

Case Assist

AI-Powered Productivity for Case Workers

 hgs | Agent X





Content

1. CASE ASSIST: SOLUTION DESCRIPTION
2. SOLUTION COMPONENTS: FEATURES & FUNCTIONALITY
3. SOLUTION DIAGRAM
4. LEARNING LIFECYCLE
5. EXAMPLES OF AFFECTED USE CASES
6. SOLUTION BENEFITS
7. SERVICE MODEL OPTIONS
8. TECHNOLOGY & INFRASTRUCTURE
9. SECURITY STANDARDS & PENETRATION TESTING

Case Assist: Solution Description



Across public services, teams are required to review large volumes of documentation such as clinical reports, application evidence, correspondence, contracts, financial records, and case files. **Manually interpreting and summarising** this information can be time-consuming, creating backlogs and introducing variation in decision quality, which increases the risk of missed or overlooked details.

Case Assist is an **AI-powered productivity solution** built on HGS's proprietary Agentic AI platform (Agent X). It is designed to streamline this effort by automatically reading, structuring, and extracting the key points from all ingested documents, mapping them to the relevant assessment or decision criteria, and producing clear, traceable summaries for staff to review. We're already helping to reduce manual workload, improve consistency and auditability, and enable teams to focus on higher-value judgement and service delivery.

The solution is hosted in a **UK based, ISO-certified cloud environment**, and uses advanced retrieval-augmented generation (RAG), and human-in-the-loop validation to ensure all outputs are grounded in source evidence and meet the standards of Government, NHS, and Local Authority services.



Intelligent case review

The Case Assist System converts unstructured data documents into a validated, structured format. It then produces AI-driven contextual outputs that help Case Workers more efficiently review, refine, and generate a final report in the format of their choice.

All documents are **securely captured and encrypted** from the moment of upload and retained only during processing.

A screenshot of a web application titled "Product Tour of CHC Plus" in the browser's title bar. The application header shows the "hgs | Agent X" logo and a user profile "ST". The main content area is titled "CHC Plus" and includes a "+ Create Client" button. Below this, a search bar contains "BP" and a dropdown arrow. The case details for "Brad Pitt" are displayed in a grid: "NHS Patient ID: 322 223 2131", "Date of Birth: 02 February 1992, 33", "RSC ASC Number: AS132313", "Address: Toronto", "Gender: female", "GP Name: Nima", and "GP Address: Toronto". A "Case Details" section is collapsed. A progress bar shows four steps: "Create Case" (checked), "Upload Documents" (checked), "Create Summary" (checked), and "Generate Checklist" (4). Below the progress bar, a "Documents" tab is active, showing a list with one item: "MSR_Sample_for_Reapplication_011717.pdf" by "stuart.staples@teamhgs.com" on "13/08/2025". An "Upload Document" button is at the bottom.


Solution Components – Features & Functionality



Case Assist is composed of five interoperable components that transform unstructured information into auditable, high-quality, review-ready outputs. The components work together with asynchronous flows designed in at the appropriate stages, enabling a sophisticated set of workflows for the end-user.

Components	Description
Vector Database	The implementation utilises an HGS hosted vector database solution for storing and retrieving document embeddings. The vector database offers the flexibility of both cloud and self-hosted deployments while providing essential features for healthcare applications: • Meta-data filtering enables classification by document type and use-case pathway • Hybrid search capabilities combine semantic similarity with keyword matching • Configurable retrieval depth balances latency against quality • Strong deletion controls ensure GDPR compliance • Asynchronous operations support improved system responsiveness
Document Parser	Document Parser serves as the primary document processing engine, transforming unstructured (or structured) documents into standardised formats: • Advanced structural understanding preserves document layout and hierarchy • Local execution option maintains privacy for sensitive information • OCR capabilities handle handwritten notes and scanned documents • Multi-format support processes varied document types (PDF, DOCX, XLSX, images) • Semantic chunking optimised for use-case specific terminology and contexts
Embedding Models	The system implements a dual-model approach for vector embeddings.
Orchestration Framework	It provides the foundation for the retrieval and processing pipeline: • Flexible architecture optimised for RAG (Retrieval Augmented Generation) • Advanced chunking strategies tailored to use-case documentation • Citation tracking maintains source references for accountability • Support for streaming responses reduces perceived latency • Mature integration with embedding and language models
Model Fine-tuning Management	The system includes capabilities for continuous improvement through: • Initial deployment using zero-shot and few-shot prompting • Progressive fine-tuning as approved output examples accumulate • Targeted minimum of 200+ examples per use-case pathway • Performance evaluation framework comparing against base models

Case Assist Solution Diagram



Technical Architecture of the Solution

Case Assist thinking and learning lifecycle



Our solution applies a structured document-processing workflow designed to deliver consistent, auditable, and accurate outputs.

Ingestion & Processing

Intelligent Indexing

Contextual Mapping

C.W. Review & Finalisation

Learning & Optimisation

Ingestion & Processing:

- Documents uploaded and stored securely
- Parser extracts structure, hierarchy, and metadata.
- OCR processes handwritten or scanned content
- Content is standardised while preserving original structure
- Detects and labels key terms and data for classification.

Intelligent Indexing:

- Text is chunked using use case optimised strategies
- OpenAI model generates vector representations
- Embeddings stored securely in an HGS database with metadata on document type and source.
- Index is created with optimised retrieval parameters

Contextual Mapping:

- System maps content to relevant domains.
- Contextual rules guide information classification
- Evidence strength is evaluated using trained algorithms.
- Source citations are maintained for traceability

C. W. Review & Finalisation:

- Generates draft from top-ranked content.
- Checks accuracy and completeness.
- Case Worker reviews, refines, & finalises content.
- Final output generated in MS Word format for download

Learning & Optimisation:

- Successful submissions are tracked for performance metrics
- Rejection patterns inform system improvements
- Fine-tuning dataset grows with approved examples
- Index maintenance refreshes based on usage patterns

Why should Government organisations invest in Case Assist?



Requirements & problem felt needs our solution can help to address.



Health & Social Care (CHC, Adult & Children's Services)

- Significant manual workload spent interpreting and summarising large volumes of information.
- Backlogs caused by manual document review.
- Risk of missing key information within extensive case documentation.
- Need for audit-ready reporting that references specific evidence.



Benefits & Citizen Services

- Need to manually review high volumes of documents submitted.
- Eligibility checks become harder when information is spread across multiple files.
- Inconsistent assess claims drive up complaints or appeals.
- Extended turnaround times, affecting service targets and citizen satisfaction.



Procurement & Commercial Services

- Need to reduce time required to review, summarise, and evaluate large volumes of contracts with high accuracy.
- Find and mitigate risks in existing and future commercial contracts.
- Improve the time it takes to triage commercial claims in existing contracts.



HMRC & Treasury Services

- Officers spend significant time manually reviewing large evidence packs submitted for relief claims, slowing case progression.
- Key financial or compliance details are difficult to surface quickly when information is dispersed across multiple formats.
- Fraud indicators or discrepancies are hard to detect manually.



Immigration & Visa Services

- Officers must review large, varied evidence packs submitted in multiple formats, which can make case preparation slow and complex.
- Need to consistently interpret evidence without spending additional time reconciling policy requirements.

Case Assist can help to address these and other requirements and equip government organisations with a faster, more consistent, and audit-ready way to handle complex case documentation.

*A Local Authority using Case Assist saw their time to complete CHC assessment reports reduce from an average of 4 hours to ~30 minutes.



Higher Productivity & Efficiency

- Faster report generation and case summarisation.*
- Significant reduction in administrative burden for case workers.
- Ability to handle higher caseloads without increasing staff numbers.



Quality Narrative & Evidence

- Improved decision quality through consistent application of domain rules.
- Standardised reporting templates across teams and departments.



Compliant with Industry Policies

- Fully aligned with GDPR, NHS, Local Authority data standards, and NCSC Cyber Assessment Framework.
- Every AI output is grounded in source material, creating legally defensible reports.



Empowered Staff

- Allows staff to focus on professional judgement rather than paperwork.
- Improves morale and reduces burnout by removing administrative load.

Service Model Options

MODEL I - DOCUMENT IMPORT

Document Import

Public/Private knowledge base
(can be limited to specific sections)

Creates searchable knowledge from
manually imported documents.

MODEL II – AGENTIC AI

Agentic AI

Integrated with data systems
(can be limited to specific processes)

Can answer queries about customer
data

MODEL III – PERSONALISED CONCIERGE

Personalised Concierge

Personas modelled to Organisation user
segments

AI responds to customer personas

Technology & Infrastructure



Sl.	Component	Technology	Version	Supported
1	Cloud Provider	AWS (UK), Azure (UK)	NA	NA
2	Datastore	MongoDB Atlas	NA	NA
3	Media store	S3	NA	NA
4	Micro-services	Uvicorn	0.34.2	NA
5	Asynchronous processing	Celery	5.5.2	TBD
6	Cache	Redis	5.3.0	TBD
7	Presentation Layer	React + Tailwind	18.2.0	TBD
8	UI/UX design Framework	Material Design	NA	TBD
9	Large language models (LLM)	OpenAI (default) or llama	GPT 4o versions or 3.1 8B	NA
10	Authentication & Authorization	Autho	NA	NA
11	Container Orchestration	ECS	TBD	Kubernetes
12	LLM hosting	EC2	g4dn family	Any compatible
13	Container Registry	AWS ECR	NA	Any compatible

Case Assist is designed to meet the stringent expectations of UK Government, NHS and Local Authority clients, with security aligned to the National Cyber Security Centre (NCSC) Cyber Assessment Framework (CAF) 14 principles.

Security Standards and Certifications

- Managed within an ISO 27001 certified Information Security Management System (ISMS), supported by Cyber Essentials Plus, ISO 9001 and ISO 22301 (business continuity and disaster recovery).
- All platform services are hosted within UK sovereign data centres (AWS London and Azure UK regions), ensuring no data is transferred outside the UK. This guarantees compliance with UK GDPR and the Data Protection Act 2018.
- Data is encrypted at rest (AES-256) and in transit (TLS 1.2+) with encryption keys managed in line with NCSC Key Management principles.
- Personnel security requires BPSS clearance for all staff, with SC clearance for administrators and privileged roles, where required (not in all cases).

Technical Controls Implemented in Case Assist

Sl.	Aspect	Standard	Version
1	Transport layer security	TLS	TLS 1.2 or higher, aligned with NCSC guidance
2	Web application security	OWASP	OWASP Top 10 applied in dev/test and monitored in production
3	Web application firewall	Akamai	Always-on WAF with ruleset tuned to NHS/LA threat profiles
4	Intrusion detection / prevention	GuardDuty / Sentinel SIEM	Integrated with CloudWatch, 24/7 SOC monitoring and alerting
6	Authentication	Auth0 / Council IdP	Supports SSO (e.g. Entra ID) and MFA; RBAC enforced
7	Authorisation	Application (Server)	Role based access control
8	Password storage	Auth0 encryption	One way salted hashing and strong key management
10	Stored document encryption	AES-256	All uploaded files encrypted at rest
11	Encryption at rest – Data	AES-256	Full database encryption, key management aligned with NCSC principles
12	Encryption – Configuration	Yes	Secrets and configuration values encrypted and access controlled
13	Container vulnerability scanning	Yes	Automated image scans at build and deploy; remediation tracked
14	Observability	CloudWatch + Sentinel	Metrics, logs and alerts integrated with SIEM
15	Audit / non-repudiation	Scrut + immutable logs	Logs retained minimum 6 months; full non-repudiation
16	Penetration testing	CHECK/CREST accredited testers	Annual ITHCs; additional tests triggered by major or cumulative minor releases
17	Vulnerability scanning	Monthly automated scans	Covers infrastructure, containers and web applications (OWASP Top 10, SANS 25)
18	Data residency	UK sovereign hosting	All services hosted in AWS London and Azure UK; no data leaves the UK

Thank you.

Name: Rob Irons

Email: rob.irons@teamhgs.com

Tel : +44 7827 940800

