



Case Study:

Automating Daily Editorial Workflows

Cumulus Media

Prepared by

 flatiron software

Executive Summary

Cumulus Media partnered with Flatiron Software to build an AI-driven daily show prep platform that combines intelligent automation with human editorial control. The system crawls up to 60 news sources, uses AI to select and summarize the most relevant articles per section, and provides a curation UI where the Cumulus team can review, adjust, and finalize the briefing before generating an editable Word document and distributing it to recipients via email.

Cumulus Media is one of the largest audio media companies in the United States, operating hundreds of radio stations and reaching millions of listeners across the country. Their on-air talent and editorial teams produce daily shows that require up-to-date, curated content from dozens of news and pop culture sources.

The Challenge

Cumulus Media's editorial and on-air teams depend on timely, well-organized content to prepare their daily shows. Monitoring dozens of news outlets, pop culture blogs, and entertainment sources manually was time-consuming and inconsistent — leaving producers without a reliable, unified briefing before each broadcast.

At the same time, Cumulus needed a solution that could be deployed securely inside their own AWS infrastructure, integrate with their existing workflows, and scale across editorial recipients without additional overhead.

Key Client Needs —

- Automate the daily aggregation of 40-60 news and pop culture sources.
- Deliver a structured, formatted daily briefing to editorial recipients every night.
- Use AI to summarize and consolidate content by section — Top Stories, Good News, Nerd News, and others.
- Build entirely within Cumulus' AWS environment for data control and security.
- Require zero manual effort from the editorial team after setup.

Solution

Flatiron Software designed and built an end-to-end AI-assisted show prep platform that puts editorial teams in control at every step. Rather than a black-box automation, the solution blends AI intelligence with a human curation layer ensuring the final briefing always reflects the team's editorial judgment, not just an algorithm's output.

Key Deliverables

1. Automated News Crawler

A modular scraping engine fetches articles daily from up to 60 sources using RSS feeds, HTML parsing, and APIs. The system handles rate limits, retries, and failures gracefully, logging all errors and achieving over 90% source coverage per run.

2. AI Summarization & Section Consolidation

Each article is summarized individually using Claude Sonnet via Amazon Bedrock, producing concise factual summaries under 150 words. A second AI pass using Claude Opus then clusters and consolidates summaries by editorial section — creating one coherent daily digest per category (Top Stories, Good News, Nerd News, and more).

3. Automated PDF Generation

A formatted daily PDF is generated from the consolidated section summaries, matching Cumulus' editorial layout. The document is timestamped, structured by section, and stored in S3 — both as a master copy and as individualized versions per recipient.

4. Individualized Email Delivery via SES

Each recipient receives a personalized email with their Word document attached, sent through Amazon SES. Deliveries are triggered on demand from the curation UI rather than on a fixed schedule, giving the team full flexibility over timing. Failed deliveries are retried automatically, logged, and accessible via the admin API for resend or inspection.

5. Flexible, Secure AWS Stack

The entire platform runs inside Cumulus' own AWS account. The AI layer is model-agnostic and orchestrated through Amazon Bedrock, supporting multiple foundation models without vendor lock-in. All sensitive data and content remain within Cumulus' infrastructure throughout the pipeline.

Delivery Process & Team Structure

Project Management & Transparency

Flatiron took full ownership of project management and delivery throughout the engagement. The team ran daily standups, maintained sprint planning with full visibility for Cumulus stakeholders, and held weekly demo and planning sessions to keep all parties aligned. Cumulus' engineering team was kept informed at every stage with minimal burden on their side.

Alignment & Quality

All code was written to Cumulus' engineering standards, with structured code reviews and documentation throughout. The AI pipeline was built as an independent module to minimize dependencies on existing Cumulus systems, allowing the project to progress in parallel without impacting any ongoing work.

AWS Services Used

- **Hosting & Compute:** Amazon ECS Fargate, AWS Lambda
- **AI & LLMs:** Amazon Bedrock (Claude Sonnet for summaries, Claude Opus for consolidation)
- **Database:** Aurora PostgreSQL (articles, summaries, deliveries, recipients, error logs)
- **Storage:** Amazon S3 (raw content and per-recipient .docx files)
- **Email Delivery:** Amazon SES
- **Scheduling:** AWS IAM, KMS, SSM Parameter Store, VPC
- **Monitoring:** CloudWatch Dashboards & Alarms, SNS alerts

Execution Insights

Challenge

Inconsistent data quality across 40-60 heterogeneous sources (RSS, HTML, APIs).

AI summarization needed to be factual and concise across highly varied article types.

Per-recipient PDF generation and delivery had to scale without performance degradation.

Admin operations (resend, recipient management, error inspection) needed a secure, code-free interface.

Solution

Built a modular scraper framework with per-source adapters, retry logic, and structured error logging into Aurora PostgreSQL.

Used Claude Sonnet (Bedrock) with tightly scoped prompts per section type, validated output length and tone before storage.

Separated master PDF generation from per-recipient copies stored in S3; SES delivery runs asynchronously with per-recipient status tracking.

Delivered a secure REST API with API key + IP allowlist authentication, covering full CRUD for recipients, delivery history, resend, and error inspection.

Outcome

- Fully automated daily pipeline — from source crawling to email delivery — with zero manual effort required from the editorial team.
- AI-generated section digests delivered to all recipients by 11 PM EST every night, consistently formatted and ready for show prep.
- Secure, scalable architecture built entirely within Cumulus' AWS account — no third-party data exposure, no vendor lock-in.
- Admin API giving the Cumulus team full operational control over recipients, delivery history, resends, and error monitoring.
- Platform future-proofed for additional AI use cases, with modular architecture ready for new sources, sections, or delivery formats.

Strategic Impact

This project demonstrated how a large media organization can harness Generative AI to eliminate repetitive editorial work at scale without compromising content quality or data security. Flatiron's independent delivery model allowed Cumulus to launch a production-grade AI system in under five months, with no disruption to existing operations and a clear roadmap for further automation.

Why Flatiron?

Flatiron delivers complex AI transformation projects at enterprise scale—fast, independently, and with no disruption to your existing teams. Our model lets you capture the value of AI while protecting your core business and product roadmap. From project management to technical delivery, we bring deep expertise, proven process, and clear accountability.

Ready to see what AI can do for your platform—without the risk or overhead?

Let's talk.

 marketing@flatiron.software

 flatiron.software

 Miami, US