support scope, monitoring coverage and access prerequisites).

- **Runbooks and incident playbooks** for common operations and recovery activities, including known error guidance where available.
- **Knowledge transfer and enablement sessions** for Buyer engineers, support staff and operational stakeholders, covering day-to-day operations, release/change procedures, monitoring dashboards and key controls.

10 Reporting and service reviews

- Operational reporting (ticket volumes, response times, resolution times, recurring issues) where support is provided.
- Service review cadence (e.g., monthly or quarterly) where a managed support retainer is in place.
- Continuous improvement backlog: prioritised improvements agreed with the Buyer (performance, cost, reliability, security posture).

11 Buyer responsibilities and dependencies

- Provide timely access to cloud tenancies/subscriptions, identity services, networks, logs and environments.
- Provide third-party licenses/subscriptions for chosen platforms and confirm vendor permissions where applicable.
- Nominate appropriate contacts for delivery decisions and acceptance sign-off.
- Provide data owners/information owners for access and classification decisions.
- Provide change windows and participate in testing and validation as required.

12 Assumptions and constraints

- Delivery is remote-first; on-site workshops/support can be provided by agreement and may incur travel costs.
- Timelines depend on access approvals, security reviews, third-party vendor dependencies, and data quality/legacy complexity.
- Out-of-scope requests are managed through change control.
- Where solutions rely on upstream systems and third-party services, performance and availability may depend on those providers.

13 Appendix A – Service Addenda

The following addenda tailor the Core Service Definition for each Lot 3 service.

13.1 Data, AI and Operational DB Transform, Implement, Train and Support

13.1.1 Service summary

End-to-end delivery and support service covering transformation, implementation, training and support for cloud data platforms, analytics workloads and operational databases.

13.1.2 Typical scope and activities

- Platform setup and configuration; connectivity and secure access enablement.
- Data engineering and transformation delivery; optimisation and performance tuning.
- Analytics enablement including dashboards and self-service adoption support.
- AI/ML enablement including MLOps patterns, automation and best practices.
- Operational handover, runbooks, SLAs and expert escalation.

13.1.3 Typical buyer outcomes

- Improved platform availability via monitoring, proactive support and rapid remediation.
- Faster incident resolution through runbooks and structured escalation.
- Accelerated adoption through training and enablement for teams.
- Optimised cloud spend via tuning and workload scheduling.

13.1.4 Platforms and tooling

- Databricks, Snowflake, Microsoft Fabric, AWS, Azure, GCP and associated services.
- Operational databases such as PostgreSQL where applicable.

13.1.5 Constraints and dependencies

- Remote-first delivery; on-site by agreement and may incur travel costs.
- Buyer provides timely access to subscriptions/tenancies, networks, identity services and data sources.
- Buyer holds required third-party licences and provides suitable test environments.
- Service hours and response times follow agreed SLAs; out-of-scope requests via change control.

13.1.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.1.7 Notes and service-specific acceptance considerations

Support tiers (Standard/Enhanced/Premium) can be selected per call-off, with operating hours and SLAs agreed per environment and business criticality.

13.2 Data & AI Agentic Platform Architecture and Design Assurance

13.2.1 Service summary

Architecture assurance service to validate and strengthen Buyer platform designs before build and migration, reducing risk and rework.

13.2.2 Typical scope and activities

- Target architecture review and assurance for data and AI/ML platforms.
- Identity and access architecture review and assurance.
- Non-functional requirements and resilience design assurance.
- Vendor/cloud service selection support and evaluation.
- Architecture decision records (ADRs) and documentation pack creation.
- Technical risk identification and mitigation recommendations.

13.2.3 Typical buyer outcomes

- Reduced rework by assuring architecture before build and migration.
- Improved security posture through early IAM assurance.
- Optimised costs through service selection and design choices.
- Smoother handover with complete architecture documentation.
- Platforms and tooling
- Databricks, Snowflake, Microsoft Fabric, AWS, Azure, GCP and associated services.
- Buyer code repositories and CI/CD tooling (read-only access for review where applicable).

13.2.4 Constraints and dependencies

- Assurance outputs are recommendations and artefacts; build and migration execution is out of scope unless separately procured.
- Buyer provides timely access to stakeholders, standards and (where applicable) read-only access to tenants and repositories.
- Dependencies on vendor services and security approvals may affect timelines; changes managed via change control.

13.2.5 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.2.6 Notes and service-specific acceptance considerations

Acceptance typically based on written sign-off of the assurance pack (review findings, ADRs, and recommended target architecture).

13.3 Analytics Modernisation, AI, Agentic, Semantic Layer and BI Enablement

13.3.1 Service summary

Modernise analytics platforms and content, implement governed semantic layers, and enable BI and agentic use-cases with secure access controls.

13.3.2 Typical scope and activities

- Discovery and current-state assessment.
- Target architecture design and assurance.
- Semantic layer modelling and governance.
- Data wrangling migration and optimisation.
- Dashboard and report migration and validation.
- Identity, access and row-level security configuration.
- Performance tuning and cost optimisation.
- AI and agentic use-case design and enablement.
- Operational handover, runbooks and knowledge transfer.
- Platform set-up for ThoughtSpot, Sigma and Power BI.

13.3.3 Typical buyer outcomes

- Reduced time and risk to migrate analytics platforms and content.
- Improved data consistency via governed semantic models.
- Improved security through standardised access controls and RLS.
- Higher adoption through training, enablement and documented runbooks.

13.3.4 Platforms and tooling

- ThoughtSpot, Power BI, Tableau, Sigma and semantic layer tooling.
- Databricks, Snowflake, Microsoft Fabric and other data backends as required.

13.3.5 Constraints and dependencies

- Buyer provides access to cloud tenancies, data sources, networks, SSO/identity and required licences.
- Migration timelines depend on data quality, legacy complexity and vendor dependencies.
- Out-of-scope changes are managed via change control.
- Ongoing 24/7 managed support is not provided unless separately contracted.

13.3.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, ThoughtSpot, Snowflake, Tableau

13.3.7 Notes and service-specific acceptance considerations

Acceptance typically includes validated migrated dashboards/reports and documented semantic models and access control configurations.

13.4 Data Engineering and Integration Services

13.4.1 Service summary

Design, build and optimise data ingestion, transformation and integration pipelines across cloud platforms with governance and operational readiness.

13.4.2 Typical scope and activities

- Data ingestion and integration pipeline design and build.
- ETL/ELT automation; orchestration patterns and scheduling.
- Data modelling and curation (e.g., lakehouse/warehouse layers).
- Data quality checks and observability.
- Security, governance and compliance controls implementation.
- Performance tuning and cost optimisation.
- Operational handover, runbooks and knowledge transfer.

13.4.3 Typical buyer outcomes

- Improved reliability and trust through data quality controls and observability.
- Reduced engineering toil through automation and reusable patterns.
- Improved performance and scalability for analytics and downstream services.
- Better supportability via runbooks and documented pipelines.

13.4.4 Platforms and tooling

- Databricks, Snowflake, Microsoft Fabric, BigQuery and associated orchestration/services.
- Integration with APIs and operational systems as required.

13.4.5 Constraints and dependencies

- Buyer provides access to source/target systems, APIs, networks and identity services.
- Third-party licences and platform quotas apply.
- Data quality issues or undocumented interfaces may require additional remediation outside initial scope.
- On-site workshops by agreement; otherwise remote-first.

13.4.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.4.7 Notes and service-specific acceptance considerations

Acceptance typically includes pipeline run evidence, reconciliation checks and operational handover pack.

13.5 Data Migration, Modernisation and Application Retirement

13.5.1 Service summary

Migrate data estates to modern cloud platforms, modernise legacy applications where required, and retire/decommission legacy components with validated cutover.

13.5.2 Typical scope and activities

- Cloud migration discovery and readiness assessment.
- Data mapping, profiling and quality remediation.
- ETL pipeline design and automation; lakehouse and data warehouse enablement.
- Cloud landing zone integration and connectivity setup.
- Application modernisation planning and refactoring support.
- Legacy application decommissioning and retirement planning.
- Cutover planning, rehearsal and migration execution.
- Post-migration validation, reconciliation and performance tuning.
- Security, governance and compliance controls implementation.
- Knowledge transfer, runbooks and operational handover.

13.5.3 Typical buyer outcomes

- Reduced migration risk through structured discovery and planning.
- Minimised downtime through rehearsed cutover and rollback.
- Improved data quality and trust through validation and reconciliation.
- Lower operating costs by retiring legacy applications.

13.5.4 Platforms and tooling

- Databricks, Snowflake, Microsoft Fabric, AWS, Azure, GCP and associated migration tooling.
- Legacy source systems as provided by the Buyer.

13.5.5 Constraints and dependencies

- Work is time-boxed to an agreed statement of work and relies on timely Buyer input and access.
- Migration schedules depend on security approvals, change windows and operational constraints.
- Buyer holds third-party licences/tenancies and provides vendor permissions.
- Undocumented legacy technical debt may require additional remediation outside initial scope.

13.5.6 Supplier type and vendors

- Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.
- Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.5.7 Notes and service-specific acceptance considerations

Acceptance typically includes reconciliation results, performance baseline evidence and signed transition-to-operations checklist.

13.6 Agent Tooling, Function Calling and Workflow Orchestration

13.6.1 Service summary

Design and implement secure agent tools and workflows, including function calling, orchestration, approvals and audit trails so agents can execute tasks reliably.

13.6.2 Typical scope and activities

- Design tool catalogues for agents, including permissions and inputs.
- Implement secure function calling with validation and allow-lists.
- Build workflow orchestration for multi-step tasks and approvals.
- Integrate with APIs, ticketing, CRM and case management systems.
- Add audit trails for tool use, decisions and outputs.
- Implement error handling, retries, timeouts and fallback strategies.
- Configure observability hooks for tracing, metrics and alerting.
- Apply human-in-the-loop controls for high-impact actions.
- Produce runbooks and handover for operations and continuous improvement.

13.6.3 Typical buyer outcomes

- Automated tasks with approvals and oversight.
- Improved reliability through validated tools, retries and fallbacks.
- Strengthened governance through audit trails and least-privilege execution.
- Reduced operational risk with human-in-the-loop controls for critical actions.

13.6.4 Platforms and tooling

- Delivered on Databricks, Snowflake or Microsoft Fabric (and integrated with Buyer systems).
- Integration with ticketing/CRM/case management tooling as required.

13.6.5 Constraints and dependencies

- Buyer must provide API access, integration endpoints, test environments and change windows.
- Tool execution requires agreed approval thresholds, audit requirements and role definitions.
- Third-party licences, rate limits and platform quotas apply.
- Reliability depends on upstream system availability and data quality; human-in-the-loop may affect end-to-end latency.

13.6.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP, ThoughtSpot

13.6.7 Notes and service-specific acceptance considerations

Acceptance typically includes tool catalogue, orchestration workflows, audit/trace evidence and operational runbooks.

13.7 AI Agentic Platform Enablement and ML Model Operations

13.7.1 Service summary

Managed cloud service enabling AI agentic platforms and ML model operations, covering environment setup, secure connectivity, CI/CD, monitoring, incident management, cost controls and lifecycle governance.

13.7.2 Typical scope and activities

- Secure landing zones, identity and networking for AI platform workloads.
- MLflow and model registry setup for versioning and approvals.
- Automated build, test and deployment pipelines for models and agents.
- Model performance, drift and data quality monitoring with alerting.
- Agent evaluation harnesses, prompt/version control and rollback procedures.
- Incident, problem and change management aligned to agreed operating hours.
- Cost tagging, chargeback reporting and optimisation across platform workloads.
- Backup, recovery and resilience testing for critical data and models.
- Lifecycle governance for models, features, access and audit evidence.
- Workspace provisioning and governance baselines for Databricks/Snowflake/Fabric.

13.7.3 Typical buyer outcomes

- Accelerated safe delivery of production-grade agents and models.
- Reduced operational risk through monitoring, governance and change control.
- Improved model reliability with drift detection and rapid rollback.
- Optimised platform spend using tagging, chargeback and remediation actions.

13.7.4 Platforms and tooling

- Databricks, Snowflake and Microsoft Fabric; AWS/Azure/GCP infrastructure as required.
- Buyer monitoring/logging stack where applicable.
- Constraints and dependencies
- Buyer provides and maintains relevant cloud tenancies/subscriptions and third-party licences.
- Service is time-bound and scoped to agreed platforms, workloads and operating hours; changes via change control.
- Performance and availability depend on underlying cloud and vendor services; some issues may require vendor engagement.
- On-site support by prior agreement and may incur travel costs.

13.7.5 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.7.6 Notes and service-specific acceptance considerations

Acceptance typically includes operational readiness, monitoring/alerting evidence, CI/CD pipeline evidence and model governance artefacts.

13.8 Operational PostgreSQL Enablement for Data Apps and AI Agents

13.8.1 Service summary

Enable and operate PostgreSQL as an operational database for data applications and AI agents, with secure access, performance, backup/recovery and operational runbooks.

13.8.2 Typical scope and activities

- PostgreSQL environment setup and secure connectivity patterns.
- Schema design support and performance optimisation.
- Backup, recovery and resilience design/testing.
- Monitoring, alerting and operational readiness.
- Access control design (roles, permissions) and auditability where required.
- Integration patterns for data apps and AI agents.

13.8.3 Typical buyer outcomes

- Reliable operational datastore for data apps and agents.
- Improved performance and stability through tuning and monitoring.
- Reduced operational risk via tested recovery procedures and runbooks.
- Better governance through access controls and audit-ready operations.

13.8.4 Platforms and tooling

- PostgreSQL on Buyer-chosen hosting (cloud managed service or Buyer-managed).
- Integration with Databricks, Snowflake, Fabric and agentic tooling as required.

13.8.5 Constraints and dependencies

- Buyer provides hosting platform/subscriptions and access approvals.
- Third-party services and quotas may apply depending on hosting selection.
- Availability depends on underlying infrastructure and Buyer change windows.

13.8.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: AWS, Azure, GCP, PostgreSQL ecosystem tools

13.8.7 Notes and service-specific acceptance considerations

Acceptance typically includes operational runbooks, backup/restore test evidence and monitoring configuration.

13.9 RAG Enablement and Enterprise Knowledge Integration

13.9.1 Service summary

Enable retrieval-augmented generation (RAG) using enterprise knowledge sources: ingest documents, index securely, and implement grounding patterns so agents answer with evidence.

13.9.2 Typical scope and activities

- Document ingestion pipelines for PDFs, emails, web pages and records.
- Secure indexing and search configuration for enterprise repositories.
- Chunking, embedding and vector retrieval patterns with governance controls.
- Grounding prompts and citation metadata to reduce hallucinations.
- Access control design for sensitive data and role segregation.
- Evaluation harnesses for relevance, accuracy, latency and cost.
- Deployment packaging for APIs, chat interfaces and embedded applications.
- Monitoring setup for retrieval quality, failures and drift.
- Operational runbooks, handover and knowledge transfer to Buyer teams.

13.9.3 Typical buyer outcomes

- Improved answer accuracy using current organisational knowledge sources.
- Reduced time spent searching policies, procedures and guidance.
- Stronger governance through controlled access and auditing.
- Measurable quality through evaluation and monitoring.

13.9.4 Platforms and tooling

- Delivered on Databricks, Snowflake or Microsoft Fabric and associated vector search/embedding services.
- Integration with Buyer identity provider and approved knowledge repositories.

13.9.5 Constraints and dependencies

- Buyer provides approved cloud tenancy, identity provider and nominated information owners.
- Knowledge ingestion depends on document availability, classification and data-sharing approvals.
- Response quality depends on source content quality and ongoing curation.
- Outputs support operational decision-making and do not constitute legal advice.

13.9.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, AWS, Azure, Microsoft Fabric, GCP

13.9.7 Notes and service-specific acceptance considerations

Acceptance typically includes evaluation results (relevance/accuracy), access control evidence and operational runbooks.

13.10 Structured Data Text-to-SQL and Semantic Enablement for Agents

13.10.1 Service summary

Enable natural-language querying of structured data by implementing governed semantics, text-to-SQL patterns and agent-safe query execution controls.

13.10.2 Typical scope and activities

- Semantic modelling, governed metrics and business vocabulary definition.
- Text-to-SQL enablement with safe query patterns and validation.
- Role-based access control and row/column-level security alignment.
- Query guardrails (allow-lists, query limits, cost controls) and audit trails.
- Evaluation harnesses for accuracy, safety and performance.
- Integration into agent workflows and chat interfaces with observability hooks.
- Operational runbooks and handover to Buyer teams.

13.10.3 Typical buyer outcomes

- Faster self-service insights through natural language access to data.
- Reduced risk via governed semantics and query guardrails.
- Improved trust and consistency via shared metrics and definitions.
- Operational transparency via auditing and monitoring.

13.10.4 Platforms and tooling

- Databricks, Snowflake, Microsoft Fabric and semantic layer tooling.
- Integration with BI platforms (e.g., ThoughtSpot, Power BI) where required.

13.10.5 Constraints and dependencies

- Buyer provides data models, access approvals and SSO/identity integration as needed.
- Accuracy depends on data quality and completeness of business definitions.
- Third-party licences and platform quotas apply.

13.10.6 Supplier type and vendors

Implementation and operational support on Buyer-selected platforms, with additional features and support beyond original supplier offerings where appropriate.

Example vendors/platforms: Databricks, Snowflake, Microsoft Fabric, ThoughtSpot, AWS, Azure, GCP

13.10.7 Notes and service-specific acceptance considerations

Acceptance typically includes measured accuracy/safety tests, access control evidence and operational runbooks.