

SERVICE DESCRIPTION FOR CATALYST

1. Overview of the Platform

Catalyst is a **Platform as a Service (PaaS)** solution designed to empower businesses to integrate artificial intelligence into their applications efficiently. It facilitates the creation of intelligent, conversational AI functionalities without necessitating extensive AI-coding expertise. Catalyst seamlessly integrates with existing platforms to deliver personalised experiences, optimise user engagement, and streamline business operations.

2. Key Features & Differentiators

- 2.1. **Fast, Effortless Integration.** Deploy complex AI capabilities in just 2–3 months —far faster than the 6+ month industry norm—without the disruption of API changes or infrastructure overhauls.
- 2.2. **Accurate and Hallucination Free.** Specialised, internal algorithms and automatic tests ensure responses are stable and free of hallucinations.
- 2.3. **Cost Effective.** The Catalyst platform leverages machine learning to learn your API and on-going optimisation to minimise AI-server usage, meaning your setup and usage costs remain as low as possible.
- 2.4. **Intuitive Development Environment.** Benefit from comprehensive libraries and user-friendly development tools that facilitate a smooth and efficient AI integration process.
- 2.5. **Simplified Maintenance.** Eliminate the need for specialised in-house AI expertise or continuous reliance on external contractors. Catalyst allows traditional developers to handle most maintenance and updates single-handedly.
- 2.6. **Dedicated Expert Assistance.** Access to a team of AI specialists readily available to provide support and guidance, helping navigate challenges and optimise AI integration strategies.

3. Core Service Features

Catalyst offers a comprehensive set of features, including:

- 3.1. **Custom AI Interfaces.** Utilise a library of AI-enabled components to add AI-driven features to existing software, enhancing user experiences and expanding application capabilities.
- 3.2. **White-Label Solutions.** Acquire a white-label version of AI solutions, adaptable to meet specific business needs while ensuring a cohesive and branded user experience.
- 3.3. **Partnership Opportunities.** Collaborate with Quik-AI to introduce AI products to without the financial obligations of rebranding or feature customisation.

4. Service Models

Catalyst is available in multiple deployment models to suit various operational needs:

- 4.1. **Custom Integration Model.** Tailored AI integration into existing software solutions, providing a bespoke AI experience.
- 4.2. **White-Label Model.** Pre-packaged AI solutions ready for branding and deployment.
- 4.3. **Partnership Model.** Collaborative approach to offer AI functionalities to end-users through strategic partnerships.

5. Infrastructure & Technology Stack

Catalyst operates on a robust **AWS-based** cloud infrastructure, leveraging:

- 5.1. **Virtualisation Technology.** AWS EC2 for compute workloads.
- 5.2. **Database.** PostgreSQL + Weaviate for structured and unstructured data management.
- 5.3. **Backend Technologies.** .NET, Java, Python.
- 5.4. **Frontend & Plugins.** Angular for CMS and components library, React for plugin architecture.
- 5.5. **Security & Compliance.** Regular audits in alignment with **Cyber Essentials certification**.
- 5.6. **AI Capabilities.** AWS Bedrock and **Cloud Sonnet** for NLP and ML applications.

6. Information Assurance & Compliance

- 6.1. **Certifications.** Compliance with **GDPR** and **Cyber Essentials**.
- 6.2. **Data Centre Security.** Hosted within highly secure **AWS London** data centres with geographical redundancy.
- 6.3. **Business Continuity & Disaster Recovery (BC/DR).** Active/Active data centres with disaster recovery sites available.

7. High Availability & Backups

Catalyst is architected for **high availability**, featuring:

- 7.1. **Redundancy & Failover.** Elimination of single points of failure with load balancing.
- 7.2. **Backup & Snapshot Management.** Configurable backup policies with **user-defined snapshots**.
- 7.3. **Recovery Time & Objectives (RTO/RPO).** Platform restoration within **1–4 hours** in case of failure.

8. Support & Monitoring

Catalyst provides **24x7x365 technical support**, including:

- 8.1. **Dedicated Engineering Team.** Proactive issue resolution.
- 8.2. **Unlimited Web and Telephone Support.** Ensuring continuous assistance.
- 8.3. **Guaranteed Response Times.**
- 8.4. **Critical Issues.** 15-minute response, 1-hour target resolution.
- 8.5. **Non-Service Affecting Issues.** 1-hour response, 24-hour target resolution.
- 8.6. **Custom SLAs.** Tailored service matrices to prioritise response times.

9. Security Measures

To ensure operational security, Catalyst implements:

- 9.1. **Multi-layered Security.** Firewalls, VPN configurations, network segregation.
- 9.2. **Automated Threat Detection.** Intrusion detection and vulnerability assessments.
- 9.3. **Continuous Security Testing.**
- 9.4. **Static Code Analysis.** At CI/CD pipeline level.
- 9.5. **Automated Unit Testing & Manual QA.** Ensuring code integrity.
- 9.6. **Beta Environment Validation.** Before deployment.
- 9.7. **Quarantine Period.** For production rollout.

10. Service Level Agreements (SLAs)

Catalyst provides a **99.95% uptime guarantee**. Compensation mechanisms are in place for service disruptions exceeding SLA thresholds.

11. Onboarding

- 11.1. The onboarding process is designed to establish a secure, scalable, and context-appropriate AI PaaS environment, enabling the Customer to operationalise generative AI capabilities using vector search and Anthropic foundation models via Amazon Bedrock.
- 11.2. During onboarding, the Service Provider will collaborate with the Customer's technical stakeholders to develop an implementation plan tailored to the target use case(s). This typically includes:
- 11.3. Provisioning of isolated compute and storage resources within a multi-tenant or single-tenant architecture, including managed vector databases (e.g., Amazon OpenSearch with k-NN or Pinecone, where supported)
- 11.4. Configuration of access roles and policies for Amazon Bedrock with scoped usage of Anthropic models (e.g., Claude 3 family), aligned to Customer-specific use and budget constraints

- 11.5. Integration of embedding pipelines for text-to-vector transformations, supporting both real-time and batch ingestion into the vector store
- 11.6. Establishment of observability components (e.g., CloudWatch metrics, log forwarding, performance tracing) for inference endpoints and vector queries
- 11.7. Optional setup of retrieval-augmented generation (RAG) pipelines and prompt chaining logic using orchestrators such as LangChain, AWS Lambda, or Step Functions
- 11.8. Security, performance, and governance requirements — including model latency expectations, API concurrency limits, and prompt auditing policies — are validated during a technical alignment session prior to go-live. All infrastructure components are provisioned using IaC (Infrastructure as Code) frameworks, ensuring repeatability and version control.

12. Offboarding

- 12.1. Offboarding ensures a secure and comprehensive disengagement from the PaaS platform, with all Customer-specific data, credentials, and compute resources decommissioned in accordance with contractual and compliance requirements.
- 12.2. The offboarding process includes:
 - 12.2.1. Revocation of access to Anthropic model endpoints via Amazon Bedrock, including deletion of any associated configuration profiles, invocation logs, or prompt templates (as applicable)
 - 12.2.2. Secure deprovisioning of Customer-specific services such as embedding functions, inference endpoints, data pipelines, and staging environments
 - 12.2.3. Validation and certification of data deletion
 - 12.2.4. Delivery of an offboarding artefact (or closure report) detailing actions taken, with logs or audit trails retained where required for compliance (e.g., GDPR, ISO 27001)
 - 12.2.5. All onboarding and offboarding activities are managed through a change-controlled process and tracked via the PaaS management console or a shared service desk. Handover documentation, support during transition, and post-offboarding access (if required) will be defined during the disengagement planning phase.

13. Service Transition & Migration Approach

Available and dependent on client requirements

14. Additional Services & Options

No additional optional services currently available.