

Creative Stack — one-page brief

Local AI video production — generation, surgical editing, 4K finishing — on hardware you control. Unreleased footage never leaves the building. Public repo: github.com/MushiSenpai/mushishi-creative-stack

What it is

- **Generation:** FLUX.2 stills, Wan 2.2 / HunyuanVideo video, shipped as six named, tested ComfyUI workflows (importable JSON).
- **Editing:** object removal (VOID) with SAM3 video-mask tracking — including the object's reflection; masked edits (SAM3 + VACE); a **forensic bridge** where a local vision LLM measures the scene first and writes the mask + fill constraints itself, so the diffusion model gets specifications to obey, not room to hallucinate.
- **Finishing:** 60 fps interpolation (RIFE) + 4K upscale (SeedVR2).

Measured (not projected)

JOB	MEASURED	NOTE
Object removal, representative clip	180 s	clean — no ghost, background reconstructed
Object removal, worst case	residue	subject behind glassware, low light — stated up front
Forensic-bridge pipeline, end to end	203 s	analysis → auto mask/fill prompts → render, unattended

What it can't do yet: liquids, heavy occlusion, and low-light targets defeat the mask — the published case studies show the failures, not just the wins. Hard footage is quoted only after a **free sample frame**, and I'll say honestly if the tech isn't there.

Engagement

- **Object removal:** Draft SGD 12/sec (min 80) · Production SGD 25/sec (min 150) · Delivery 4K SGD 40/sec (min 240), revision included.
- **Turnkey pipeline deployment** for studios/agencies — the whole stack on your hardware, with the workflow catalogue and runbook.

Madhan Kumar Reddy · Singapore · reddy.madhankumar.sg@gmail.com · theinvalid.me — every number above is a measured run on one RTX 5090 (32 GB); benchmarks and failure logs are public in the repo. Generated 2026-07-02.