G-Cloud 15

Service	Statamic CMS & AI Safeguarding
Document	Service Definition

CONTENTS

SERVICE CONCEPT

MODERN CMS REQUIREMENTS

Public sector organisations rely on digital services as primary channels for providing authoritative information, supporting access to services, and meeting statutory obligations around accessibility, inclusion, and transparency. These services must be usable by the widest possible audience, regardless of accessibility needs, language, or digital confidence, and must operate reliably under public scrutiny.

Modern public digital delivery is no longer limited to a single website. Information is expected to be consistent and accessible across multiple channels, including websites, search, assisted digital services, and emerging interfaces. This requires content to be managed once and reused safely and consistently, rather than duplicated across systems or sites. A modern CMS must therefore support multi-site and multi-channel delivery from a single, authoritative content base, while maintaining governance and control.

This service provides a modern content management system designed specifically for that context. It supports accessible, multilingual, and user-centred public digital services, and enables organisations to manage large and complex content estates efficiently. The CMS is designed for use by non-specialist publishers, while embedding the structures and controls required to maintain consistency, compliance, and quality over time.

EMERGING USE OF AI

The service also recognises that many public organisations operate in domains where information is complex and user need is uncertain. This includes organisations with significant public communication responsibilities, such as local authorities, regulators, and public agencies that support people navigating complex legal, regulatory, or eligibility-based processes. In these contexts, websites often become a primary route for both information provision and enquiry handling.

AI-assisted capabilities within the service are intended to support this reality. They can help users find relevant information, understand guidance, and orient themselves within complex subject areas, reducing avoidable one-to-one enquiries while maintaining accuracy and trust. These capabilities are particularly suited to organisations that must balance broad public communication with the cost and capacity constraints of individual support.

At the same time, the service acknowledges that the use of AI in public-facing digital services introduces material risks if left unmanaged. AI within the service is therefore operated within a defined governance and assurance framework, ensuring its use remains bounded, transparent, auditable, and subject to human oversight. AI is treated as an assistive capability operating within the CMS, not as an autonomous system or decision-maker.

The service is intended for public organisations that need to handle complex public communication efficiently, improve access to information across multiple channels, and manage high volumes of public interaction, while maintaining accountability, compliance, and public trust.

SERVICE OVERVIEW

This service provides a modern, headless content management system designed to support the delivery of public digital services across multiple channels. The CMS acts as the authoritative source of public information, enabling organisations to manage content once and deliver it consistently across websites, search services, assisted digital channels, and emerging interfaces.

The CMS is API-centric and designed for multi-site and multi-channel use. This allows organisations to operate multiple public-facing websites or services from a shared content base, reducing duplication and inconsistency while supporting accessibility, multilingual delivery, and structured content reuse. The service is intended to support high-performance public communication at scale, rather than isolated or single-channel websites.

STATAMIC PRO CMS

Statamic Pro is a modern content management system (CMS) built on the PHP framework Laravel. Unlike many traditional CMS platforms that rely on relational databases, Statamic uses a *flat-file by default* approach for content storage, enabling content to be stored in human-readable formats such as Markdown, YAML, and JSON while still providing the option of database storage when needed. This architecture improves performance, simplifies deployment and backup, and increases security by reducing the attack surface associated with database-driven systems.

Statamic Pro delivers a combination of developer flexibility and editorial usability. It includes a modern, intuitive control panel and supports over 40 fieldtypes for flexible content modeling, enabling structured and customizable content schemas. The platform's administrative interface features live preview, drag-and-drop navigation building, asset management, revision history, workflow controls, and form management — features that are part of the Pro edition and extend beyond the base CMS.

The system supports both *traditional full-stack* and *headless* implementations via REST API and GraphQL, making it suitable for multi-channel content delivery and integration with JavaScript frameworks or decoupled front-ends. Multi-site and multi-language support are included, allowing multiple websites or linguistic variants to be managed within a single installation.

Because Statamic is packaged as a Laravel application, developers can extend and customise the CMS using Laravel's ecosystem, including integration with middleware, models, and services. Native Git integration enables version control workflows, which simplify development, collaboration, and deployment processes.

Statamic's architecture and feature set position it as a flexible, scalable CMS that balances developer control with editorial usability. Its flat-file model, headless capabilities, and extensibility make it a suitable choice for modern digital services that require performance, custom content structures, and future-proof architectures.

AI GOVERNANCE FRAMEWORK

Alongside the CMS, the service includes an embedded AI governance framework that enables the safe and controlled use of AI-assisted capabilities within the CMS environment. This framework exists to allow organisations to realise the operational benefits of AI—such as improved usability, more efficient content management, and better access to complex information—while managing the risks associated with AI use in public-facing digital services.

The governance framework is designed to ensure that AI-assisted capabilities are bounded, transparent, and subject to human oversight. It is not tied to a single AI tool, vendor, or model, and can accommodate change over time as AI technologies evolve. This allows organisations to adopt AI where it adds value, without introducing unmanaged risk or dependency.

Alignment with organisational AI policies, security requirements, and assurance practices is a natural outcome of this approach. The framework provides the controls and operating model needed to support safe AI use in contexts where accuracy, accountability, and public trust are critical.

Together, the CMS and AI governance framework provide a single, coherent service. The CMS establishes the structure, authority, and consistency of public information, while the governance framework ensures that any AI-assisted use of that information is managed appropriately for public sector contexts. The service does not replace transactional systems, case management platforms, or back-office tools; it is designed to operate as the public-facing front door to existing organisational systems.

SECTION 3 — TECHNOLOGY ARCHITECTURE

OVERVIEW

The service is built on a modern content management architecture designed to support long-lived, accessible, and multi-channel public digital services. The architecture separates content management from presentation and delivery, allowing organisations to manage public information as a governed asset that can be reused and adapted over time.

This approach supports flexibility, resilience, and efficiency, while reducing duplication and operational overhead across complex digital estates.

HEADLESS CONTENT MANAGEMENT

The service uses a headless content management approach, where content is created, governed, and stored independently of how it is presented to users. Content is managed centrally and delivered to websites and other interfaces as required, rather than being tightly coupled to a single front-end implementation.

This allows public organisations to change or redesign websites, introduce new interfaces, or adapt to evolving user needs without restructuring content or governance models.

What this means for clients

- Content remains usable and authoritative across redesigns and service changes.
- Presentation changes do not require content migration or re-authoring.
- Content can be treated as a long-term organisational asset rather than a by-product of a specific website.

API-CENTRIC ARCHITECTURE

The service is API-centric by design. Content and structured data are exposed through APIs as the primary means of delivery, rather than being embedded directly into presentation layers.

This enables the CMS to operate as part of a wider digital service architecture, supporting integration with other systems and enabling content to be reused across multiple services and channels.

What this means for clients

- The CMS can integrate cleanly with existing systems such as payments, case management, and CRM platforms.
- Digital services can be composed from multiple systems without duplicating content or logic.
- The architecture supports future change without locking organisations into a single delivery model.

MULTI-SITE SUPPORT AND DIGITAL ESTATE RATIONALISATION

The architecture supports the operation of multiple public-facing websites and services from a single CMS instance. Sites can share content models, templates, media assets, and governance controls, while still allowing for appropriate variation where required.

This enables organisations to rationalise fragmented digital estates and reduce the cost and complexity associated with maintaining multiple CMS installations.

What this means for clients

- Multiple websites and service sites can be managed consistently from one system.
- Governance, accessibility, and security controls are applied uniformly.
- Operational overhead is reduced by minimising duplicated platforms and processes.
- Organisational change can be accommodated without replatforming or site proliferation.

OMNICHANNEL CONTENT DELIVERY

The architecture supports delivery of public information across multiple digital channels from a single, governed content base. Content can be delivered to websites, search interfaces, assisted digital services, and conversational or AI-assisted interfaces without being duplicated or re-authored.

This ensures that public information remains consistent and authoritative regardless of how users choose to access it.

What this means for clients

- Users receive consistent information across channels.
- Accessibility and governance standards apply equally across delivery mechanisms.
- New channels can be introduced without restructuring content or compromising control.

FUTUREPROOF

By separating content, presentation, integration, and delivery concerns, the architecture supports long-term sustainability of public digital services. Each layer can evolve independently within defined boundaries, reducing risk and avoiding tightly coupled systems that are difficult to change.

This architectural approach supports incremental improvement rather than periodic re-platforming.

What this means for clients

- Reduced risk of large-scale technical debt accumulation.
- Greater ability to adapt services over time within constrained budgets.

- A platform that supports continuous improvement rather than disruptive replacement.
- Enablement of longer term omni-channel and sophisticated digital comms strategies.

FACILITIES FOR ADMIN USERS

The service provides a set of administrative facilities designed to support staff who are responsible for maintaining public digital services, but who are not specialist web publishers or developers. These facilities focus on enabling organisations to manage content, structure, and governance independently, without routine reliance on technical support, while maintaining appropriate control and assurance.

CONTENT CREATION AND MANAGEMENT

The service supports structured content creation through defined content types, fields, and reusable components. Content is created and managed within controlled structures rather than free-form pages, separating content from presentation and reducing the risk of layout or design issues.

Content can be drafted, edited, previewed, and updated through an administrative interface designed for routine use. Changes to text, images, documents, and page structure can be made without modifying code or templates.

What this means for administrative users

- Staff can update pages, guidance, and service information themselves without returning to a developer.
- Subject matter experts can focus on accuracy and clarity rather than technical constraints.
- Routine changes can be made quickly, reducing delays between policy or service change and public communication.
- Confidence in making updates reduces avoidance behaviours and reliance on workarounds such as PDFs or email requests.

PUBLISHING WORKFLOWS AND COLLABORATION

The service supports multi-user content workflows, allowing multiple contributors to draft, review, and approve content before publication. Publishing states make it clear whether content is in draft, under review, scheduled, or live.

Preview functionality allows users to see how content will appear before it is published. Publishing actions are controlled through permissions, ensuring that only authorised users can make content public.

What this means for administrative users

- Content can be prepared collaboratively without confusion over ownership or status.
- Reviews and approvals are handled within the system rather than through email or offline processes.
- The risk of accidental publication is reduced.
- Governance processes are followed because they are built into everyday work, not imposed externally.

GOVERNANCE, VERSION CONTROL, AND AUDITABILITY

The service maintains a full history of changes to content, including who made changes and when. Previous versions of content can be viewed and restored where necessary. Publishing activity is logged to support accountability and review.

These controls are applied automatically as part of normal system use rather than requiring additional administrative effort.

What this means for administrative users

- Mistakes can be corrected quickly without technical escalation.
- Teams can work confidently, knowing changes are reversible.
- Organisations can demonstrate accountability and traceability if content is challenged.
- Governance does not slow down publishing or centralise control unnecessarily.

MEDIA AND ASSET MANAGEMENT

The service includes a central media library for managing images, documents, and other assets. Assets can be reused across multiple pages and sites, with metadata such as titles and alternative text managed consistently.

This supports accessibility requirements and reduces duplication of files across the content estate.

What this means for administrative users

- Assets are easier to find and reuse, saving time.
- Accessibility requirements are easier to meet and maintain.
- Teams avoid creating multiple inconsistent versions of the same document or image.
- Content quality improves without increasing administrative burden.

AUTOMATION AND QUALITY GUARDRAILS

The service includes automation and guardrails that support content quality and consistency. This includes support for search optimisation, structured metadata, and content validation derived from content models and templates.

Automation is used to prompt, guide, and enforce consistency, not to make publishing decisions automatically.

What this means for administrative users

- Editors are supported to publish content that is findable and well-structured without specialist knowledge.
- Common errors are prevented before publication.
- Quality standards are maintained even when publishing is decentralised.
- Staff spend less time checking compliance and more time improving content.

LANGUAGE & LOCALISATION

The service supports the management of content in multiple languages or locales, with relationships maintained between source and translated content. Language variants can be governed independently while remaining aligned.

Responsibilities for translation and review can be assigned appropriately.

What this means for administrative users

- Statutory language obligations can be met within a single system.
- Translated content does not drift or become outdated unnoticed.
- Language responsibilities are clear and manageable.
- Multilingual publishing does not require separate sites or duplicate processes.

FORMS AND SURVEYS

The service supports the creation and management of forms for collecting structured information from public users. Forms can be used to support enquiries, feedback, registrations, and other data capture needs associated with public digital services and we support internal Statamic Forms and other third party forms systems.

Forms can be configured and managed by administrative users within the CMS, allowing organisations to create and update forms without developer involvement. Validation, structured data capture, and routing of submissions are supported to enable integration with downstream systems.

Where organisations prefer to use external form or survey platforms, the service also supports integration with third-party form services. This allows clients to continue using established tools for surveys or specialist data collection while maintaining a consistent public-facing experience through the CMS.

What this means for administrative users

- Forms can be created and updated internally without relying on developers for routine changes.
- Different approaches to data capture can be supported within the same service, depending on organisational needs.
- Existing investments in external form or survey tools can be retained.
- Governance and control over how data is collected and routed is maintained.

INDEPENDENCE FROM DEVELOPERS

The service is designed so that the majority of day-to-day content and publishing activity can be carried out by internal teams. Developers are typically only required for substantial structural changes, new integrations, or changes to underlying templates, not for routine updates.

What this means for administrative users

- Organisations retain control over their own content.
- Ongoing costs are reduced by limiting developer involvement to genuinely technical work.
- Publishing teams can respond quickly to change.
- Digital services remain sustainable over time without continual external dependency.

FOR PUBLIC USERS

The service provides facilities that support members of the public in accessing information and engaging with public services through digital channels. These facilities are designed to ensure that public information is accessible, understandable, and findable for users with a wide range of needs, abilities, and levels of digital confidence.

ACCESSIBILITY AND INCLUSIVE ACCESS

The service is designed to support accessible public access by default. Content is delivered through structured templates and patterns that support compliance with accessibility standards, ensuring that public users can access information regardless of disability, assistive technology, or interaction method.

Accessibility is treated as a foundational requirement rather than a retrofit, supporting consistent use across pages, services, and channels as content evolves over time.

What this means for public users

- Information can be accessed using assistive technologies.
- Content remains accessible as it is updated and expanded.
- Users are not excluded due to design, layout, or content structure choices.

SEARCH & NAVIGATION

The service supports multiple ways for users to find information, recognising that public users often do not understand organisational structures or service boundaries. Content is structured to support clear navigation, effective on-site search, and consistent presentation across topics and services.

The service can include enhanced search capabilities, including AI-assisted search (e.g. Alhgolia and , Typesense) to help users locate relevant information using natural language queries rather than exact terminology.

What this means for public users

- Users can locate information without needing to understand how services are organised internally.
- Relevant content can be surfaced even when users are unsure what to search for.
- The effort required to find accurate and authoritative information is reduced.

Note: Specific AI-assisted search approaches will be defined as part of service configuration.

LANGUAGE AND COMMUNICATION SUPPORT

The service supports the delivery of public information in multiple languages and formats, allowing organisations to meet statutory language obligations and communicate effectively with diverse audiences.Content can be managed in multiple languages while remaining aligned with source material, reducing inconsistency and error.

What this means for public users

- Information is available in appropriate languages.

- Translated content remains consistent and reliable.
- Users are less reliant on informal or third-party interpretations.

FEEDBACK AND INTERACTION

The service supports the collection of feedback and structured input from public users, allowing organisations to understand where information is unclear or where users experience difficulty.

Feedback mechanisms can be integrated into digital services without forcing users into complex contact or escalation routes.

What this means for public users

- Users can provide feedback easily.
- Issues can be identified and addressed over time.
- Interaction supports service improvement without unnecessary friction.

USE OF THIRD PARTY GENERATIVE AI

The service supports the use of third-party generative AI tools to provide AI-assisted public user interfaces, such as conversational chat interfaces and AI-supported search and navigation.

These interfaces operate over content managed within the CMS and are constrained to approved purposes and datasets. AI tools are configured to respond only using information that the organisation has chosen to make publicly available through the website or digital service.

The use of generative AI is governed through the embedded AI governance framework, which defines:

- the definition of the use case
- identified risks and mitigations
- algorithmic registrations
- the scope of permitted AI use
- the content and data sources AI tools may access
- thresholds for human escalation or intervention
- logging and audit of AI interactions

AI interfaces are not autonomous decision-making systems. They are used to support information access and guidance only, and do not replace statutory processes, formal advice, or human judgement. The service allows organisations to adopt generative AI interfaces in a controlled manner, while managing risks related to accuracy, bias, inappropriate responses, and public trust.

ACCESSIBILITY

The service supports the delivery of accessible public digital services that meet Web Content Accessibility Guidelines (WCAG) requirements. Accessibility is addressed through a combination of CMS configuration, structured content design, publishing controls, and front-end implementation practices.

The CMS supports the creation of structured, standards-based content that can be rendered in accessible formats across different devices and assistive technologies. Content models, templates, and

publishing workflows are designed to encourage consistent use of headings, labels, alternative text, and semantic structure, reducing reliance on individual authors to manage accessibility manually.

Accessibility considerations are embedded into publishing processes, including preview and review stages, allowing issues to be identified before content is made public.

Subject to client approval, the service can also include deployment of an accessibility support overlay such as **accessiBe**, which provides automated, user-controlled accessibility adjustments. This is offered as a supplementary measure to enhance usability for some users, not as a replacement for WCAG-compliant design and content.

The service does not assume that accessibility can be solved solely through tooling. It supports compliance through a combination of platform capability, implementation approach, and organisational governance.

What this means for public users

Public users are able to access information and services in ways that better suit their individual needs, devices, and assistive technologies.

Accessible content structures and consistent presentation help users:

- navigate content more easily
- understand information without unnecessary barriers
- use assistive technologies effectively

Where an accessibility overlay is deployed, users may also benefit from additional controls such as text resizing, contrast adjustments, and navigation aids, which can improve usability for some individuals. This approach supports inclusive and equitable access to public information and services, helping organisations meet statutory accessibility obligations while improving the overall experience for a wide and diverse public audience.

Safeguards are in place to ensure that accessibility enhancements complement, rather than obscure, underlying compliance with recognised accessibility standards.

INTEGRATIONS

API CENTRIC

The service is designed to integrate with existing organisational systems and external services through an API-centric architecture. It is intended to operate as the public-facing layer for digital services, while allowing core operational systems to remain authoritative.

The CMS exposes structured content and data through APIs, enabling integration with a wide range of internal and third-party systems. This allows organisations to connect public-facing content, forms, and interactions to existing platforms without duplicating data or functionality.

Typical integration scenarios include:

- **Payment services**, where transactional processing is handled by an external payment service provider and the CMS provides the public interface and content.

- **Customer relationship management (CRM) systems**, such as widely used commercial or public-sector CRM platforms, where user interactions or submissions are routed for follow-up or processing.
- **Case management systems**, where the CMS supports information capture and guidance, and cases are created or updated in downstream systems.
- **Email, notification, and communications services**, enabling confirmations, updates, and alerts to be sent as part of digital service journeys.
- **Form and survey platforms**, where organisations choose to use established external tools for surveys or specialist data capture, while maintaining a consistent public-facing experience through the CMS.
- **Business reporting and analytics tools**, allowing insight to be drawn from content usage, interactions, and service performance.
- **Third-party systems and services** that expose APIs, allowing integration to be tailored to organisational needs.

The service does not seek to replace specialist systems such as CRMs, case management platforms, payment engines, or survey platforms. Instead, it provides a structured, governed interface through which public users interact with services, while operational processing and data ownership remain with existing systems.

Integration approaches are designed to be flexible and adaptable, allowing organisations to connect the CMS to their current technology estate and to accommodate change over time as systems evolve.

ANALYTICS

The service supports the collection and exposure of usage and interaction data generated through the CMS and AI-assisted components. This includes data relating to content usage, user interactions, and AI-assisted activity, such as searches, conversational interactions, and escalation events.

The CMS integrates with external analytics and business intelligence tools via APIs, rather than providing a proprietary analytics solution. This allows organisations to use their preferred analytics platforms and reporting approaches.

The service can be connected to commonly used analytics and BI tools, including web analytics platforms and data visualisation services such as Google Analytics, Firebase, Looker, and equivalent tools. Integration approach and configuration are dependent on organisational requirements and existing analytics environments.

Analytics tooling, dashboards, and interpretation of data are not intrinsic to the service and are typically provided through external platforms or separate services.

SECURITY

The service is designed to meet public sector security expectations through a combination of architectural choices, operational controls, and ongoing management. Security is treated as an outcome of how the service is designed and operated over time, rather than as a one-off configuration or compliance exercise.

The approach focuses on reducing exposure, limiting complexity, and ensuring that security controls remain effective throughout the life of the service.

SECURITY BY DESIGN AND REDUCED ATTACK SURFACE

The CMS is designed with a lightweight architecture that minimises unnecessary system components and avoids reliance on large, unmanaged plugin ecosystems. Functionality is delivered through controlled configuration and development rather than through the extensive use of third-party plugins.

Compared to plugin-heavy CMS platforms, this approach significantly reduces the number of external components that must be assessed, patched, and monitored over time. This lowers the overall attack surface and reduces exposure to vulnerabilities introduced through third-party extensions.

What this means for clients

- Reduced security risk associated with unmanaged or outdated plugins.
- Fewer external dependencies requiring ongoing security assessment.
- A more predictable and controllable security posture over the life of the service.

ACCESS CONTROL

Access to the service is governed through role-based access control, aligned to the principle of least privilege. Roles can be defined to separate responsibilities such as content creation, review, approval, and administration.Administrative access is restricted to authorised users and aligned to defined responsibilities.

What this means for clients

- Clear separation between editorial, approval, and administrative functions.
- Reduced risk of accidental or unauthorised changes.
- Easier demonstration of governance and accountability.

AUTHENTICATION AND MULTI-FACTOR AUTHENTICATION

Administrative access to the service is protected using secure authentication mechanisms, including multi-factor authentication. Where required, the service can integrate with organisational identity and access management approaches to align with enterprise authentication policies.

What this means for clients

- Reduced risk of credential compromise.
- Alignment with common public sector security baselines.
- Consistent access controls across internal systems.

VULNERABILITY MANAGEMENT AND PATCHING

The service is operated under a proactive vulnerability management regime. This includes regular monitoring, vulnerability scanning, and timely application of security updates.The platform architecture supports controlled updates and predictable patching without reliance on large numbers of third-party components.

What this means for clients

- Known vulnerabilities are identified and addressed in a timely manner.

- Security maintenance does not rely on ad-hoc or manual intervention.
- Operational burden associated with security management is reduced.

AUDIT, LOGGING, AND TRACEABILITY

Access to the service and changes made within it are logged to support auditability and operational oversight. This includes visibility of who accessed the system and what changes were made.Audit and traceability mechanisms support incident investigation, assurance activity, and compliance review where required.

What this means for clients

- Clear evidence is available to support investigations or assurance requests.
- Accountability for changes is maintained.
- Security controls can be demonstrated, not just asserted.

DATA PROTECTION AND DATA RESIDENCY

Service data is hosted in line with client requirements and applicable data protection legislation, with data residency typically within the UK or EEA. Data is protected using industry-standard encryption for data at rest and in transit.Client data remains under client ownership, and secure processes are in place for data access, extraction, and deletion on service exit.

What this means for clients

- Alignment with statutory data protection obligations.
- Clarity over where data is stored and how it is protected.
- Confidence that data can be securely recovered or removed as required.

SECURITY OVER THE LIFE OF THE SERVICE

The security model is designed to remain effective over multi-year contracts without increasing operational complexity. Security controls are embedded into standard operation rather than relying on exceptional processes or specialist intervention.This supports long-term sustainability and consistent compliance as the service evolves.

What this means for clients

- Security posture does not degrade as the service grows or changes.
- Governance effort remains proportionate.
- Security remains manageable within normal operational capacity.

AI GOVERNANCE (AI GUARDIAN)

The service includes an embedded AI governance layer that supports the safe, controlled, and accountable use of AI-assisted capabilities within public digital services. This governance layer applies consistently across all AI components used within the CMS, including public-facing AI, internal AI tools, and AI services provided by third parties.

AI governance within the service is designed to be operational and repeatable. It provides the structures, controls, and processes required to manage AI use over its full lifecycle, from initial proposal through deployment, monitoring, and ongoing review.

SCOPE OF AI COVERED BY THE GOVERNANCE LAYER

The AI governance layer applies to all AI-assisted capabilities operating within the service, regardless of whether those capabilities are public-facing, internal, or provided by third-party AI services.

Where third-party generative AI models are used, they operate within an Ember-designed retrieval and control component (the Ember RAG module). This component forms part of the governance layer and is responsible for constraining, supervising, and evidencing AI behaviour within the CMS environment.

The Ember RAG module is configured per use case and performs three core governance functions:

First, it constrains AI operation to approved purposes and content. The module restricts the information and context available to the AI, ensuring that outputs are grounded only in authorised content and that AI behaviour remains within the defined scope of the service. This prevents AI from operating beyond its intended use or accessing unapproved data sources.

Second, it governs interaction thresholds and human escalation. The module defines conditions under which AI-assisted interactions must be escalated for human review or intervention. These thresholds are configured to reflect the sensitivity, complexity, or uncertainty of a given use case and ensure that AI does not operate beyond agreed limits of autonomy.

Third, it records and evidences AI activity. The module maintains records of AI interactions and outcomes, providing a basis for audit, assurance, and post hoc review. This supports governance, monitoring, and continuous improvement of AI-assisted capabilities over time.

The same governance layer also applies to internal AI-assisted functions used to support content management activities, such as summarisation, translation, tagging, or quality checks. These internal capabilities are subject to the same constraints, oversight, and evidence requirements as public-facing AI use.

By applying a single governance layer across all AI components, the service ensures that AI-assisted behaviour is controlled, reviewable, and accountable, regardless of where or how AI is used within the CMS.

AI OBJECTIVES AND USE-CASE DEFINITION

Before any AI-assisted capability is deployed within the service, the governance layer supports a structured process to define and approve the intended use of AI. This ensures that AI is introduced deliberately, with clear purpose, defined boundaries, and understood risks.

The service supports early-stage consideration of AI objectives, including the problem AI is intended to address, the benefits it is expected to deliver, and the context in which it will operate. As part of this stage, proposed AI use is reviewed against the organisation's AI policy and governance principles, ensuring alignment with agreed expectations and constraints.

Following this, the service supports evaluation of how AI will be applied within the CMS environment. This includes consideration of appropriate AI components, how they will be constrained and governed, and how they will interact with content, users, and existing systems. Decisions made at this stage define the permitted scope of AI operation.

Where required, the service supports the completion of AI-specific impact assessment activity, distinct from data protection impact assessments. This allows organisations to consider potential impacts on users, fairness, transparency, and public trust, and to identify risks that must be mitigated before deployment.

Outcomes from this stage inform the configuration of the AI governance layer, including:

- approved AI use cases and exclusions
- conditions under which human review or escalation is required
- expectations for monitoring, bias detection, and performance evaluation
- review points to reassess effectiveness, risk, and continued appropriateness over time

Records of decisions and assessments are maintained to support audit, assurance, and accountability.

MONITORING, BIAS DETECTION, AND PERFORMANCE MANAGEMENT

Once AI-assisted capabilities are live, the service supports ongoing monitoring and management to ensure that AI behaviour remains aligned with approved objectives, constraints, and governance expectations.

Monitoring is treated as a routine operational activity rather than a one-off assurance exercise. The service supports continuous visibility of AI behaviour to enable early identification of issues and informed decision-making.

The service provides mechanisms to observe and review AI performance, including patterns of use, accuracy indicators, escalation frequency, and user feedback. This allows organisations to understand how AI is being used in practice and whether it is operating within agreed parameters.

Bias and inclusivity management are supported through structured testing and review. The service supports repeatable bias testing approaches that assess AI outputs across a range of representative user contexts. Results are reviewed to identify disparities in tone, accuracy, or completeness, and to inform corrective action where required.

HUMAN ESCALATION

Human escalation and fallback behaviour are monitored as part of normal operation. Escalation events are recorded and reviewed to confirm that thresholds are functioning as intended and that sensitive or high-risk interactions are appropriately redirected to human support.

The service maintains structured records of AI interactions, escalations, and feedback to support audit, assurance, and retrospective review. These records provide evidence of how AI has behaved over time and support accountability and transparency.

Performance management activity informs ongoing refinement of AI configuration, content, and governance controls. Where monitoring identifies issues or emerging risks, AI behaviour can be adjusted, constrained, or suspended in a controlled manner.

PERIODIC REVIEW AND LIFECYCLE MANAGEMENT

The service supports formal, periodic review of AI-assisted capabilities as part of ongoing performance management. These reviews are intended to assess whether AI use remains appropriate, effective, and aligned with the objectives agreed at the point of approval.

Reviews are typically conducted on a quarterly basis and consider AI performance against the original objectives defined during the AI objectives and use-case definition stage. This includes assessment of whether AI capabilities are delivering the intended benefits, operating within agreed constraints, and continuing to justify their use in the given context.

Periodic review activity considers:

- performance against agreed objectives and success criteria
- findings from monitoring, bias detection, and escalation activity
- identified deficiencies or unintended effects
- opportunities to improve effectiveness, usability, or control
- continued alignment with organisational AI policy and governance expectations, including where policies have evolved since initial deployment

Reviews may result in decisions to maintain the current AI configuration, to modify or enhance AI-assisted capabilities, to impose additional constraints, or to suspend or withdraw AI use where it is no longer appropriate.

The service supports recording outcomes of periodic reviews to provide an auditable record of decisions and to inform future governance activity. This ensures that AI use remains subject to active oversight and does not persist by default once deployed.

HOSTING & SUPPORT

HOSTING OPTIONS

The CMS and AI governance components are hosting-agnostic by design and can be deployed within any suitable modern hosting environment. This includes client-managed infrastructure, third-party managed platforms, or public cloud environments specified by the client. The service does not require a proprietary hosting model and does not impose dependency on Ember-provided infrastructure.

This allows organisations to align hosting arrangements with their own technical standards, security requirements, and procurement preferences.

DEPLOYMENT AND ENVIRONMENT CONFIGURATION

To make the service operational, the CMS and AI governance components must be deployed and configured within a hosting environment. The service includes technical deployment and configuration capability to support installation into a client-specified environment.

Deployment activities are limited to the technical configuration required to operate the service and may include:

- installation of the CMS and AI governance components
- configuration of production and non-production environments
- security and access configuration
- configuration of integrations required for core operation

Deployment does not include delivery of a wider digital project, service design, bespoke development, or content production. Any such activities are outside the scope of this service unless separately procured.

EMBER-PROVIDED HOSTING (OPTIONAL)

Where clients require hosting to be provided as part of the service, Ember offers managed cloud hosting using **Amazon Web Services (AWS)**. Ember-provided hosting is suitable for public digital services at local, regional, and national levels.

Hosting environment

- Deployed within AWS, selected for scalability, reliability, and suitability for public sector workloads
- Configured to support secure operation, resilience, and high availability
- Supports production and non-production environments as required

Data location and residency

- Service data is hosted in AWS data centres located in the UK and/or EEA, as agreed with the client
- Data residency requirements are confirmed during onboarding and reflected in deployment configuration
- All data storage, processing, and backups are managed within the hosting environment

Security and information assurance

Ember-provided hosting operates within an Information Security Management System (ISMS) certified to:

- ISO/IEC 27001
- Cyber Essentials Plus
- NCSC Cloud Security Principles

Security controls include:

- encryption of data at rest and in transit
- role-based access control aligned to least-privilege principles
- multi-factor authentication for administrative access
- segregation of client environments and data
- comprehensive logging and monitoring to support auditability

Resilience, backup, and recovery

- Hosting environments are designed with redundancy and monitoring to support service availability
- Automated backups are performed in line with Ember's business continuity policies
- Recovery Time Objectives (RTO) and Recovery Point Objectives (RPO) are defined to minimise disruption and data loss

Vulnerability management

- Regular internal and external vulnerability scanning, including independent third-party testing
- Identified vulnerabilities are assessed, prioritised, and remediated in accordance with Ember's security policies

SUPPORT, AVAILABILITY, AND SERVICE LEVELS

The service includes support for the CMS and AI governance components, covering incident management, fault resolution, and assistance with operation of the service within the agreed scope.

Service availability

- The service is designed to operate at **99%+ availability**
- A **minimum availability commitment of 99%** is provided as standard where Ember provides hosting

Support hours and access

- Support is provided via Ember's service desk during standard business hours, Monday to Friday
- Support requests can be logged 24/7 via email or a support portal and are triaged by priority

Incident response targets

Support requests are prioritised based on impact and urgency, with the following standard response and resolution targets:

- Priority 1 — Critical impact
  Response within 1 business hour; resolution targeted within 3 business hours
- Priority 2 — Significant impact
  Response within 4 business hours; resolution targeted within 1 business day
- Priority 3 — Minimal impact
  Response within 1 business day; resolution targeted within 4 business days

SCOPE OF SUPPORT

Support applies whether the service is hosted by Ember or deployed into a client-specified environment and covers the CMS and AI governance components only. Support does not include bespoke development, content creation, or delivery of wider digital services unless separately procured.

OFFBOARDING AND SERVICE EXIT

OFFBOARDING PROCESS

Offboarding is initiated following formal notification from the client of their intention to terminate the service. Ember confirms the offboarding plan with the client and coordinates activities to ensure an orderly and secure transition.

The offboarding process includes:

- agreement of timescales and responsibilities
- support for data extraction and transfer
- secure decommissioning of the client's service instance (where hosted by Ember)
- contract and account closure

Where required, Ember provides reasonable assistance to support the client's transition to alternative arrangements.

ACCESS TO DATA ON EXIT

Clients retain access to their service data for the duration of the contract and during the agreed offboarding period.

Ember supports the extraction of client data in standard, machine-readable formats (such as CSV or JSON), covering data generated through use of the service, including:

- website content and media
- configuration data
- AI governance logs and audit records (where applicable)

Guidance and support are provided to ensure data can be accessed and transferred in a usable form.

DATA DELETION AND DISPOSAL

Following confirmation that data extraction has been completed, Ember securely deletes remaining client data from the service in accordance with:

- ISO/IEC 27001-aligned data sanitisation procedures
- applicable data protection legislation

No client data is retained beyond the agreed retention period following service termination.

ACCOUNT CLOSURE AND ACCESS REVOCATION

As part of offboarding:

- access to the service and associated systems is revoked
- administrative credentials are disabled

- any remaining service accounts are closed

Where any credentials or assets have been provided as part of the service, Ember coordinates their secure return or disposal as appropriate.

FEEDBACK AND TRANSITION SUPPORT

As part of offboarding, Ember may request feedback on service delivery and the offboarding process to support continuous improvement. Additional transition support may be provided where agreed.