

STOLEN 67

a Kogane Bay story — an AI-generated short film

RUNTIME 8:14	DELIVERY 1920×804 · 2.39:1	LANGUAGE English	SUBMITTED September 2026
-----------------	-------------------------------	---------------------	-----------------------------

1 STATEMENT

Every frame of visual content in this film is generated by artificial intelligence. There is no photographed footage, no live-action plate, no stock library material and no photographed performer anywhere in the work. All voices are synthesised. All music is generated. The only non-generative picture elements are seven title and type cards, which are rendered deterministically from code (see §3).

This document lists every model used, what each was used for, and how. Generation logs supporting every statement here are retained and available on request under §4.7.

2 MODELS AND TOOLS

Counts are model invocations recorded across the production, including takes that were reviewed and rejected. They are not counts of shots in the finished film.

ROLE	MODEL	PROVIDER	INVOCATIONS
Still images character anchors, location plates, shot keyframes	seedream-v5-pro seedream-v5-pro-edit	Venice	339
Still images — earlier phase	gemini-3-pro-image gemini-3.1-flash-image	Google	122
Still images — incidental	nano-banana-pro-edit	Venice	3
Moving image every shot in the finished film	seedance-2-5-image-to-video seedance-2-0-image-to-video	Venice	316
Voice 46 of the 47 narration cues in the finished film	tts-elevenlabs-turbo-v2-5	Venice	69
Voice — earlier phase one cue survives into the film	tts-minimax-speech-02-hd	Venice	140
Music score and source cues	elevenlabs-music	Venice	38
Evaluated, not used in the film	flux-3-first-last-frame-to-video	Venice	2
Evaluated, not used in the film	gemini-omni-flash-preview	Google	1

Takes from evaluated-and-rejected models are quarantined on disk and are not present in the delivered master. Every model above is reached through one of two API providers, Venice and Google; no model was invoked directly from its originating vendor. Counts are drawn from 465 per-invocation run logs.

3 HOW THE MODELS WERE USED

The film was produced through a purpose-built pipeline, not through ad-hoc prompting. Character and location descriptions live once, in source files; prompts and manifests are compiled build artefacts. Every render carries a seed derived from its identifier and prompt hash, so any frame in the film can be reproduced.

- **Anchored identity.** Each character exists as one locked anchor image. Every later frame of that character is an edit conditioned on that anchor rather than a fresh generation, which is how the film holds a consistent cast across 72

shots.

- **Keyframe, then motion.** Every shot is composed first as a still image, reviewed, and promoted; the video model then animates that approved frame. No shot is generated from text alone.
- **The look is applied once, in code.** Grain, halation, colour grade and the 2.39:1 crop are applied identically to every shot by the finishing stage. They are never requested from a generative model, because 72 individually convincing grain structures read as false the moment they are cut together.
- **Type is authored, not generated.** All titles, cards and on-screen graphics are drawn deterministically from code at the delivery frame; no title, credit or caption was produced by a generative model. The models do render incidental text *inside* the footage where a scene contains it — a diegetic security-camera timestamp in one three-second insert — and that text is a product of the video model rather than authored type.
- **Audio is assembled, not generated whole.** Voice, music and ambience are generated separately and mixed deterministically to a measured target of -18.0 LUFS integrated, -1.1 dBTP.

4 HUMAN AUTHORSHIP

The screenplay, the story, the characters and every line of dialogue and narration were written by the director. No model wrote narrative content.

Each shot was designed, framed, prompted, reviewed and either selected or rejected by a person. The edit — shot order, duration, rhythm, sound placement — is entirely human. Rejected takes outnumber used takes by a wide margin, and that selection is the authorship: the models executed instructions and produced candidates; a person decided what the film is.

5 COMPLIANCE

- **No real-person likeness or voice** (§4.5). Every character is a synthetic identity created for this film from generated anchor images. No performer was photographed or recorded. All speech is model-synthesised; no voice was cloned from a real person's recordings by us.
- **No partner marks** (§6.6). No Venice or MoonPay name, logo or mark appears anywhere in the work. Tooling is credited in this disclosure, not on screen.
- **Consistency.** The anchor discipline described in §3 exists specifically to prevent drift across the film.

6 EVIDENCE RETAINED

Available in full on request under §4.7: an append-only provenance record of **917 permanent entries** covering every import, promotion, normalisation and finish; **465 per-invocation run logs** recording the model, parameters, seed and cost of every call made during production; all source block files, compiled prompts and manifests; and every rejected take, retained rather than deleted.

I confirm that the information in this disclosure is complete and accurate to the best of my knowledge, that all visual content in **STOLEN 67** is AI-generated, and that no real person's likeness or voice appears in the work.

SIGNATURE — DIRECTOR / SUBMITTER

DATE