📝 Release Notes

New Feature: Video Remix Agents (Open Beta)
Scaling your video ad production just got easier. We are thrilled to introduce the Video Remix Agents (Open Beta), a suite of three AI-powered agents designed to transform raw video content into polished, performance-optimized short-form videos for platforms like TikTok, Reels, and Shorts.

The End-to-End Remix Pipeline: Whether you are combining glossy brand assets with raw UGC, these agents automate the heavy lifting so you can easily remix your existing content into **10x the amount of social spots**.

🔄 The Pipeline & Standalone Use Cases

To successfully generate a video, these agents must always be run sequentially in a specific order: Splitter > Script > Stitch.

1. Video Splitter Agent (Stage 1 or Standalone)
While this is the first step in the remix pipeline, **it is also an incredibly powerful standalone tool. What it does:** It dynamically scrubs your raw files to detect physical shot changes, camera angle shifts, and precise cut points. If you just need to rapidly extract the best scenes from a massive raw shoot without storyboarding, you can use this agent on its own.

- **The Output:** It generates a scene manifest with timestamps, file mappings, and scene metadata. This manifest serves as the foundational input for the next stage.

Standalone Use Case: If you just need to rapidly extract the best scenes from a massive raw shoot without storyboarding, you can use this agent on its own to log and clip your footage.

2. Video Remix Script Agent (Stage 2)
Acting as your AI performance creative director, this agent takes the scene manifest from the Splitter and writes structured storyboards.

- **With a Video Template:** You can optionally select a video template to guide the structure. When a template is selected, the agent follows it exactly: matching the scene count, layer structure, and timing. *Note: The chosen template must have only one video layer per scene.*
- **Without a Template:** If no template is selected, the system uses intelligent defaults. It automatically creates 4-6 scenes with 1 video layer each, with a maximum duration of 60 seconds.

- **The Output:** It creates structured storyboards with scenes, layers, text overlays, and timing. It also includes a compliance checklist at the end of the storyboard output.

Standalone Use Case: If you already have a manual list of video clips and timestamps, you can feed them directly into this agent to rapidly brainstorm narrative structures and generate multiple storyboard variations for A/B testing without committing to rendering them yet.

3. Video Stitch Agent (Stage 3)
This is the last step that produces the rendered output.

- **What it does:** It reads the scene references, timecodes, and layer information from the generated storyboard.
- **The Output:** It takes the storyboard output and stitches the referenced clips/scenes into a cohesive final video.

Standalone Use Case: If your creative team prefers to manually write their own scripts or storyboards using exact timestamps from existing assets, you can feed that manual script directly into the Stitch Agent to automatically assemble and render the final video, bypassing the AI scripting phase entirely.

🛠️ Two Ways to Work: Chat vs. Workflows

Depending on your production goals, you can test these agents via Chat (interactive) or Workflows (node-based automation).

- **Workflow Mode (Best for Automation & Scale):** Workflows provide a visual, node-based canvas. You connect agent nodes with lines to orchestrate automatic data flow between nodes. Crucially, Workflows support larger collections, allowing you to process massive batches of raw video at once and scale your production.
- **Chat Mode (Best for Exploration):** This provides an interactive, conversational experience. In Chat, you manually select agents per message and transition between stages. The Chat mode supports up to 14 videos per session.