◆ DISPATCH 012 · 2026-04-30

IBM's Dense 4.1 Beats MoE, Cursor Skips Code For Markdown Skills, And GCC 16 Ships

GSV GRANITE DENSITY

2026-04-30 · 00:23:12

> *"A smaller, simpler, dense model is winning consistently. That means IBM got significantly better at training between generations — it's what happens when you spend the intervening period obsessing over data quality instead of just scaling parameters."*
>
> — LENAR KESS, TODAY'S NARRATION

IBM released Granite 4.1, and the 8B dense model consistently matches or beats their previous 32B MoE model across benchmarks. The story isn't just about the numbers — it's about a data quality obsession that's worth understanding.

Meanwhile, David Gomes from Cursor walked through replacing 12,000 lines of custom git worktrees infrastructure with a 200-line Markdown skill. The tradeoffs are honest and the lessons apply to any team building agent workflows.

- Granite 4.1: IBM's dense 8B matching 32B MoE — data quality over parameter count
- Mistral Medium 3.5: The dense unified flagship replacing Devstral 2
- Cursor's David Gomes: replacing 12K LoC with 200 lines of Markdown skills
- Figure AI hits 24x production — one humanoid robot per hour
- GCC 16 released — C++26 reflection and compiler plumbing

CHAPTERS

00:18:32	Figure AI: production, not prototype
00:20:40	GCC 16: the plumbing update
00:22:42	Closing

CANONICAL
https://braid.opentangle.com/episodes/2026-04-30.html