

Cut editorial video-review time from hours to minutes — with an editor sign-off gate and full source traceability on every AI summary

Thomson Reuters

We built a production pipeline that ingests broadcast video, transcribes it and produces editorially-formatted summaries — cutting analyst review time from hours to an estimated few minutes per video, with a mandatory editor sign-off gate and traceability from every summary back to the source footage and transcript.

THE CHALLENGE

Editorial teams were spending hours manually reviewing long-form video to extract key moments, quotes and summaries for downstream products — a bottleneck at exactly the moments news breaks. In a newsroom, an ungrounded or misattributed AI summary is a reputational and accuracy risk, so a generic summarisation tool that could not cite its source or be reviewed before publication was unusable. The pipeline had to preserve editorial voice and keep a human in control of what ships.

WHAT WE BUILT

- Designed a serverless ingestion pipeline on AWS to handle large video files with chunked, fault-tolerant, observable processing
- Combined Amazon Transcribe for speech-to-text with foundation models on Bedrock for editorial-style summarisation against a custom prompt-and-style guide
- Made every summary traceable back to its transcript timestamp and source footage, so editors can verify any claim before it is used
- Added a human review surface with an explicit sign-off gate, so editors accept, edit or reject AI output and feed corrections back into the pipeline

OUTCOMES

- Cut analyst review time from hours to an estimated few minutes per video, with an editor sign-off gate and source traceability on every summary
- Pipeline scales to peak news cycles without manual intervention and with full processing observability
- Editorial voice preserved through prompt engineering plus a reviewer feedback loop
- No AI summary reaches a downstream product without human sign-off

STACK