AI in Media Supply Chain and Distribution Workflows

TPFL Consulting – Tom Pflaum

Tom Pflaum

How Agentic AI is Transforming Media Supply Chain & Distribution Workflows

For years, media supply chain and distribution workflows have relied on a familiar mix of manual effort and rigid automation.

Operators and engineers spend countless hours downloading files, checking specs, verifying content, crafting FFmpeg commands, transcoding assets, and performing manual QC — all while referencing dense policy documents and delivery specifications.

Even when we automated with scripts (Python + FFmpeg, cloud-based Step Functions and MediaConvert) or workflow engines (low code/no-code) the systems remained rigid. Policies were hardcoded, exceptions required human intervention, and manual quality control was still the final gatekeeper.

Don’t just optimize the assembly line — build an army of experts.

At the recent SMPTE Bits-by-the-Bay conference, I presented a more powerful alternative. Here’s a summary of that talk.

Full Presentation Slides

A Real-World Example: Local Ad Ingest

A local ad for Lucy Chevrolet arrives via email: "I put the Lucy Chevrolet spot for next week in your Dropbox account..."

In the traditional manual workflow, an experienced operator would:

- Receive the email
- Download the file
- Analyze it with MediaInfo
- Verify the content and promo details
- Reference house format specs (MXF, 29.97i, 5.1 audio, etc.) and transcode policies (e.g., 3:2 pulldown for 24fps sources)
- Manually build and run an FFmpeg command
- Place the finished MXF file in the ad inventory

Over the past 20 years, multiple technologies have been developed to automate the process:

- Using scripting and ffmpeg
- Using no code/low code workflow engines
- Using step functions (AWS lambda) for cloud deployments

The policies were hardcoded into the scripts or the workflow specifications.

Enter Agentic AI

What if we could close the loop with AI — not just for transcoding, but for intelligent orchestration of the entire workflow?

Using the OpenClaw agentic framework, the LLM gains real skills:

- Monitors email/Dropbox

- Downloads new ads
- Analyzes technical specs
- Applies the right transcoding rules
- Transcodes the video
- Places files in the correct inventory folder

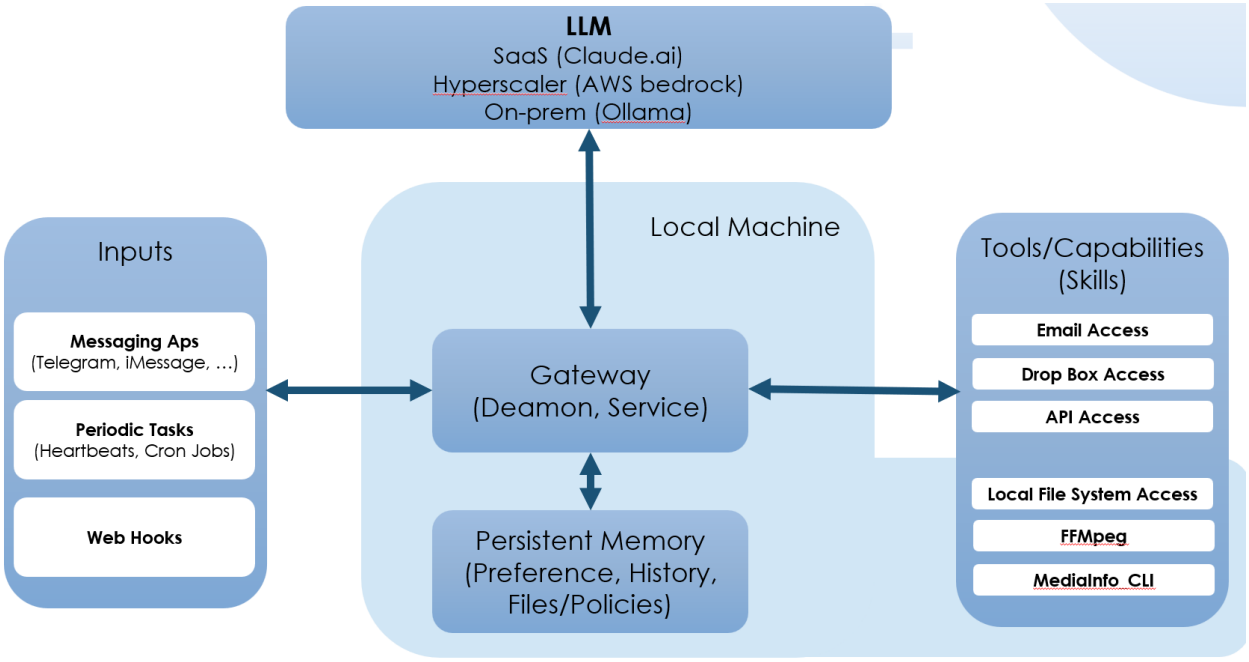

The LLM provides the intelligence, OpenClaw provides the skill and the application specific knowledge.

LLMs Actually Know Video Tech

LLMs demonstrate impressive knowledge of video workflows and FFmpeg:

Example 1 (Claude Opus): Asked to trim 30 seconds from a DNxHD file (chronicle.mov) and output MP4. The model correctly noted DNxHD isn't supported in MP4, suggested + performed a transcode to H.264, then fixed an audio issue (IPCM → AAC).

Example 2 (24fps → 29.97i): GPT-4o provided a solid telecine command with 3:2 pulldown. Gemma 4 was more cautious but offered valid alternatives. Results vary by model, but knowledge is strong — and you can always document edge cases in your policy files.

Below are some chats I had with various LLMs showcasing deep understanding of video technology and tool usage:

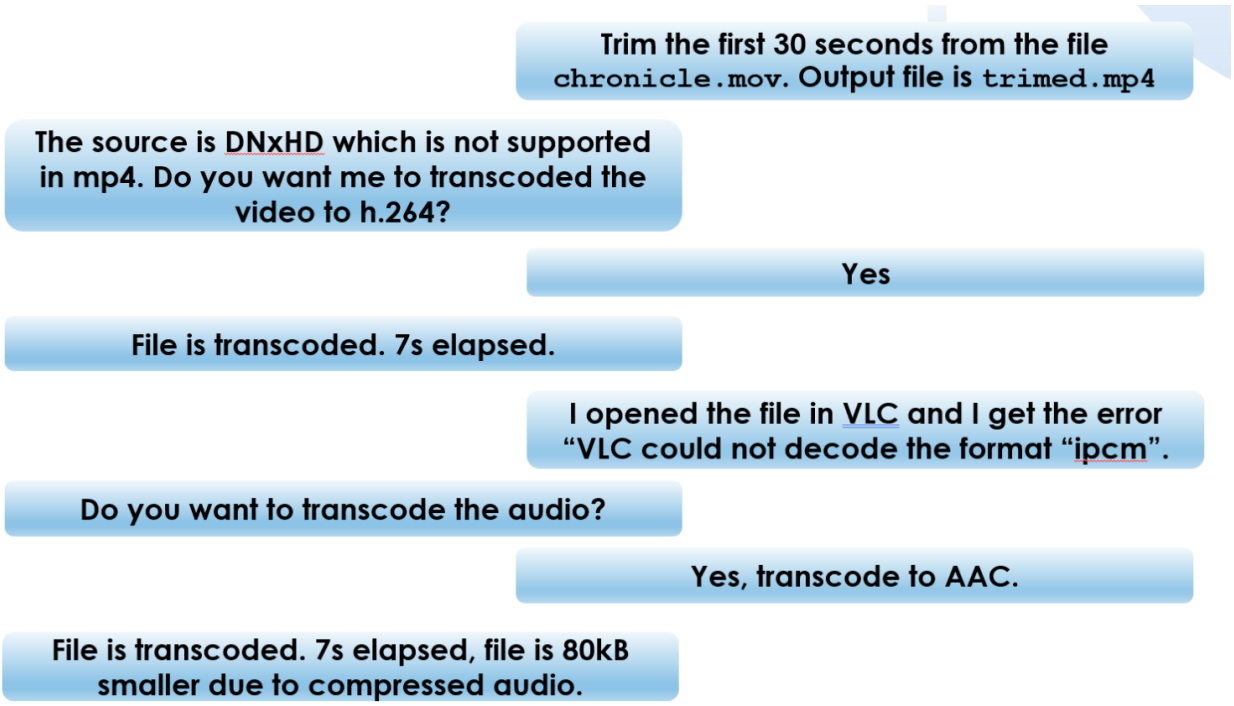

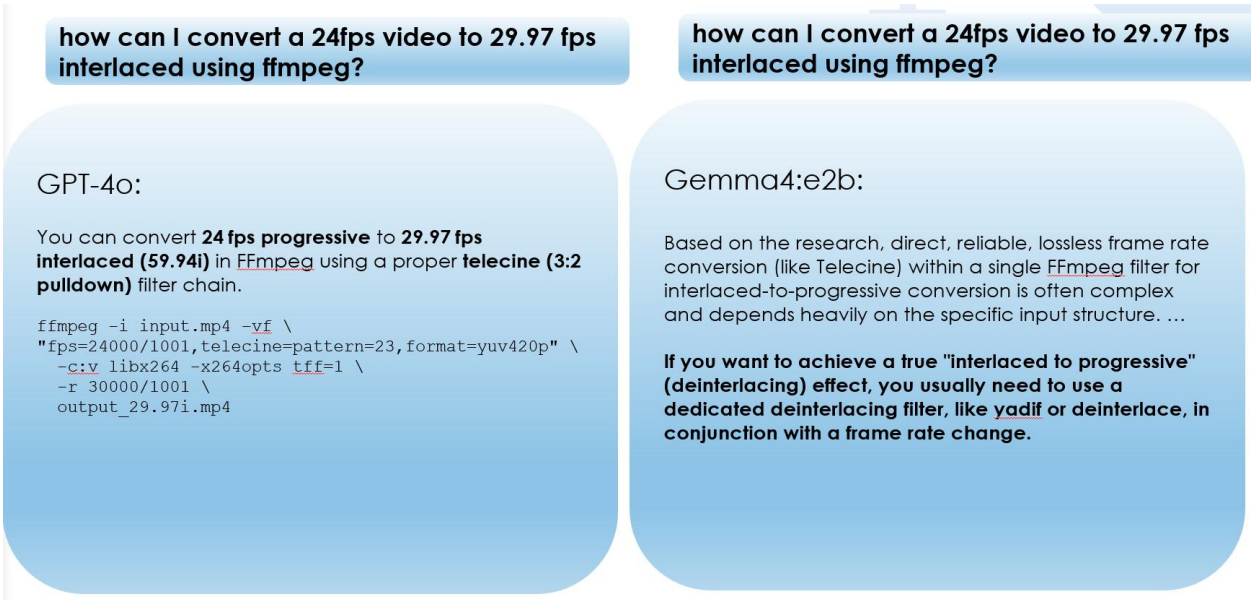

Full Automation Setup using OpenClaw

Policies live in persistent memory as human readable Markdown files. Using OpenClaw's cron functionality, the agent runs every 5 minutes with a prompt like:

“If there is a new email from an advertiser saying they uploaded a file to Dropbox, download the file, transcode it to house format, and place it in the ad inventory.”

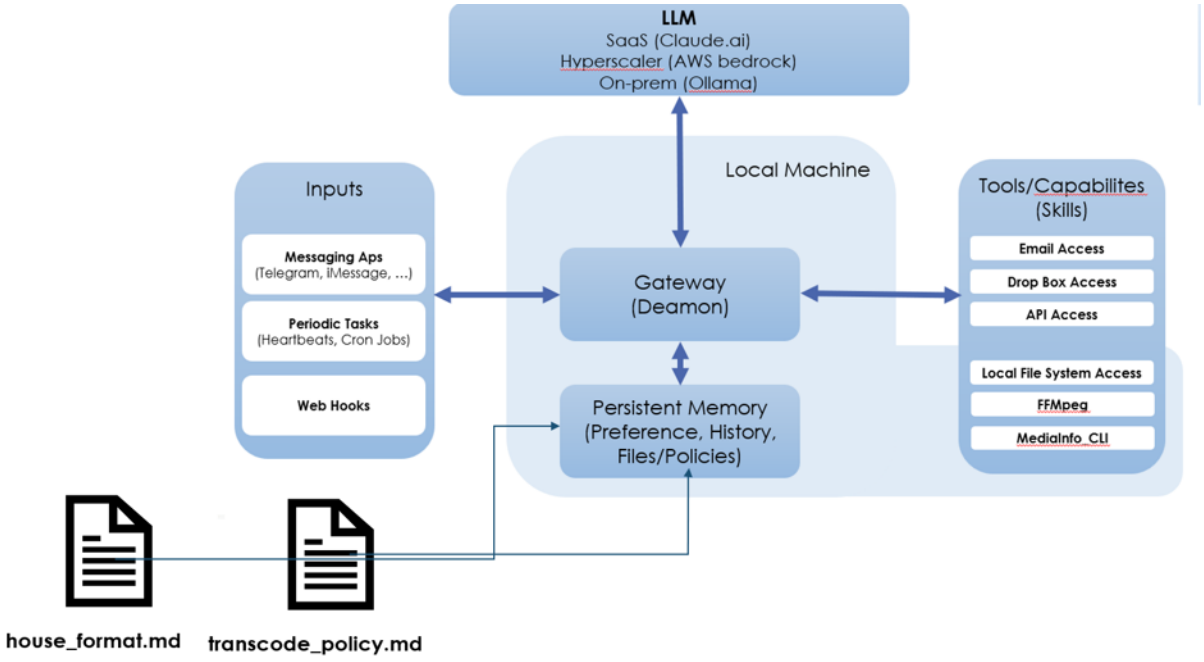

Why OpenClaw?

- Open source with strong community traction
- Excellent pre-built skills and session/memory management
- Direct filesystem + shell access
- Works with virtually any LLM (Claude, Gemini, Ollama, Bedrock, OpenRouter, etc.), on-prem, in the cloud or SaaS.
- Can be deployed 100% locally

Alternatives include LangChain, Temporal, or cloud agent platforms, but OpenClaw excels for rapid prototyping with local tools.

Summary

LLMs + agentic frameworks like OpenClaw can successfully orchestrate complex media supply-chain tasks.

Don’t just optimize the assembly line — build an army of experts.

In the next article in this series, I will be "building out" the agent to manage things like scaling, reporting and semantic QC. I will also address which AI model to choose.