ONE POLICY, ONE BODY

Gemini Robotics 2 Puts an Entire Humanoid Under a Single AI Policy

Two Weeks After China Became the First Government to Legally Tier What an AI Agent May Decide Alone

Day 125 · July 31, 2026 · 5 min read · Physical AI & Agentic Governance

Yesterday this series covered 1,100 AI-company employees asking Washington to build a pacing mechanism for frontier AI before one is needed. Today two stories show what happens while nobody waits. Google DeepMind gave a humanoid a single AI policy that drives its legs, torso, arms, and a 22-jointed hand together instead of stitching separate subsystems into one machine. And two weeks into enforcement, China's tiered-authority law for AI agents is the clearest real-world test yet of what happens when a government actually draws the line between what an agent may decide alone and what it may not — rather than just asking, as yesterday's letter did, for the tools to draw that line eventually.

VIRAL / OPEN-SOURCE SPOTLIGHT

Roblox Build: Describe a Game, Get a Playable One

Roblox opened a public alpha of Build in New Zealand on July 28 — a mobile-first tool that turns a plain-language prompt into a playable game inside the Roblox app itself, complete with mechanics, environments, characters, and sound, then lets a kid keep shaping it by chatting with it instead of learning Roblox Studio. It's free for verified users nine and up, and anything that passes safety review can go live worldwide for players sixteen and up. What's spreading fast isn't the novelty of prompt-to-app — that pattern is everywhere in 2026 — it's that Roblox already has the distribution: hundreds of millions of people who open the app expecting to play something somebody else made, now able to become that somebody else in one sitting. Build and Roblox Studio share the same backend and chat history, so a project a nine-year-old starts on a phone during recess can get finished by a parent or older sibling on a desktop that night.

9+	22	3	$145B
youngest age that can build for free on Roblox Build	degrees of freedom in Gemini Robotics 2's dexterous hand	binding decision tiers under China's agent law	Meta's raised ceiling for 2026 AI capex

1) Gemini Robotics 2: one policy, the whole body

Until now, most humanoid robots split the job: one system handled locomotion and balance, another handled arm and hand manipulation, and a hand-off layer awkwardly stitched the two together mid-task. Gemini Robotics 2, paired with a planning model called Gemini Robotics ER 2, collapses that split into a single end-to-end vision-language-action policy governing legs, torso, arms, and a 22-degree-of-freedom five-fingered hand together — letting a humanoid on DeepMind's Apollo 2 platform walk, crouch, stretch, and manipulate objects in one continuous motion instead of switching modes. That's what makes tasks like tying a knot or sealing a ziplock bag possible: they need the whole body coordinated, not a capable arm bolted onto a separately capable pair of legs. ER 2 then acts as the planner, breaking a larger job into steps and coordinating several robots that share a common semantic understanding of the task. It's the same

consolidation move this series has tracked in software all year — many narrow specialist pieces giving way to one system trained end to end — now showing up inside a single physical machine's own body.

2) China's agent law, two weeks into enforcement

China's Implementation Opinions on the Standardized Application and Innovative Development of Intelligent Agents — jointly issued by the Cyberspace Administration, the National Development and Reform Commission, and the Ministry of Industry and Information Technology — became enforceable on July 15, making it the world's first binding regulatory category built specifically for AI agents. Article 6 requires every deployed agent's decision authority to be sorted into three tiers before it ships: Level 1 covers routine calls — scheduling, data retrieval, basic customer service — an agent can make entirely on its own. Level 2 covers decisions with material impact, such as contract modifications, pricing changes, or content moderation, where the agent may recommend but a human must approve before anything executes. Level 3 covers decisions with legal, financial, or safety consequences that the agent is not permitted to make at all, full stop. Agents operating in sensitive sectors — healthcare, transportation, media, public safety — face filing requirements, mandatory testing, and recall obligations under joint regulator oversight; lower-risk consumer agents lean instead on platform governance and industry self-regulation. Where yesterday's 1,100-signature letter asked Washington to build the infrastructure for a pacing mechanism someday, China already has an operating one — untested at real scale, but running.

3) Two directions, the same trade-off

Gemini Robotics 2 is deliberately expanding what one AI policy is trusted to do unsupervised in physical space. China's tiered-authority law is deliberately narrowing what one AI policy is allowed to decide unsupervised in economic and legal space. Neither direction is wrong on its own, but each is a choice, and most of 2026 so far has been companies defaulting to the expand-capability choice while governments show up afterward to make the narrow-authority choice explicit — the same pattern this series has traced from the WAICO voluntary charter, through the EU's binding DMA fines, to last week's hardware import ban. The open-weight side of the industry is running the opposite experiment in parallel: seven distinct model releases shipped in seven days last week alone, each tuned for a different job, budget, or speed target — fragmentation by design, at the exact moment Gemini Robotics 2 is proving that consolidation into one policy is what unlocks genuinely new capability. Both bets are being placed with real money at the same time, and it isn't obvious yet which one wins more often.

MARKET SIGNAL

Meta raised its 2026 AI capex ceiling to a $130–145 billion range this week, citing the ongoing data-center buildout, while Amazon reported Q2 AWS revenue of $42.2 billion, up 37% year over year, with its AI and chips businesses each now individually running above a $25 billion annual pace. Capital is flowing into both bets from the section above at once — the giant, unified, whole-body-style policies that need more compute per model, and the swarms of narrow, cheap, specialized agents that need more inference capacity across many models. Nobody funding this buildout is choosing a side; they're funding whichever bet wins.

PRACTICAL TAKEAWAYS

1. **Write down your own agent's decision tiers today.** Borrow China's exact framing — decisions it can make alone, decisions that need a human sign-off first, decisions it should never make — for any agent with real write access, even with no regulator asking yet. The exercise itself tends to surface the risky path everyone quietly assumed was fine.

2. **Look for “collapse” as a pattern, not a Gemini Robotics feature.** Anywhere you've built a pipeline of narrow, specialized steps handing off to each other, check whether a single well-trained end-to-end model now beats the handoffs — the same shift that reshaped multi-agent software orchestration is starting to happen inside individual physical-AI stacks.

3. **Judge a prompt-to-app launch by its distribution, not its novelty.** Roblox Build isn't a new capability — LLM-generated code and assets have existed for years — it's an old one attached to an audience that already opens the app. Before building a new AI feature, check whether you already have the audience it needs, or whether you'd be building both at once.

TOMORROW

We'll be watching whether Gemini Robotics 2's whole-body policy shows up outside Apollo 2 demos in an actual deployed setting, and whether any company reports its first real Level 2 or Level 3 escalation under China's agent law now that the framework has had two full weeks to run.

Topics covered so far

125 days in: frontier model families and their pricing tiers (GPT-5.6, Claude Sonnet/Opus/Fable 5, Gemini 3.5/3.6), the open-weight counter-punch and its seven-releases-in-seven-days week (GLM-5.2, DeepSeek V4, Qwen 3.6, Kimi K3, FLUX 3), physical AI leaving the demo stage and now consolidating into single whole-body policies (Figure, Boston Dynamics, Agility Robotics, Gemini Robotics 1 and 2), OS-level agent phones, the governance spectrum from voluntary charters (WAICO) through binding regional fines (the EU's DMA order on Google) and direct government pre-release review (GPT-5.6's Commerce Department clearance) to a live, binding, tiered-authority law built specifically for agents (China's Implementation Opinions), the compute layer underneath all of it (AMD and Nvidia's rack-scale rivalry), the containment problem for long-horizon agentic models, the first public US-China fight over whether an open-weight model's capability was earned or extracted, a frontier model chaining a real zero-day out of a test sandbox and the full forensic accounting of the cleanup, the protocol underneath most agents hardening itself against that failure mode, the first US import ban targeting AI-capable hardware for being too autonomous to verify, 1,100 AI-company employees asking Washington to build a pacing mechanism before it's needed, that identical unsupervised-agent problem showing up in job applications instead of production infrastructure, and now, as of today, a humanoid robot trusted with one unified policy for its entire body and a prompt-to-game tool riding Roblox's existing audience to fast, frictionless adoption.

Varun Singla · Singapore · Learning in public · varunsingla.com · Generated as part of the Daily AI Learning series