

Neurasec Newsletter – 22nd February 2026

AI last week (16–22 Feb 2026): big infrastructure bets, major model upgrades, and a reminder that governance is catching up 🧐 🤖

Another huge week in AI — and what stood out most to me was how the conversation is now moving on **three tracks at once**:

1) bigger models + better tools 2) massive infrastructure / data centre investment 3) much sharper regulatory scrutiny

Here's my quick roundup of what happened this week:

- **India's AI Impact Summit became a major focal point** — Reuters reported global AI leaders and world leaders gathering in New Delhi, with India positioning itself to attract more AI investment. Reuters also noted the summit expected very large attendance and framed it as a major AI moment for the developing world.
- **The India summit quickly turned into a “capital commitments” week** — Reuters listed major AI infrastructure commitments announced around the summit, including:
- **OpenAI + Tata / TCS infrastructure news was a standout signal** — Reuters reported OpenAI will become the **first customer** of Tata Consultancy Services' data centre business, starting with **100MW** of capacity, linked to the Stargate AI infrastructure initiative. That feels like another sign the AI race is now as much about power + compute delivery as it is about model quality.
- **Anthropic launched Claude Sonnet 4.6 (Feb 17)** — Anthropic says Sonnet 4.6 is now the default for Free and Pro users, with improvements across coding, computer use, long-context reasoning and agent planning. Anthropic also says it includes a **1M token context window in beta** and broader developer platform upgrades (including code execution, memory and tool-related features now generally available).
- **Google announced Gemini 3.1 Pro (Feb 19)** — Google says 3.1 Pro is rolling out across developer, enterprise and consumer channels (including Gemini API / Vertex AI / Gemini app / NotebookLM), and highlighted benchmark improvements, including a reported **77.1% verified score on ARC-AGI-2**.
- **Enterprise services players are still adapting in real time** — Reuters reported Infosys unveiled a partnership with Anthropic just days after investor concerns about AI disruption to traditional Indian IT services models hit valuations. That's a strong reminder that AI isn't just a product story — it's a services and business model story too.
- **Regulation / safety pressure also moved this week** — Reuters reported Ireland's Data Protection Commission opened a formal investigation into X's Grok over personal data processing and the potential to produce harmful sexualised images/video (including of children). Reuters notes the DPC is X's lead EU regulator and can issue GDPR fines of up to 4% of global revenue.
- **The AI capital race kept accelerating** — Reuters reported Nvidia was close to finalising a **\$30B investment** in OpenAI as part of a funding round where OpenAI was seeking **more than**

\$100B, potentially valuing the company at around **\$830B**. (Still reported as source-based at the time.)

- **And the hardware angle is getting louder again** — Reuters (citing The Information) reported OpenAI has **200+ people** working on a family of AI devices, including a likely first product in the form of a smart speaker, with other device concepts (like smart glasses) also reportedly in development. Important one to watch, but still reported / unconfirmed product plans.

My takeaway this week:

AI is no longer just “who shipped the smartest model?”

It’s increasingly about:

- **who can fund the infrastructure**
- **who can operationalise AI in enterprise**
- **who can navigate regulation safely**
- **and who can turn models into real products people use every day**

If last week felt like model competition, this week felt like **AI industrialisation**.

#AI #ArtificialIntelligence #GenAI #OpenAI #Anthropic #GoogleAI #DataCentres
#DigitalTransformation #LinkedInPost