

TECH & AI CASE STUDY - DIGITAL PUBLISHING - 2025

A GenAI Editorial Workflow That Tripled Content Output

Based on a delivery engagement with Quillhouse Media

180 to 540
MONTHLY OUTPUT

87%
EDITOR FIRST-PASS ACCEPTANCE

-58%
DRAFT TIME

+34%
ORGANIC TRAFFIC

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 Quillhouse Media 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 Quillhouse Media - 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

Quillhouse Media is a digital publishing company producing 180 articles a month across 14 verticals — consumer finance, home and garden, technology, health and others — for an audience it reaches almost entirely through organic search and its newsletters. The business model is volume with standards: coverage depth drives search authority, search authority drives traffic, and traffic drives the advertising and affiliate revenue that funds the newsroom. Its 26 staff writers are supported by a bench of freelancers, and its editorial standards are genuinely high — every piece passes an editor, and the publication's credibility with readers and search engines alike depends on that gate holding at any volume.

Growth targets set for 2025 demanded triple the output. The commercial logic was sound — competitors were publishing faster in Quillhouse's mid-tail verticals and renting the search authority Quillhouse had spent years building — but the arithmetic was not: hiring the equivalent writing capacity had been priced at \$1.9M a year the company did not have. The newsroom needed a way to publish three times more without becoming a content farm.

02 Where Things Stood

The pipeline bottlenecked in two predictable places: research and first-pass structure. Writers spent hours assembling sources and outlining before a draft even began, and editors then spent roughly 60% of their time on formatting and SEO hygiene — metadata, internal links, heading structures, image alt text — rather than on the judgment work of angle, argument and accuracy. The most expensive people in the pipeline were doing the least differentiated work, and every added article multiplied the tax.

Quality problems were already visible at current volume. Near-duplicate content across verticals — the same comparison article published twice with different headlines — had dented search rankings on topics the company should have owned. Brand voice drifted noticeably between staff writers and freelancers, so the publication read like fourteen different publications depending on the day of the week. And the previous year's experiment haunted every planning conversation: leadership had trialed raw AI drafts without quality gates, published a batch to meet a traffic target, and retreated publicly after corrections and an uncomfortable internal postmortem. That failure was, in hindsight, the most useful asset in the company — it had proven precisely which failure modes any future pipeline would need to prevent.

McKinsey has estimated generative AI could add \$2.6 trillion to \$4.4 trillion in value annually across analyzed use cases, much of it in content and marketing workflows, and Gartner has predicted that around 30% of outbound marketing messages from large organizations will be synthetically generated by 2025, up from just 2% in 2022. Quillhouse's leadership believed the direction. Their question was narrower and harder: could the economics work at their scale without surrendering the editor's veto?

03 The Challenge

The challenge was to triple throughput while making quality gates tighter, not looser — the opposite of how scaling usually feels inside a newsroom. The pipeline had to automate the undifferentiated 80% of the work (research assembly, first-pass structure, formatting, SEO hygiene) while keeping the differentiated 20% (angle, argument, voice, publish decisions) in human hands. It also had to solve the two failure modes the prior AI experiment had exposed: factual drift and voice drift. A model that writes plausible sentences is table stakes; a system that refuses to invent claims and sounds like Quillhouse in every vertical is the actual product.

There were commercial constraints as well. Search engines were growing less tolerant of thin, duplicated or unreviewed content, and Reuters Institute surveys reportedly find back-office and workflow automation rank among publishers' top AI priorities, ahead of fully automated newsroom content — a signal the industry had learned the same lesson Quillhouse had. The pipeline therefore had to make every draft more reviewable, not less: editors needed to see what changed and why, fast, at three times the volume.

04 Discovery and Diagnosis

Emerson Technology began with a three-week discovery across the editorial workflow, content archive and analytics. The audit mapped where hours actually went — research assembly, first-pass drafting, editing, formatting, SEO hygiene, publication — and confirmed the two bottlenecks with data from writers' own time tracking. The archive audit then quantified the duplication problem: clusters of near-identical articles across verticals, several of which had been cannibalizing each other in search results. Those clusters became the specification for a research service that would see the whole catalog before proposing any new piece. Interviews with editors after the failed AI year produced the most valuable requirements. Editors did not object to machine drafting; they objected to reviewing a finished blob. What they asked for was a draft that arrives pre-structured, with its claims flagged and sourced, its deviations from the style guide visible, and its AI contributions marked — a diff, not a manuscript. The diagnosis reframed the build: the product was not an article generator but an editorial pipeline with automated quality gates and a human sign-off that stays mandatory at any volume. That framing — editors as the mandatory last step, not a rubber stamp — became the design principle everything else served.

05 The Solution

Emerson Technology designed and built an editorial pipeline inside a Next.js workspace, with automated gates before every human review and a human gate before every publication. The architecture has three layers.

Research and briefing

A research service clusters search demand and competitor gaps across the 14 verticals, then assembles a brief with angle options for editors to approve before anyone drafts. The service sees the whole catalog, which is how near-duplicates get caught at proposal stage rather than after publication — the system will not propose a brief that overlaps an existing piece without an explicit differentiation angle. Briefs include target queries, required sources and the questions readers actually ask, so writers and the drafting model start from the same evidence. Editors report the briefing step alone changed their week: they approve angles in minutes and spend their judgment on whether the angle is worth publishing, not on assembling background.

Grounded drafting

Draft generation uses a GPT-4-class model with retrieval over Quillhouse's style guide, fact base and archive through pgvector RAG, so output sounds like the publication rather than a chatbot. Grounding is the difference the failed experiment lacked: the model draws on approved sources and the archive's established positions, which makes it both more accurate and more recognizably Quillhouse, and it invents far less when it is fetching rather than remembering. Prompts are versioned in git with test cases on representative briefs for every vertical, so a prompt change that quietly degrades a vertical's voice surfaces in the test suite instead of three weeks later in a reader email.

Quality gates and human sign-off

Every draft passes automated quality gates before an editor opens it: plagiarism screening against the web and the archive, a brand-voice classifier scored per vertical, banned-claim refusals drawn from a codified list of

superlatives and regulated phrases the system must never generate, and fact-check flags that require linked sources for every checkable claim. Editors then review a structured diff showing exactly what changed and which passages were machine-contributed, and nothing publishes without human sign-off — a rule enforced by the pipeline itself, not by policy. Temporal orchestrates the queue so drafts move through gates with retries and state tracking, and performance data flows back monthly into prompt and brief tuning. The gates matter most precisely when volume scales past reviewer capacity: automated screens protect quality exactly when human attention is most stretched.

06 The First 90 Days

Weeks one and two built the banned-claims list with the editorial leads — the fastest, cheapest risk control in the whole system — and versioned every existing prompt with per-vertical test cases. Weeks three through six piloted the full pipeline on two verticals with contrasting voices, running generated drafts through the gates and into editors' queues alongside normal assignments. The first-pass acceptance metric became the pilot's scoreboard from day one: volume without acceptance just moves work downstream, so the team tracked how many drafts editors accepted without rework, and iterated briefs and retrieval until the number climbed.

Weeks seven through ten extended the pipeline to eight verticals, added the structured-diff review experience based on pilot feedback, and set the publication rule in stone: no human sign-off, no publish, enforced in code. Writers were repositioned rather than replaced — briefs give them the boring 80% pre-assembled, and they own the judgment calls, reporting that their work shifted from assembling sources to sharpening arguments. Weeks eleven and twelve completed the rollout across all 14 verticals, turned on the monthly performance loop correlating traffic and engagement with draft patterns, and retired the freelance voice-drift problem with a style guide the retrieval layer actually enforces. By the end of the quarter the pipeline was publishing at triple the old cadence with the same headcount, and the editor's veto had survived intact — the quality gates caught more issues per piece as volume rose, not fewer.

07 The Results

Within two quarters, output tripled from 180 to 540 pieces a month with unchanged editorial headcount — the growth target met without the \$1.9M hiring bill. First-pass editor acceptance reached 87%, meaning the overwhelming majority of drafts arrive publishable after one review cycle, and average draft time fell 58% from brief to editor-ready. Organic traffic grew 34% across the two quarters as coverage depth expanded into mid-tail topics the newsroom had never had the capacity to serve, and the near-duplicate clusters stopped forming because the research service sees the whole catalog before proposing anything.

METRIC	BEFORE	AFTER
Monthly output	180 pieces	540 pieces
Editor first-pass acceptance	-	87%
Draft time	-	-58%
Organic traffic	-	+34%

The qualitative changes compounded the quantitative ones. Not a single article publishes without editor sign-off, and the gates hold under triple volume because they are code rather than vigilance. Writers report the pipeline now handles the boring 80% and they handle the judgment calls, which has shown up in retention as well as output. The brand voice reads like one publication across all 14 verticals, and the monthly performance loop has begun steering briefs toward topics the data says readers want — the editorial workflow has become an asset that compounds, rather than a cost center that scales linearly with headcount.

08 What Your Business Should Do

The pattern behind Quillhouse's result is repeatable, and the order of operations matters. Gate every draft automatically: run plagiarism, brand-voice and factual-consistency screens before a human ever opens the first draft, because automated gates protect quality precisely when volume scales past reviewer capacity. Ground generation in your own archive with retrieval over style guides, approved sources and past coverage — models grounded in your corpus sound like your brand and invent far less. Keep editors as mandatory sign-off, with structured diffs that highlight AI contributions clearly, because accountability staying human protects both editorial standards and your search rankings. Automate research, not judgment: use models to cluster topics, summarize sources and draft structure, while angle, argument and voice remain editorial decisions — the ones that actually differentiate a publication.

Build a banned-claims list codifying the claims, superlatives and regulated phrases the system must never generate, because explicit refusals are cheaper, faster and safer than post-publish corrections and legal exposure. Measure acceptance, not output: track first-pass editor acceptance alongside raw volume every week, since output without acceptance just moves work downstream and burns out exactly the people you needed most. Version prompts like code, with test cases on representative briefs for every vertical, so prompt regressions surface in tests rather than weeks later as off-brand, inaccurate content. And feed performance back to drafting by correlating traffic, engagement and conversion with draft patterns monthly — that loop turns your editorial workflow into a compounding asset competitors cannot quickly copy.

Ask the questions Quillhouse's editors now review quarterly: where your content process actually bottlenecks — ideas, drafting, editing or publishing; what tripled output would be worth if quality held, and what breaks in your pipeline first; which quality checks live in editors' heads today that could be automated safely; whether readers could detect AI-assisted drafts now, and whether it would change their trust; how defensible your brand voice is when ten new writers join in one quarter; and what facts, claims or topics your publication should refuse to generate under any circumstances. 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 estimated generative AI could add \$2.6 trillion to \$4.4 trillion in value annually across analyzed use cases, much of it in content and marketing workflows. - McKinsey, 2023

Gartner has predicted that around 30% of outbound marketing messages from large organizations will be synthetically generated by 2025, up from just 2% in 2022. - Gartner, 2022

Reuters Institute surveys reportedly find back-office and workflow automation rank among publishers' top AI priorities, ahead of fully automated newsroom content. - Reuters Institute

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. Gate every draft automatically**

Run plagiarism, brand-voice and factual-consistency screens before a human ever opens the first draft. Automated gates protect quality precisely when volume scales past reviewer capacity.

OWNER

TARGET DATE

STATUS

--	--	--

☐**2. Ground generation in your archive**

Retrieve from style guides, approved sources and past coverage with RAG. Models grounded in your own corpus sound like your brand and invent far less.

OWNER

TARGET DATE

STATUS

--	--	--

☐**3. Keep editors as mandatory sign-off**

Require human approval before anything publishes, with structured diffs that highlight AI contributions clearly. Accountability stays human, which protects both editorial standards and your search rankings.

OWNER

TARGET DATE

STATUS

--	--	--

☐**4. Automate research, not judgment**

Use models to cluster topics, summarize sources and draft article structure. Angle, argument and voice remain editorial decisions — the ones that actually differentiate a publication from its competitors.

OWNER

TARGET DATE

STATUS

--	--	--

☐**5. Build a banned-claims list**

Codify the claims, superlatives and regulated phrases the system must never generate. Explicit refusals are cheaper, faster and safer than post-publish corrections and legal exposure.

OWNER

TARGET DATE

STATUS

--	--	--

☐ **6. Measure acceptance, not output**
Track first-pass editor acceptance rate alongside raw volume every week. Output without acceptance just moves work downstream and burns out exactly the people you needed most.

OWNER

TARGET DATE

STATUS

☐ **7. Version prompts like code**
Store prompts in git with test cases on representative briefs for every vertical. Prompt regressions are silent otherwise and surface weeks later as off-brand, inaccurate content.

OWNER

TARGET DATE

STATUS

☐ **8. Feed performance back to drafting**
Correlate traffic, engagement and conversion with draft patterns on a monthly cadence. That loop turns your editorial workflow into a compounding asset competitors cannot quickly copy.

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. Where does your content process actually bottleneck — ideas, drafting, editing or publishing?

Q2. What would tripled output be worth if quality held, and what breaks in your pipeline first?

Q3. Which quality checks live in editors' heads today that could be automated safely?

Q4. Could your readers detect AI-assisted drafts now, and would it change their trust if they could?

Q5. How defensible is your brand voice when ten new writers join in one quarter?

Q6. What facts, claims or topics should your publication refuse to generate under any circumstances?

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 Quillhouse Media'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