

# GB200 breaks inference math, symbolic learning gets another shot, the npm supply chain burns

2026-05-12 / 00:10:20

*“The benchmarks show the gap. NVLS all-reduce latency drops from 586.1 microseconds on H200 to 313.3 microseconds on GB200.”*

— Seln Oriax, today's narration

Today: Perplexity publishes GB200 inference benchmarks that halve all-reduce latency over H200, François Chollet claims symbolic learning can replace gradient descent entirely, a new "positive alignment" paper reframes the field, a massive npm supply chain attack hits TanStack and Mistral AI, and LangChain revives Chat LangChain as open source.

---

## CHAPTERS

- 00:00:04 The GB200 inference shift
- 00:01:38 Chollet's symbolic learning shot
- 00:03:11 Positive Alignment and the future of safety
- 00:05:32 npm supply chain burn
- 00:07:36 LangChain's return to open source
- 00:08:50 Clooney, Hanks, and the Human Consent Standard

