ENGINEERING CASE STUDY

Modernizing a 15+ Year-Old Vital Records Workflow with AI

High-accuracy document extraction integrated into an established North Carolina government software system.

98%	95%	45%
Birth-field accuracy	Death-field accuracy	Estimated labor reduction

THE MODERNIZATION CHALLENGE

A legacy application workflow, originally designed and built more than 15 years before the AI integration, relied on two independent manual transcriptions of each vital record. The opportunity was to add intelligent automation without replacing the application, removing established quality controls, or disrupting staff workflows.

ENGINEERING CONTRIBUTION

Train a domain-specific document extractor, map complex outputs into legacy record structures, and enhance the existing review screen with prefilled values and source-field highlighting.

Deployment: pilots at two organizations; production at one. | Approximate project period: 2025.

01 The problem and the constraints

Vital records demand highly accurate structured data. The established workflow used two independent transcriptions and adjudication of discrepancies. Scanned certificates presented additional difficulties: noisy or degraded text, variable scale and positioning, imperfect field labels, and death-certificate fields containing multiple data elements.

02 How the extraction approach evolved

Early coordinate-based prototype

Recognized labels and cropped expected text regions for transcription. Fixed-location assumptions were vulnerable to differences in scale, alignment, and image quality.

Trained document understanding

Google Document AI Custom Extractor used manually labeled examples to learn certificate fields. Label recognition, fuzzy matching, and spatial context improved field identification.

03 Integrated architecture

01
Scanned certificate
Birth or death image

→

02
Custom Extractor
Learned field recognition

→

03
Transformation logic
Parse / map legacy fields

04
Prepopulated review
Highlight image region

→

05
Independent entry
Compare and adjudicate

WHY THE INTEGRATION MATTERED

The legacy application retained ownership of record storage, operator workflows, and discrepancy resolution. AI supplied structured values to those workflows rather than requiring a full replacement platform.

04 Measured extraction performance

Each certificate type used approximately 1,000 training images and 1,000 separate test images. Accuracy was measured by field-level exact match against ground-truth values; it was not a document-level pass rate.



Residual errors were commonly associated with text recognition. More complex death-certificate fields also required downstream parsing into individual structured data elements.

05 Estimated change in human effort

WORKFLOW	BIRTH	DEATH
Two manual transcriptions	10 min	20 min
AI review + one manual transcription	5.5 min	11 min
Estimated labor-effort reduction	**45%**	**45%**

Times are reconstructed estimates, not instrumented production benchmarks. They exclude shared adjudication effort and automated processing time. Designed capacity was at least 100 births and 100 deaths per hour; sustained production throughput was not established.

06 Why human review stayed in the loop

The accuracy achieved made it reasonable to explore using AI extraction as one independent transcription without a first-pass human reviewer. The deployed approach instead retained a short review of AI-populated fields, followed by separate human transcription and discrepancy adjudication.

This protected established accuracy practices and helped ease employee concerns that automation would replace their roles. Focusing a field highlighted the matching region on the source image, turning a blank-form entry task into a rapid, visually guided verification task.

REUSABLE ENGINEERING LESSON

Modernize the capability, not necessarily the entire system. A 15+ year-old government workflow gained learned document extraction, structured data transformation, and a faster review experience while retaining the application's proven operational controls.

Evidence note: This retrospective case study reports approximate historical accuracy, reconstructed effort estimates, and an engineering throughput target. Final recorded-data accuracy after human verification was not separately quantified.