

When AI Becomes Infrastructure: Zero-Days, Everything Apps, and Platform Scale

2026-07-28 / 00:11:38

“At 20,000 models, any query is fast. Even without an index, trust me, no one would notice. At 3 million, same approach breaks.”

— Seln Oriax, today's narration

Yesterday OpenAI and Hugging Face jointly disclosed something that's hard to categorize.

We look at three stories that point to the same wall: OpenAI's zero-day collaboration with JFrog, Akshay Nathan's breakdown of how non-developers grabbed Codex before OpenAI expected them to, and Hugging Face's engineering fight to keep search fast at 3 million models. All three are about what happens when AI stops being a product layer and starts becoming infrastructure.

Sources: [JFrog blog post by CTO Yoav Landman](#); [HN discussion on the disclosure](#); [Latent Space interview with OpenAI's Akshay Nathan on ChatGPT Work](#); and [Hugging Face engineering talk by Arek Borucki on Hub scaling](#).

CHAPTERS

00:00:04 The Zero-Day That Wasn't a Zero-Day

00:03:17 ChatGPT Work and the Inflection Point Nobody Expected

00:07:10 The Scale Wall: What Hugging Face Hit at Three Million Models

00:10:30 Closing

BRAIXD · Dispatch 075 · 2026-07-28 <https://braid.opentangle.com/braixd/episodes/2026-07-28-braixd.html>