CATALYST Service Description

Framework-ready Schedule

Version: 1.1

Date: January 2026

Supplier: Adarga Limited

Service at a Glance

Service name	CATALYST
Service type	AI-enabled data fusion and decision support platform (per-tenant managed service)
Primary users	Analysts, investigators, operations teams, and data/engineering teams supporting mission outcomes
Core capabilities	Fusion engine, structured store, ingestion pipeline, API, audit
Delivery model	Per-tenant subscription with banded capacity and optional deployment modules
Typical use cases	Cross-source entity and event fusion, triage and prioritisation, investigations support, and operational dashboards

Overview

CATALYST is a per-tenant platform for ingesting, normalising, fusing and serving multi-source information to support analysis and operational decision-making. It combines scalable ingestion, a structured store, and a fusion engine to link entities, events and content across datasets, while providing strong auditability for government use.

CATALYST is designed to integrate with existing customer data stores and systems. Deployments can start with a standard configuration and extend through optional delivery modules where bespoke integration or new ingestion modalities are required.

Core Service Components

Fusion Engine

Performs cross-source fusion and enrichment to connect related entities, events and artefacts. Supports configurable fusion rules and workflows aligned to customer operating models.

Structured Store

Stores normalised, queryable representations of fused data to enable reliable retrieval, analytics and downstream applications. Designed to support traceability from fused outputs back to source evidence.

Ingestion Pipeline

Ingests and processes data from customer sources via file, API, and system integration patterns. Provides scheduling, validation, transformation and monitoring of ingestion runs.

API

Provides programmatic access for customer applications, workflows and integrations. Supports controlled access patterns to enable secure use by downstream systems.

Audit

Captures auditable records of access and key platform actions. Designed to support governance, assurance and operational oversight.

Deployment Model

CATALYST is supplied per tenant. A tenant is an isolated instance of the service with dedicated configuration and logical separation of data and access.

Tenants can hosted on-premises or in the cloud and are agnostic to cloud provider.

Environments

Tenants can be delivered as a single production environment or as multiple environments (for example: development, test, and production) as required by the customer.

Where multiple environments are used, environments are configured to support safe change management and release processes.

Capacity Dimensions (Commercial)

CATALYST recurring pricing is typically shaped by three measurable dimensions:

- Data under management (decimal TB) per tenant. Derived outputs are excluded from quota.
- Parallel processing pipelines per tenant (one included; additional pipelines are optional).
- Additional ingestion modalities beyond the standard set (unstructured documents, audio, video, image).

Implementation and Onboarding

A typical deployment includes the following activities, delivered as fixed packages and/or via SFIA day rates depending on scope:

- Installation and baseline configuration (including standard hardening and smoke testing).
- Connection to customer data stores (credentials, network paths, scheduling and monitoring configuration).
- Integration to customer systems (for example identity and access management, logging/monitoring, and operational tooling).
- Operational readiness activities (runbooks, handover, and agreed support processes).

Acceptance and Handover

Acceptance criteria are agreed at call-off and typically cover functional validation, security controls evidence (as applicable), operational monitoring, and successful end-to-end data flows for the agreed scope.

Optional Delivery Modules (Bespoke)

Where required, CATALYST can be extended through optional delivery modules. These are delivered as fixed-price packages where scope is bounded, or on a time and materials basis using the SFIA rate card where requirements are uncertain or discovery is needed.

- Bespoke application build (customer-facing workflows or specialist user interfaces).
- Bespoke data pipeline build (new or complex ingestion and transformation logic).
- Bespoke application integration (non-standard integration patterns or complex interface constraints).

- Bespoke modality fusion/extension (new source-data classes requiring new extraction, normalisation or enrichment logic).

Additional Ingestion Modality Definition

An additional ingestion modality is a distinct, non-standard ingestion and normalisation pathway implemented within a tenant to acquire and process a new class of source data that is not included in the standard modality set.

Standard modalities included: unstructured documents, audio, video, image.

A modality is considered distinct where it requires one or more of:

- a new parsing or extraction approach
- a new schema mapping or normalisation approach
- a new enrichment strategy specific to that data class
- a new ingestion interface pattern (for example API polling, webhook/event-driven, bulk file drop, database extract, warehouse connector)

Not a modality (for avoidance of doubt):

- a new endpoint, table, or dataset within an already-enabled modality
- configuration of credentials, endpoint URLs, schedules, or source selection where the underlying modality is already enabled
- downstream applications, dashboards, analytics, and business logic built on top of ingested data

Security, Privacy and Governance

CATALYST is designed for government-aligned security and assurance practices. Specific controls and evidence are agreed at call-off, based on the target environment and customer requirements.

Typical security and governance features and practices include:

- Tenant isolation and controlled administrative access.
- Audit logging for key actions and access patterns.
- Role-based access control aligned to customer operating models.
- Secure integration patterns with customer identity, logging and monitoring services (where applicable).
- Data handling aligned to customer policies, including retention and deletion requirements where agreed.

Data Protection

CATALYST processes customer data on behalf of the customer. Data protection roles (controller/processor), handling instructions, and any required impact assessments are agreed at call-off.

Service Management and Support

Support arrangements are selected at call-off. A standard support offer is available, with optional enhanced support tiers.

Support typically includes:

- Incident management and service request handling.
- Platform monitoring and operational reporting (as applicable to the chosen delivery model).
- Release and change management aligned to agreed environments and maintenance windows.
- Knowledge transfer and runbook documentation for operational teams.

Service Levels

Service levels (for example support hours, response targets, and availability) are agreed at call-off and documented in the order form or service schedule.

Roles and Responsibilities

Area	Supplier responsibilities (typical)	Customer responsibilities (typical)
Service delivery	Configure and operate CATALYST as agreed; deliver deployment activities; provide documentation and handover.	Provide stakeholders and SMEs; agree acceptance criteria; support governance and approvals.
Access and identity	Integrate with agreed identity controls where applicable; implement platform roles.	Provide identity service inputs, groups/roles mapping, and access approvals.
Data sources	Configure connectors and ingestion runs for	Provide access to data sources, network routes, credentials, and data

	agreed sources; monitor ingestion.	dictionaries where needed.
Security and assurance	Provide agreed security controls evidence and support assurance activities within scope.	Define assurance requirements, classification constraints, and provide security approvals.

Pricing and Ordering

Pricing is provided in the associated Pricing Schedule and SFIA Rate Card. Recurring charges are per tenant per year, with banded capacity and optional add-ons. One-off deployment services and optional delivery modules are ordered as fixed packages and/or via SFIA day rates.

To order CATALYST, the buyer typically selects:

- Core Platform Subscription (per tenant, per year)
- Data Under Management band (per tenant, per year)
- Any additional pipelines (per additional pipeline, per tenant, per year)
- Any additional modalities (per additional modality, per tenant, per year)
- Any required deployment services and optional delivery modules

Assumptions and Constraints

Assumptions (typical):

- Customer provides timely access to environments, networks, and data sources required for delivery.
- Customer provides operational and security stakeholders for approvals and assurance activities.
- Data readiness (formats, dictionaries, access permissions) is sufficient to support the agreed scope and timelines.

Constraints (typical):

- Non-standard hosting, unusual assurance requirements, or highly bespoke integrations may require additional discovery and delivery effort.
- Performance and throughput are influenced by source system characteristics and data quality; final sizing is confirmed during delivery.

Exit and Data Portability

At contract end or on request (subject to agreed terms), the supplier will support export of customer data and configuration artefacts where applicable, and will delete customer data from the tenant in line with agreed retention and deletion requirements.