

THE ARCHITECT'S STILLNESS

Mission Control video-model case study | Published May 9, 2026



THE ARCHITECT'S STILLNESS

10 Spark video models | 3 completed | 7 runtime failures | prompt-fidelity review

Mission Control case study | Published May 9, 2026

Executive Read

Ten Spark video models were run against the same psychological-horror cinematic prompt. Three produced MP4s. Seven failed before reviewable output due to memory pressure, loader incompatibilities, timeout, or model-specific runtime API mismatches.

Best completed creative result: CogVideoX 5B. It captured the dark dome/cathedral atmosphere better than the other completed clips, but still missed the prompt's required cyborg brain, child/doll beat, soldier mind-control cutaway, and reflected-eye finale.

Models	Completed	Failed	Best Prompt Match	Avg GPU
10	3	7	3.5/10	85.7%

Prompt Fidelity Review



Model	Status	Latency	Prompt	Verdict
Wan 2.1 T2V 1.3B	completed	343.94s	0.5/10	Rendered a valid MP4 but the image is effectively black, so it fails the visual and narrative req
CogVideoX 2B	completed	209.27s	2.5/10	Produces a coherent Gothic/industrial surface with cracked-glass energy, but it is almost static
CogVideoX 5B	completed	714.53s	3.5/10	Best completed clip: it suggests a cracked dome/cathedral mood, but remains mostly environ
Mochi 1 Preview	failed	20.73s	n/a	CUDA out of memory during model load/generation.
SkyReels V2 DF 14B 540P	failed	7200.01s	n/a	Exceeded the 7,200,000 ms HF video runtime timeout at 960x544, 49 frames.
LTX Video 13B Distilled	failed	48.88s	n/a	CUDA out of memory despite supplied source keyframe.
HunyuanVideo I2V	failed	300.56s	n/a	Runtime tensor/indexing failure after source-image handoff; likely adapter/source conditioning
SkyReels V1 Hunyuan I2V	failed	5.75s	n/a	Repository layout mismatch: missing model_index.json for the current diffusers loader path.
SkyReels V2 I2V 14B 540P	failed	607.73s	n/a	Tensor channel mismatch in image-conditioning path; model adapter does not match runtime
Stable Video Diffusion XT	failed	7.71s	n/a	Runtime API mismatch: guidance_scale was passed to a pipeline that does not accept it.

Failure Analysis

Mochi 1 Preview: CUDA out of memory during model load/generation.

SkyReels V2 DF 14B 540P: Exceeded the 7,200,000 ms HF video runtime timeout at 960x544, 49 frames.

LTX Video 13B Distilled: CUDA out of memory despite supplied source keyframe.

HunyuanVideo I2V: Runtime tensor/indexing failure after source-image handoff; likely adapter/source conditioning incompatibility.

SkyReels V1 Hunyuan I2V: Repository layout mismatch: missing model_index.json for the current diffusers loader path.

SkyReels V2 I2V 14B 540P: Tensor channel mismatch in image-conditioning path; model adapter does not match runtime input shape.

Stable Video Diffusion XT: Runtime API mismatch: guidance_scale was passed to a pipeline that does not accept it.

Operational Findings

Vitals covered 292 samples from 2026-05-08T23:30:11.518Z to 2026-05-09T02:08:58.363Z.

Average GPU utilization was 85.7%, p95 was 96.0%, and max observed utilization was 96.0%.

Peak temperature was 83.0C and peak observed power was 93.66W.

A timed-out SkyReels remote process survived the SSH timeout and briefly overlapped with a later I2V job. The stale PID was killed manually; the queue then completed cleanly.

Recommendations

- Add low-memory presets for Mochi, LTX, and 14B SkyReels before public default use.
- Exclude unsupported kwargs per pipeline, especially Stable Video Diffusion guidance_scale.
- Split complex cinematic prompts into keyframe-first workflows with explicit source frames for each major beat.
- Fix image-conditioning adapters for Hunyuan/SkyReels I2V before judging creative quality.
- After SSH timeouts, explicitly terminate the remote process group to prevent stale GPU jobs.