
Case Study: Designing an AI-Assisted Editorial System for Scalable Enterprise Communications

Context

In a fast-moving enterprise environment, communications teams are often expected to produce a high volume of content across multiple channels — newsletters, campaign emails, inserts, and executive communications — while maintaining consistent voice, accuracy, and strategic alignment. At the same time, resources are limited and timelines are tight.

The challenge was not simply writing faster, but scaling **quality, clarity, and consistency** without sacrificing editorial judgment.

The Challenge

The team needed to:

- Produce multiple recurring monthly assets across audiences and formats
- Maintain a consistent brand voice while adapting tone for different use cases
- Reduce production time without introducing risk or inaccuracies
- Allow for human editorial oversight and iteration
- Support reuse of content across channels without duplication fatigue

Traditional workflows were time-intensive and difficult to scale, especially for content that shared overlapping themes but required different framing.

My Approach

Rather than treating AI as a shortcut, I designed an **AI-assisted editorial system** focused on structure, constraints, and human review.

The goal was to create a repeatable workflow that:

- Separated **inputs** (source material, audience context, purpose) from **outputs**
- Applied clear editorial constraints around tone, length, and intent

- Enabled content reuse while avoiding redundancy
- Preserved human judgment at key checkpoints

I also created a guidebook to go with it so new team members or other teams could use it (and they have).

At a high level, the system followed this logic:

Source Content → Audience Context → Editorial Constraints → Draft Output → Human Review

This modular approach allowed the same core material to be shaped into multiple formats — newsletters, campaign emails, or inserts — while remaining accurate, relevant, and on-brand.

Safeguards and Editorial Controls

Because the work operated in an enterprise environment, the system was intentionally designed with guardrails:

- Clear separation between factual source material and generated language
- Structured inputs to reduce ambiguity and hallucination risk
- Mandatory human review before any final use
- Flexibility to override AI output when nuance or judgment was required

The system functioned as an **editorial assistant**, not a replacement for editorial thinking.

Results

The AI-assisted workflow delivered measurable impact:

- Reduced content production time by **20+ hours per month**
- Enabled the creation of **seven consistent monthly assets**
- Improved clarity and cohesion across related communications
- Was adopted by additional teams to support similar workflows

Most importantly, the system allowed the team to scale output **without diluting voice or intent**.

What This Demonstrates

This project reflects my approach to AI and content strategy more broadly:

- Designing systems, not one-off solutions
- Applying AI responsibly within enterprise constraints
- Balancing efficiency with editorial quality and trust
- Translating complex information into clear, human-centered communication

Details have been generalized to protect proprietary information.