

# #SORA WON'T SAVE YOUR PITCH. HERE'S WHAT WILL.

Let's start with an **honest take** most explainer studios won't give you:

if you can pick up Veo 3, Runway Gen-4, or whatever AI video model is winning benchmarks this week and produce something that lands with your investors or enterprise buyers, **you absolutely should**. That's the right call. Go do it.

This post exists to explain exactly when and why that stalls — and what the **hybrid model that actually works** looks like in 2026.

THE CFO

## BUDGET REVIEW

Video Production	\$120,000
AI Tools & Credits	\$28,000
Creative Direction	\$45,000
Design & Story	\$38,000
Edits & Revisions	\$26,000

TOTAL \$257,000

ROI > HYPE

THE FOUNDER

▶ Veo 3

Generating...

73%

Runway

Gen-4

Rendering...

62%

⌘ Sora

Enhancing...

48%

Prompt

"Cinematic shot of our product transforming the future..."

GENERATE

STORY?  
PROOF?  
IMPACT?

INVESTOR  
FEEDBACK

— TOO VAGUE  
— MORE PROOF  
— SHOW ROI  
— WHY NOW?

V3  
PENDING

DECK v7

WIP

DECK v6

WIP

DECK v5

SHIP  
PITCH  
REPEAT

NARRATIVE FLOW (DRAFT)

# THE TOOLS ARE GENUINELY EXTRAORDINARY.

THAT'S NOT THE PROBLEM.



The AI video generation market has **matured fast**. The options available in 2026 are **materially better** than what existed two years ago.



Veo 3 generates **synchronized audio** alongside video.



The current version of Kling handles **longer clips** at subscription price points that would have seemed **implausible** in 2023.



Runway Gen-4 gives you **creative control** that professional users have noted offers meaningful **creative flexibility**.



And then there's the Sora story, which is worth telling clearly. According to OpenAI's public announcements, the company **discontinued** the Sora app in early 2026.

The product generated impressive demos but **struggled to convert that novelty** into sustained, practical use. Public reporting suggested early interest did not translate into **lasting engagement** for users trying to do something specific and repeatable with the output.

THE TAKEAWAY ISN'T THAT AI VIDEO IS BROKEN. IT'S THAT A TOOL ISN'T THE SAME AS A WORKFLOW — AND A WORKFLOW ISN'T THE SAME AS A STORY.



# WHAT ACTUALLY STALLS MOST FOUNDERS

Here's a pattern worth examining. A founder with a genuinely complex product — deep tech, B2B SaaS, AI/ML, biotech, quantum sensing, whatever — spends a weekend with a video generation tool. They produce a few clips. The visuals look fine. Maybe impressive.

**But something is off.**

**The output feels generic.**

It doesn't actually explain the product. It just... looks like other explainer videos.



So they try again with a better prompt. Still off. Third attempt, they give up.



*That's not a tool failure.  
That's a narrative failure.  
And the tool made it faster.*




AI amplifies whatever brief you give it. A confused brief produces a confused video — **at machine speed.**

The gap between what these tools can generate and what actually persuades a buyer or investor isn't in the rendering pipeline. It's upstream, in the work that has to happen before anyone opens a video editor: figuring out who you're actually talking to, what they already believe, what language they use, and **what single idea they need to leave the room holding.**



That work is hard. It's not automated. And most founders, understandably, **haven't been trained to do it.**



# THREE SPECIFIC PLACES PURE AI OUTPUT STILL BREAKS FOR COMPLEX PRODUCTS

Even with the best available models and a solid brief, there are **three recurring failure modes** for complex B2B tech explainers.

## 1 UI MOCKUPS.




If your product has a dashboard, a workflow interface, or any screen-based interaction that a buyer needs to understand, AI video tools will frequently hallucinate it. The generated UI will look plausible and be completely wrong.

**Brand colors will shift. Buttons will appear that don't exist. Data will be nonsense.** For a buyer evaluating your product, that's not a visual inconvenience — it's a trust problem.

Manual motion graphics or direct screen capture are among the most reliable solutions here.

## 2 CHARACTER CONTINUITY.



Character consistency limitations apply directly to B2B explainers. You can't reliably generate a character who looks exactly like an existing persona across a series of clips. If your narrative follows a buyer through a workflow — which is often the clearest way to explain a complex product — maintaining visual consistency across scenes requires either **extensive prompt engineering, heavy post-production work, or a manual fallback.**


## 3 BRAND-LOCKED VISUALS.



If your brand has specific colors, logo treatments, or visual identity rules that matter (and for investor-facing and enterprise sales materials, they often matter a lot), raw AI output will drift. Every generation is a coin flip on whether the brand cohesion holds.

**Manual augmentation remains the most dependable approach to ensuring brand consistency,** though model fine-tuning is an emerging alternative worth watching.

NONE OF THIS MEANS AI VIDEO TOOLS AREN'T USEFUL. IT MEANS THE USE CASE DETERMINES THE **WORKFLOW — NOT THE OTHER WAY AROUND.**



# THE HYBRID MODEL THAT ACTUALLY WORKS IN 2026

Here's what senior creative direction plus an agentic production stack looks like in practice.

## 1 SENIOR CREATIVE DIRECTION FIRST

A senior creative director handles the narrative architecture first. This means a real discovery conversation with the founder or team — understanding the product deeply enough to find the **human-scale version** of a technical truth. Not dumping it down. Going deep enough that the explanation becomes **accurate and clear**. This is the layer that **can't be automated**, and it's the layer that determines whether anything downstream is worth shipping.

## 2 AI PRODUCTION HANDLES VELOCITY

Once the story exists, AI production handles velocity. Scene generation in **Veo or Runway**, voiceover through **ElevenLabs**, storyboard iteration at **compute cost** rather than agency day rates. What might previously have required a multi-person motion team can now be executed by **one studio operator** working across an **agentic production stack**.

## 3 MANUAL MOTION GRAPHICS FILLS THE GAP

Manual motion graphics step in wherever AI output fails the quality bar: **UI mockups** built to spec, **character consistency** across scenes, **brand-locked** visuals that have to be exactly right. The deliverable is the **receipt**, not the methodology. If AI gets it done, great. If it doesn't, manual production **fills the gap** without ceremony.

## THE RESULT

A 60-90 SECOND EXPLAINER DELIVERED IN TWO TO THREE WEEKS, AT A FRACTION OF TRADITIONAL AGENCY RATES.



1

2

3



# THE HONEST VERSION OF THE ROI QUESTION



**There's a real objection worth addressing directly:** if AI video tools are this accessible, why pay anyone?



**Honest answer:** for many use cases, you shouldn't. High-volume social content, internal training videos, quick product updates — these are often entirely reasonable **DIY territory in 2026**, and the tools that handle them are good enough that the cost comparison isn't close.



**The calculus changes when the stakes change.**

A pitch to a **Series A investor**.  
 An **enterprise sales cycle** where the buyer won't read a white paper.  
 A **product launch** where the homepage video is the first impression for a market that's never heard of you. An **IR communication** explaining a complex technology thesis to fund managers who need to approve a capital allocation.



**In those moments,** the cost of a generic output — or a clear output built on the wrong story — **isn't the production fee. It's the deal that doesn't close. The round that gets passed on. The buyer who stays confused and stays away.**



**AI VIDEO TOOL**  
 Generate video from prompt

SCENE 12  
 AI GENERATED

VOICEOVER  
 OPTION 03

REVISE HOOK AGAIN?

WIP

DECK v7

FUEL & CAFFEINE

NEEDS BETTER STORY

PITCH DECK v3  
 INVESTOR DECK v5  
 PRODUCT LAUNCH v2  
 IR UPDATE v1

CONVINCE ME THIS MOVES THE OUTCOME, NOT JUST THE PIXELS.


BUDGET REVIEW	
ITEM	COST
DIY (AI TOOLS)	\$
IN-HOUSE	\$\$
AGENCY / STUDIO	\$\$\$
OUTCOME RISK	?????
TOTAL IMPACT ON OUTCOME	?


SKEPTRICAL BY DEFACNT.  
 CONVINCED BY EVIDENCE.

ENTERPRISE DUE DILIGENCE  
 OUTCOMES > OUTPUTS


\* Disclosure: The remainder of this section describes our own services at Infrairs. We built and run [Supramono](#), [CenterOS](#), [Evotron Studio](#), and [VirtualSpace](#) — and we used the same hybrid approach described in this post to explain each of them. We're sharing those as examples of the workflow in practice; you can judge the outputs directly rather than take our word for it.

# WHAT TO DO NEXT




 If your product is complex and your pitch isn't landing the way you need it to, start with the diagnosis.

**What's actually confusing your buyers? Is it the product, the message, or the medium?**

 If you want a structured answer to that question in one week, our **Clarity Workshop** is \$1,500 NZD and delivers a 30-second pitch and 60-second storyboard draft built around your actual buyers.

-----

If you already know the story and just need it produced fast, a **Single Explainer** is \$3,500-\$8,000 NZD delivered in two to three weeks.

 And if you want to try Veo or Runway first, we'll tell you exactly which prompts to start with.

**If you stall on the third attempt, you know where to find us.**



**TALK TO US AT INFRAIRIS →**



# Infrairis

<https://startups.infrairis.com>