

TECH & AI CASE STUDY - MORTGAGE LENDING - 2026

Document AI That Cut Mortgage Underwriting Time by 70%

Based on a delivery engagement with Brightlane Mortgage

-70%

UNDERWRITING TURNAROUND

6.5 to 1.2 hrs

DATA ENTRY PER FILE

98.6%

EXTRACTION ACCURACY

+38%

LOANS PER UNDERWRITER

A fillable practice edition: the full story, 8 actions,
6 reflections and a 90-day plan - work it with your team.

Free - no email required
emersontechnology.com

MAKE THIS PRACTICE BOOK YOURS

Work it, do not just read it

This practice book turns the Brightlane Mortgage case study into a working session for your team. Read the full story first - what happened, what was built and what it changed. Then tick the 8 actions you can commit to, write short answers to the 6 reflections, and finish with a 90-day plan. Every field in this PDF is fillable: open it in Adobe Acrobat Reader, Chrome, Edge or macOS Preview, type directly into the boxes and save.

YOUR NAME	ROLE
COMPANY	TEAM / DEPARTMENT
DATE	EMAIL (OPTIONAL)

MY GOAL FOR THIS PRACTICE SESSION

- 01 The story** - What happened at Brightlane Mortgage - before, during and after the project.
- 02 The results** - The numbers that changed, plus industry benchmarks from public research.
- 03 The 8-action checklist** - Tick boxes as your organization completes each step. Owner and date fields included.
- 04 The 6 reflections** - Fillable questions that pressure-test your own roadmap.
- 05 My 90-day plan** - Three fillable phase cards that turn this case study into your roadmap.

RUN IT AS A 90-MINUTE TEAM SESSION

25 min READ Read the story together. Ask everyone to note the one number that surprised them most.	40 min PRACTICE Work the 8-action checklist and the 6 reflections. Assign an owner and a date to every action you commit to.	25 min COMMIT Fill in the 90-day plan. One goal per phase, one number that tells you it worked. Book a review date before you leave the room.
--	--	---

THE STORY

What happened, and why it worked

01 The Company

Brightlane Mortgage is an independent lender funding about \$1.4B in residential loans a year through purchase and refinance channels, working with a correspondent network and a retail loan-officer force that competes for borrowers in some of the fastest-moving markets in the country. Its underwriting operation is the heart of the business: every file must move from a sprawling borrower document package to a clean, decision-ready loan record, and the speed of that transformation is effectively the company's competitive position. In mortgage lending, the lender who can answer a complete file fastest wins the borrower, because borrowers shop on rate for a week and on certainty for the remaining forty days.

The company employed a seasoned corps of processors and underwriters whose expertise was real but whose time was consumed by work beneath their expertise: opening documents, finding figures, and retyping them into the loan origination system. Leadership had watched volumes swing hard with rate cycles — refinance waves one quarter, purchase dominance the next — and had lived the staffing whiplash of hiring contractors for surges they could not train fast enough, then standing the capacity down when the wave passed.

02 Where Things Stood

Underwriting ran on manual document review. Processors opened tax forms, pay stubs, bank statements and employment letters — a 94-page average package per file — then retyped every relevant figure into the loan origination system by hand. The data entry alone consumed 6.5 hours per file, spread across days of interruption, and every keystroke was an opportunity for a transposition error that would surface later as a rework loop, a stalled condition or a compliance finding. Average turnaround sat at 11 days from complete file to underwritten decision, and the clock mattered: roughly one in five borrowers drifted to faster competitors mid-process, costing an estimated \$1.8M a year in lost commissions.

The error tax was its own line item. Keystroke-driven rework consumed another \$900K a year in overtime and quality-control spend, and the rework loops hit morale as hard as margin — a processor who spent Thursday retyping a bank statement spent Friday explaining why Thursday's figure disagreed with Wednesday's. Industry context made the internal pain feel less like an outlier and more like the norm: ICE Mortgage Technology's Origination Insight Reports have shown the average US mortgage taking more than 40 days to close, underscoring how much manual work remains across the industry, and McKinsey Global Institute research suggests automation could cut the cost of credit checks by up to 30% for banks, with document-heavy tasks among the most automatable work.

Seasonality turned a hard job into an impossible one. When refinance volume spiked, contractor underwriters arrived with two weeks of training and inherited the same 94-page packages, and error rates followed. Quality-control sampling caught the worst of it, but the QC team was sized for steady state, so the very periods that generated revenue were the periods that generated exceptions.

03 The Challenge

The challenge was to remove the retyping without removing the judgment — automating the extraction and validation of income, asset and employment data while keeping credit decisions, exceptions and conditions with licensed humans. The technical difficulty was not OCR in general; it was the long tail of real documents: regional pay stub formats, self-employment tax returns with schedules, bank statements with asset transfers that need explanation, and employer letters that say the right thing in the wrong order. Extraction quality on the easy 80% proves little; the economics live in how the system handles the difficult remainder.

The organizational difficulty was trust and compliance. Underwriters had seen automation initiatives before that moved work around without reducing it, and compliance leadership needed a defensible record of where every number came from — model version, source document, reviewer action. JPMorgan Chase has reported that its COIN platform reviews commercial credit documents in seconds, work that previously consumed an estimated 360,000 lawyer-hours annually; that example impressed Brightlane's board, but the board's conclusion was the right one: the value was proven only when extraction accuracy could be measured, audited and held steady at volume.

04 Discovery and Diagnosis

Emerson Technology began with a four-week discovery across Brightlane's document intake, origination system and underwriting workflow. The first exercise was an honest catalog: every document type the company receives, including regional and lender variants, mapped with who touches it and where it stalls. The second was timestamp instrumentation — logging handoffs across the existing process before automating anything — which produced a finding leadership did not expect: the largest delays were not in underwriter decision time but in document intake, rework loops and waiting on missing items, exactly the stages nobody had been measuring. Baseline metrics made the ROI case obvious and revealed which stages actually created delay rather than where teams assumed it sat.

The diagnosis also quantified the extraction opportunity. Sampling recent files showed that the core income, asset and employment figures lived in a manageable set of document types, and that most fields were extracted correctly even by naive OCR — the failures concentrated in classification errors (a schedule C read as a W-2) and in normalization (annualizing a bonus correctly). That pointed the architecture at the right problem: routing each page to the right extraction schema, validating figures against each other, and routing low-confidence fields to humans. The team also established a gold-standard set of annotated documents during discovery, so every future accuracy claim would be measured against labeled truth rather than asserted.

05 The Solution

Emerson Technology designed and deployed an intelligent document processing pipeline on AWS that turns borrower packages into validated, decision-ready data with a full audit trail. The architecture has three layers.

Extraction and classification

Borrower files arrive through a secure Next.js portal, where Amazon Textract handles OCR and a fine-tuned LayoutLM classification model in PyTorch routes each page to the right extraction schema — income, assets, employment or insurance. The classification layer is where generic document AI usually fails, so it was trained on Brightlane's own document catalog, including the regional and self-employment variants the discovery had catalogued. Python microservices normalize the extracted figures — annualizing bonuses, mapping employer names, structuring asset histories — and cross-check pay stubs against W-2s and bank statements, flagging any field below a confidence threshold. The pipeline starts by lifting data out of documents rather than making credit calls, because extraction carries lower risk, builds trust and delivers most of the time savings in document-heavy

operations.

Human review of exceptions

A review console shows underwriters every extracted value beside its source snippet, so a reviewer verifies rather than retypes — and low-confidence fields route to reviewers automatically while high-confidence fields flow straight through. This confidence-threshold routing keeps accuracy high while cutting review volume by an order of magnitude, and it concentrates human attention on exactly the ambiguous cases where judgment earns its salary. Underwriters approved the model of working quickly once they saw it: the console does not ask them to trust the machine, it asks them to check it, which is what they were already doing by hand at ten times the cost.

Orchestration and audit trail

Temporal orchestrates the multi-day underwriting workflow with automatic rechecks for missing documents — a stalled file waiting on an updated bank statement gets chased by the system, not by a processor remembering at 5pm. Every extraction lands in Postgres with a full audit trail recording the source image, model version and reviewer action, feeding compliance reports and the LOS integration. Accuracy is measured against the gold-standard annotated set on every release, because extraction drift is silent otherwise and compliance teams will eventually ask for evidence. The audit trail was designed in from day one rather than retrofitted — in a regulated industry, retrofitting auditability costs far more than building it in.

06 The First 90 Days

Weeks one and two completed the document catalog and the timestamp instrumentation, then stood up the portal and Textract pipeline in a staging environment. Weeks three through six piloted extraction on two loan programs with the highest volume, running the pipeline in parallel with manual processing: every extracted field was compared against the processor's retyped values, and disagreements fed the classification model's tuning loop. The gold-standard set grew weekly, and extraction accuracy was reported against it on every release — a discipline that made the go/no-go conversations factual rather than rhetorical.

Weeks seven through ten deployed the review console to processors and underwriters, with a 45-minute onboarding each and a one-page playbook for the exception workflow; reviewers graded the console's snippets as the feature they had not known they needed. Temporal orchestration went live in the same window, and the automatic rechecks immediately surfaced how many files had been quietly aging in "waiting on documents" purgatory. Weeks eleven and twelve extended coverage across channels, connected the LOS integration for production flow, and published the first weekly cycle-time dashboard by channel and loan type — visibility that keeps operations, sales and technology aligned on where the next bottleneck has moved. By the end of the quarter the pipeline had processed its heaviest weekly volume to date without a single contractor hire.

07 The Results

Within six months, the numbers moved decisively. Average underwriting turnaround fell from 11 days to 3.3 — a 70% reduction — and manual data entry dropped from 6.5 hours per file to 1.2, converting processor time from keystrokes into verification and exception handling. Document extraction accuracy held at 98.6% across incoming packages, measured continuously against the gold-standard set, and completed loans per underwriter rose 38% without adding headcount.

METRIC	BEFORE	AFTER
Underwriting turnaround	11 days	3.3 days (-70%)
Data entry per file	6.5 hours	1.2 hours

METRIC	BEFORE	AFTER
Extraction accuracy	-	98.6%
Loans per underwriter	-	+38%

The qualitative changes compounded the quantitative ones. Fall-out from slow processing declined noticeably as borrowers stopped waiting three weeks for an answer, protecting the commission revenue that mid-process drift had been eroding. The team absorbed its heaviest refinance season without hiring contractors, which changed the seasonality conversation from staffing panic to capacity planning. Borrower complaints about paperwork status dropped sharply as portal updates became automatic, and the QC team now audits exceptions by exception rather than by sampling blind. When a regulator asks where a figure came from, the answer is a source snippet, a model version and a reviewer name — in seconds, not days.

08 What Your Business Should Do

The pattern behind Brightlane's result is repeatable in any document-heavy operation, and it starts before any model is chosen. Map every document type you receive — catalog each form, statement and letter in your intake flow, including regional and lender variants, because extraction quality depends on knowing exactly what arrives and who touches it. Instrument current turnaround times first, logging timestamps at every handoff, so baseline metrics make the ROI case obvious and reveal which stages actually create delay rather than where teams assume it sits. Automate extraction before decisions: start by lifting data out of documents, not making credit calls, since data extraction carries lower risk, builds trust and delivers most of the time savings.

Keep humans on exceptions only, routing low-confidence fields to reviewers with source snippets attached — confidence-threshold routing keeps accuracy high while cutting review volume by an order of magnitude. Build an audit trail from day one, storing every extracted field with its source image, model version and reviewer action, because regulated industries will demand it and retrofitting auditability later costs far more than designing it in. Orchestrate long workflows with a durable engine such as Temporal rather than scripts, since retry logic, timers and state tracking prevent files from stalling when an external party is slow. Measure accuracy against a labeled gold-standard set on every release, because accuracy drift is silent otherwise and compliance teams will ask for evidence — and publish cycle time weekly by channel and loan type in one shared dashboard, because visibility keeps operations, sales and technology aligned on where the next bottleneck has moved.

Ask the questions Brightlane's leadership now reviews quarterly: how many hours your team spends retyping data that already exists inside documents you receive every day; which stage of your process customers would call slowest, and whether your internal metrics agree; what a 70% faster turnaround would be worth to you in closed deals, referrals or retention; whether you could trace every figure back to its source document if a regulator audited one file tomorrow; where keystroke errors enter your workflow today and what they truly cost annually; and which process you would automate first if accuracy could be proven on a small pilot. Work through the eight actions in this practice book with your team, and revisit those reflection questions each quarter to keep the pipeline honest.

REAL-WORLD PROOF & BENCHMARKS

McKinsey Global Institute research suggests automation could cut the cost of credit checks by up to 30% for banks, with document-heavy tasks among the most automatable work. - McKinsey Global Institute

JPMorgan Chase has reported that its COIN platform reviews commercial credit documents in seconds, work that previously consumed an estimated 360,000 lawyer-hours annually. - Reuters, 2017

ICE Mortgage Technology's Origination Insight Reports have shown the average US mortgage taking more than 40 days to close, underscoring how much manual work remains. - ICE Mortgage Technology

THE PRACTICE PAGES

The 8-action checklist

Work through the list with your team. Tick each box as your organization completes the step, and fill in an owner and a target date so every action has a name and a deadline.

☐**1. Map every document type you receive**

Catalog each form, statement and letter in your intake flow, including regional and lender variants. Extraction quality depends on knowing exactly what arrives and who touches it before any automation begins.

OWNER

TARGET DATE

STATUS

--	--	--

☐**2. Instrument current turnaround times first**

Log timestamps at every handoff before automating anything. Baseline metrics make the ROI case obvious and reveal which stages actually create delay rather than where teams assume it sits.

OWNER

TARGET DATE

STATUS

--	--	--

☐**3. Automate extraction before decisions**

Start by lifting data out of documents, not making credit calls. Data extraction carries lower risk, builds trust, and delivers most of the time savings in document-heavy operations.

OWNER

TARGET DATE

STATUS

--	--	--

☐**4. Keep humans on exceptions only**

Route low-confidence fields to reviewers with source snippets attached. Confidence-threshold routing keeps accuracy high while cutting review volume by an order of magnitude in most deployments.

OWNER

TARGET DATE

STATUS

--	--	--

☐**5. Build an audit trail from day one**

Store every extracted field with its source image, model version and reviewer action. Regulated industries will demand it, and retrofitting auditability later costs far more than designing it in.

OWNER

TARGET DATE

STATUS

--	--	--



6. Orchestrate long workflows, don't script them

Use a durable workflow engine such as Temporal for multi-day processes. Retry logic, timers and state tracking prevent files from stalling when an external party is slow.

OWNER

TARGET DATE

STATUS

--	--	--



7. Measure accuracy against a labeled set

Maintain a gold-standard set of annotated documents and test extraction against it on every release. Accuracy drift is silent otherwise, and compliance teams will ask for evidence.

OWNER

TARGET DATE

STATUS

--	--	--



8. Publish cycle time weekly

Report underwriting turnaround by channel and loan type every week in one shared dashboard. Visibility keeps operations, sales and technology aligned on where the next bottleneck has moved.

OWNER

TARGET DATE

STATUS

--	--	--

AFTER THE CHECKLIST: OUR TOP 3 PRIORITIES

From the eight actions above, write down the three you will start first - and the single number you expect each one to move. Three priorities with owners beat eight without any.

PRIORITY 1 NUMBER IT SHOULD MOVE	PRIORITY 2 NUMBER IT SHOULD MOVE	PRIORITY 3 NUMBER IT SHOULD MOVE
--	--	--

THE PRACTICE PAGES

Reflections for your team

Use these questions in your next planning session and type short answers into the boxes. The gaps you find are usually the fastest wins.

Q1. How many hours does your team spend retyping data that already exists inside documents you receive every day?

Q2. Which stage of your process would customers describe as slowest, and do your internal metrics actually agree?

Q3. What would a 70% faster turnaround be worth to you in closed deals, referrals or retention?

Q4. If a regulator audited one file tomorrow, could you trace every figure back to its source document?

Q5. Where do keystroke errors enter your workflow today, and how much do they truly cost you annually?

Q6. Which of your processes would you automate first if accuracy could be proven on a small pilot?

THE PRACTICE PAGES

My 90-day plan

Finish the session by committing to a plan. Three phases, one goal each, one number that tells you it worked. Bring this page to your next leadership review - or to a free consultation with Emerson Technology, and we will pressure-test it with you.

Days 1-30 Focus: audit and alignment

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 31-60 Focus: build and prove

PRIMARY GOAL**KEY ACTIONS****HOW WE MEASURE SUCCESS****OWNER**

Days 61-90 Focus: scale and measure

PRIMARY GOAL

KEY ACTIONS

HOW WE MEASURE SUCCESS

OWNER

Tip: save a copy per team so each group keeps its own answers. Revisit the reflections every quarter - the questions stay the same even as the numbers move.

BOOK THE REVIEW BEFORE YOU CLOSE THE BOOK

Plans without a review date drift. Before you close this practice book, agree on when you will come back to these pages, who will run that session, and when you will check the numbers again. If you want an outside pair of eyes, bring this page to a free consultation with Emerson Technology and we will pressure-test the plan with you.

REVIEW DATE	REVIEWED BY	NEXT CHECK-IN
------------------------------------	------------------------------------	--------------------------------------

Ready for a transformation like Brightlane Mortgage's?

Emerson Technology designs and builds the data, AI, cloud and automation systems behind results like these. On a free consultation call we will map this exact playbook to your business - your data, your constraints, your numbers.

Book a free consultation

support@emersontechnology.com

or start online: emersontechnology.com/contact