

Sub-second search across decades of video that shipped as an official AWS solution

Thomson Reuters — Software Development Engineer (AWS)

Journalists at Thomson Reuters could search video archives by headline and date, but never by meaning, so decades of footage sat unused. I led the build that made it searchable by intent in under a second, published as an official AWS solution and showcased at IBC 2024.

38→4 min

ANALYST REVIEW
TIME

<1s

SEMANTIC SEARCH

2.8x

ANALYST
THROUGHPUT

IBC 2024

OFFICIAL AWS
SOLUTION

THE CHALLENGE

Media organisations sit on enormous video archives, decades of broadcasts, interviews and press conferences, but the metadata is shallow: headlines and dates, sometimes a transcript, rarely anything searchable by *meaning*. A journalist hunting "comments by central bankers about inflation expectations between 2022 and 2024" had nothing useful to query, so a valuable asset went unused and reporters lost hours.

WHAT I DID

- Ingest pipeline. Take a video, extract audio, transcribe with timestamps, chunk by topical boundaries, embed each chunk for semantic search.
- Retrieval and summarisation. A query hits OpenSearch (vector + lexical); the top chunks plus context go to a Bedrock model, which produces a summary citing each clip with its in-video timecode.
- A reference UI. A React front-end demonstrating the search-and-summarise loop, built so customers could adopt it as an AWS solution and customise it.

OUTCOME

Journalists could search decades of footage by meaning in under a second, turning an unused archive into a working reporting tool and saving the hours previously lost to manual hunting. The solution was published as an official AWS reference architecture and showcased at IBC 2024, the broadcast industry's flagship event in Amsterdam. Internal AWS teams have used it as the basis for further media-vertical work.

STACK