THE JEV LAUNCH, SIX DAYS LATER · RECEIPTS ONLY

I read the launch like a fraud examiner.

Hype and proof arrived together, as they always do. I sorted them. Six slides, sources at the end.

The scorecard:

Manufactured.

Awesome-list spam **within 6 days** of launch

A waitlist built as a **hype mechanic**

Reddit, 230-point thread: **"astroturfed marketing campaign"** — top comment agreed, 304 upvotes

~30 promo posts a day across platforms

Smoke. Documented.

Validated.

2,024 real production decisions replayed: 97.3% vs 94.6% · 151ms vs 4s · **$0.17 vs $5.51**

Vercel swapped it into production: **5–18×** faster classification

A neutral referee confirmed the reranking win — zero training

Fire. Also documented.

Who ran the race?

The 193.6×/444.6× headlines were **the vendor's own race** — against an average of rival frontier models, on workflows the vendor picked.

They admitted it. The neutral referee found where it **loses**: big specialized classification jobs.

Aim ≠ landing.

Both columns true.
That's normal.

This is every launch now. The skill isn't sniffing fraud or swallowing hype — it's reading receipts.

Read the receipts,
not the launch.

WHAT I'D BUILD ON THE VALIDATED HALF

Routing. Guardrails. Classification.

The plumbing decisions your agent makes thousands of times a day — now cheap enough to stop rationing. That's an architecture conversation, and it's my favorite kind. My inbox is open.

Erkin Tek · ships AI