

Internal memo — Tencent ARC has not published official Wan 2.5 documentation (as of 2 Oct 2025). All notes below come from Instavar pilots on NDA hardware. Please keep this draft internal until a public release lands.

Why we trialled the preview

We generate regulated 1080×1920 @ 30 fps advisor videos. Motion-controlled B-roll is the bottleneck: Wan 2.2 Animate + VACE MV2V deliver usable clips but require manual ambience and depth fixes. The Wan 2.5 preview hinted at two upgrades worth testing:

- Native ambience alongside the video, saving Foley passes.
- More stable dolly/slider/choreographed moves with less geometry warping.

Pilot setup at a glance

- **Hardware:** dual NVIDIA L40S pod, 64 GB VRAM per card.
- **Runtime:** 25–40 s per 10 s clip (comparable to Wan 2.2).
- **Aspect ratio:** locked to 9:16.
- **Control surface:** provisional MCP method `wan25.generate_broll` behind a feature flag; inputs validated to ≤ 12 s duration and a small enum of camera/audio presets.

What we observed

Clip characteristics

- Solid up to ≈ 12 s; anything longer drifts or ghosts.
- Texture retention beats Wan 2.2 on fabrics, lighting, and reflections.
- Ambient stem renders $\sim 92\%$ of the time. When it drops, it goes silent for the whole clip.
- Slider/dolly/crane tokens hold their intent across seeds better than Wan 2.2 replacement mode.
- Ambience ships mono at roughly -18 LUFS. We still layer licensed music and compliance VO in Remotion.

Timeline schema tweak

We added a provisional generator entry:

```
{
  "tool": "wan25",
  "mode": "t2v",
  "prompt": "morning sun across a glass-walled trading floor, advisors reviewing screens",
  "camera_path": "slider-right",
  "mood_audio": "subtle city ambience",
  "duration": 9,
  "seed": 182903,
  "qa": {
    "clip_reject_on": ["motion_glitch", "ambient_dropout"],
    "fvd_budget": 320
  }
}
```

camera_path and mood_audio map to the prompt tokens we saw responding in the preview build. Everything is versioned in `timeline.json` (so we can revoke the feature quickly if the API surface shifts).

QA gates we enforced

- \times \times LatentSync \times \times stays in play. Wan 2.5 ambience does not carry dialogue; human speech still routes through HunyuanVideo-Avatar + Azure TTS.
- \times \times Ambient watchdog: \times \times rerun if RMS falls below -50 LUFS for >300 ms mid-clip.
- \times \times FVD cap: \times \times keep below 320 to avoid structural drift on overlays.
- \times \times Seed logging: \times \times every job pins `seed` + QA verdict for later analysis.

Outcomes so far

- **14 pilots** **13 acceptances** after QC (92%).
- **Interpolation load down 46%**. Only four clips needed RIFE to bump 24 fps footage to 30 fps.
- **Editorial savings ≈40 minutes** per final deliverable by skipping manual ambience.

What still blocks production rollout

- The control tokens are undocumented; Tencent could change them without notice.
- Ambience dropouts require monitoring and fallback Foley.
- Compliance needs explicit C2PA labelling for synthetic ambience before we can ship externally.

Immediate follow-ups

1. Capture 30 additional pilot clips to tighten QA stats (FVD distribution, ambience reliability, reviewer sentiment).
2. Test Wan 2.5 on localized billboard replacements to see if it can displace the Wan 2.2 Animate VACE combo.
3. Keep watching Tencent ARC channels for an official release; once a spec exists we can promote this memo to a customer-facing post.

Internal: Author — Production Engineering. Do not forward outside the advisory team until the preview NDA lifts.