444.6× CHEAPER · 193.6× FASTER · FOOTNOTE BELOW

444× cheaper. Read the footnote.

This week's AI launch printed two gigantic numbers. Both true. Both worth reading twice. Six slides.

Bring your skeptic hat.

RACE ONE

The vendor's own track.

THE NUMBERS

193.6× faster · 444.6× cheaper — Jev against an **average of rival frontier models**, on workflows **the vendor chose itself**.

THE CREDIT

They say so themselves, in the fine print. **That honesty matters.**

RACE TWO

The neutral referee.

WHERE IT WON

Independent test (Parallel Web Systems): Jev **matched a trained reranker with zero training**. Genuinely impressive.

WHERE IT LOST

Big, specialized classification jobs — **dedicated models still win** on cost and speed when you own the data.

Which number survives production?

THE PATTERN

Vendor numbers open meetings. **Independent numbers sign contracts.**

THE SCAR TISSUE

I've bought infrastructure on both kinds. **Only one kind survived contact with production.**

Self-reported numbers point where the vendor aims.

Independent tests tell you where it lands.

Ask for the referee.

YOUR TURN

Before your team rebuilds on this week's 444×: who ran the race — and who refereed it?

Sources in the first comment: typesafe.ai launch claims · Parallel Web Systems independent eval (Sept 18).

Erkin Tek · ships AI