

Technical Architecture Brief

Eligibility Copilot — AI-Assisted Decision Support for Public Benefits Intake

SOLUTION OVERVIEW

Eligibility Copilot is a human-supervised AI decision-support system designed for public benefits intake operations. The system assists caseworkers in processing applicant documentation, applying multi-program eligibility rules, and producing auditable determination recommendations. Human review and override capability are structural requirements of the system, not optional features.

PROCESSING PIPELINE

1. Document Ingestion

Applicant or caseworker uploads supporting documents (IDs, income statements, residency verification). Documents are passed to the OCR pipeline for processing.

2. OCR & Extraction (AWS Textract)

AWS Textract performs field-level extraction on uploaded documents. Output is structured JSON containing identified fields, values, and confidence scores. Supports multi-format document inputs.

3. LLM-Assisted Analysis

Extracted structured data is passed to a RAG-based LLM pipeline. The language model applies the relevant program eligibility rules (SNAP, Medicaid, CHIP) to the applicant data, generating a recommendation with criterion-level citations and an overall confidence score.

4. Human Review Interface

The caseworker receives the AI recommendation in a structured review interface showing the recommendation, supporting citations, confidence level, and source document excerpts. The caseworker confirms or overrides.

5. Audit Event Logging

Every event — document upload, AI recommendation, caseworker action, override rationale — is written to the audit log with worker ID, timestamp, and record identifier. The log is tamper-evident and exportable.

6. Governance Artifact Generation

On a configurable schedule, the system generates AIA documentation, OMB AI inventory entries, and parity gap reports across demographic dimensions including language (English/Spanish).

INFRASTRUCTURE SUMMARY

| | |
|------------------|--|
| Frontend | React application deployed on Vercel (CDN-distributed, HTTPS enforced) |
| Backend API | Python/FastAPI containerized application deployed on Render |
| OCR Engine | AWS Textract — FedRAMP-authorized service |
| LLM / Analysis | LLM-based RAG pipeline with program rules knowledge base |
| Database | Structured relational store for audit logs, case records, and governance artifacts |
| Authentication | Role-based access control; caseworker and supervisor roles |
| Language Support | English and Spanish — bilingual document processing and UI |

AUDITABILITY & OVERSIGHT FEATURES

- Complete longitudinal audit trail: AI input, AI output, caseworker action, timestamp, and worker ID on every record
- AI recommendation is stored independently from final determination — enabling accuracy analysis and retrospective review
- Override events captured with mandatory rationale field — supporting quality assurance and pattern analysis
- Exportable audit log in structured format for IG review, legal discovery, or internal oversight
- Parity gap dashboard tracking outcome consistency across language groups and demographic dimensions
- Auto-generated AIA and OMB AI inventory documentation aligned with current federal guidance
- All AI outputs include confidence score and criterion-level citation — no opaque recommendations

DEPLOYMENT NOTES

Eligibility Copilot is currently deployed in a production configuration accessible to demonstration and pilot participants. The system is designed for agency evaluation with the following considerations:

- Current deployment uses cloud-hosted infrastructure; on-premises or GovCloud deployment available for agency requirements
- No PII is retained beyond the active session scope in the current configuration — agency data governance requirements addressed in pilot design
- FedRAMP path: core infrastructure (AWS Textract) is already FedRAMP-authorized; full system FedRAMP assessment is on the production roadmap
- Integration with existing case management systems via API is supported; specific integration scope defined during pilot scoping
- Pilot deployment is designed to run in parallel with existing workflows — no displacement of current processes required

Technical inquiries: kent@thehentexgroup.com | thehentexgroup.com