



Driving Secure Digital Transformation

Data Agent

Enterprise AI Data Bridge

Service Definition Document

January 2026

Data Agent – Enterprise AI Data Bridge

Summary

Data Agent unifies enterprise data access for AI agents, LLMs and automation. It provides governed natural-language access to CRM, ERP, finance, HR and databases across on-premise, private, sovereign and hybrid environments. It removes bespoke integrations, supports multi-agent ecosystems and enforces compliance, auditability and full data sovereignty.

Features

- Natural-language access to enterprise data
- Unified multi-source data federation
- MindsDB-enhanced architecture with governance
- MCP, A2A and REST agent connectivity
- Flexible on-premise or sovereign deployment
- Semantic understanding of business terms
- Deterministic, auditable data interactions
- Multi-agent integration and orchestration support
- Enterprise-grade governance and compliance
- 100+ pre-built enterprise connectors

Benefits

- Removes integration complexity and engineering burden
- Accelerates AI adoption safely and quickly
- Reduces hallucinations with verified source data
- Ensures compliance and full auditability
- Protects data sovereignty and residency
- Enables consistent multi-agent workflows
- Unifies systems and resolves schema conflicts
- Supports global multi-language access
- Delivers rapid ROI through faster integrations
- Scales as AI strategy evolves

Service Description

Overview

Data Agent is a secure, governed access layer connecting enterprise systems to AI agents, LLMs and automation platforms. Rather than building bespoke integrations across CRM, ERP, HR, finance, operational systems and databases, organisations use the Data Agent as a single, policy-aware gateway to enterprise data.

AI models and automation workflows communicate with the Data Agent using natural language or machine interfaces. The platform identifies data sources, enforces authentication and policy, resolves business terminology and returns consistent, structured results. This removes system-specific complexity and ensures that all AI workflows rely on trusted, reconciled and compliant data.

The Data Agent federates data across CRMs, finance suites, ERPs, HR systems, APIs, warehouses and custom applications, providing a unified view across previously siloed systems. Its semantic engine interprets synonyms, abbreviations, domain language and multilingual queries, enabling accurate retrieval for both technical and non-technical users.

Optimised for agentic AI architectures, it integrates with RAG agents, SQL agents and orchestration platforms, combining structured, unstructured and multi-source data to support analytics, reporting, operational insights and compliance workflows.

Security and sovereignty are central to the platform. It supports offline, on-premise, private and sovereign deployments, with GDPR-aligned controls, encryption, auditability and policy enforcement. Every response is deterministic, source-backed and traceable, reducing AI risk and strengthening governance.

The platform integrates seamlessly with chat platforms, automation frameworks and developer tools through REST, MCP and A2A interfaces. With full observability, audit trails and deterministic execution, the Data Agent forms the foundation for enterprise-scale, compliant AI transformation.

Architecture

SQL The platform uses a layered architecture beginning with a Secure Access Layer governing authentication, RBAC/ABAC, encryption and audit controls. The Interface Layer provides REST APIs, MCP channels and A2A machine interfaces for conversational AI, automation pipelines and agent frameworks.

The Data Integration and Federation Engine connects to CRMs, ERPs, HR systems, finance applications, warehouses, APIs and databases, applying policy and normalising responses. The

Context Modelling Engine captures semantic relationships, business terminology and synonyms to interpret natural-language queries with precision.

A Deterministic Execution Pipeline ensures reproducible and compliant execution of all operations across APIs, databases and federated sources. The Normalisation and Fusion Layer reconciles schema variations, resolves conflicts and prepares multi-source outputs for humans, LLMs and agents.

An Audit and Observability Plane captures full lineage, decision paths and operational metrics. The Deployment Layer supports on-premise, private cloud, sovereign cloud, hybrid and offline models with containerised, Kubernetes-ready components for scalable and resilient operation.

Incident Management

Viewdeck classifies incidents raised at its service desk using the following P1 – P3 priority. Incidents raised with the Service Desk will be triaged and responded to accordingly in priority order.

- P1 Total loss of service to all users.
- P2 Some loss of critical service for some or all users.
- P3 Limited loss of service or work around possible limiting loss experienced.

Priority	Response type (during standard business hours)	Target resolution (Mon-Fri 9-5)
P1	30 minute response, sustained effort using all necessary and available resources until service is restored.	4 hours
P2	4 hour response to assess the situation, staff may be interrupted and taken away from lower priority jobs.	1 working day
P3	One working day response using standard procedures and operating within the normal frameworks.	3 working days