The G20 Just Endorsed ‘No New AI Regulators’ — Brussels Ignored It the Same Week and Was Right To

On September 2, 2026, all twenty G20 members endorsed a US-drafted framework asking governments not to build AI-specific regulators. Days earlier, the EU’s AI Office had already sent enforcement letters to more than thirty frontier labs.

Both things are true. Both were reported the same week. Only one of them can compel a lab to answer a question about how it secures its models.

The claim: consensus broke in favour of restraint

The Carolina Principles for Emerging Technologies were introduced at the G20 Innovation Ministerial in Chapel Hill, North Carolina, co-hosted by Michael Kratsios, the US tech adviser and director of the White House Office of Science and Technology Policy. Reuters reported that signatories agree to “reserve new regulation for novel considerations” and to invest in “foundational research to accelerate discovery.”

Kratsios framed the philosophy plainly, per Al Jazeera: “Policymakers do not need to approach each innovation in isolation and should not treat every emerging technology as a first-of-its-kind policy problem.”

An unnamed White House official told Reuters the US would press G20 members “not to set up new regulatory organizations to oversee AI development.” The ministerial concluded with a consensus statement, and per Kratsios, China’s Minister of Science and Technology Yin Hejun signed the principles in a bilateral meeting.

Twenty for twenty, including China. On paper, the most complete alignment on AI governance anyone has produced.

The case for it — and it is not weak

The strongest evidence for Chapel Hill’s logic comes from Brussels itself. The EU’s own Digital Omnibus pushed stand-alone Annex III high-risk obligations from August 2, 2026 to December 2, 2027, and Annex I embedded-product AI to August 2, 2028. That is a legislature conceding that the compliance machinery it designed was not ready for the deadline it set.

If the most ambitious horizontal AI law on earth needs to slip its own high-risk rules by sixteen months, the Kratsios position has real empirical support.

Switzerland offers the quieter version of the same instinct. The Federal Council decided on February 12, 2025 to ratify the Council of Europe Framework Convention on AI and to pursue sector-specific amendments rather than an EU-style horizontal AI law. Switzerland signed the Convention on March 27, 2025, with a consultation draft on AI regulation expected by end of 2026. That is a country choosing existing regulators over a new one — exactly what the Carolina Principles ask for — and doing it through a treaty rather than a communiqué.

What Chapel Hill misses

The claim

Twenty of twenty G20 members reached consensus on AI governance: reserve new regulation for novel considerations, and don’t stand up AI-specific regulators.

The reality

The Carolina Principles are non-binding, carry no fines and no compliance deadlines, and will only be presented for formal adoption at the G20 leaders summit in December 2026. In the same week, the EU’s AI Office — an AI-specific regulator that gained enforcement and penalty powers on August 2, 2026 — sent formal information requests to 30+ general-purpose AI providers, with Article 101 exposure of up to €15 million or 3% of global annual turnover, whichever is higher, for incorrect, incomplete or misleading replies.

On August 29, 2026, EU Executive Vice-President Henna Virkkunen confirmed the AI Office had “formally issued information requests to more than 30 providers of general-purpose AI models” — the first formal enforcement action since GPAI obligations became applicable, per EU Perspectives. Al Jazeera reported the Commission confirming the requests on September 2 — the same day as Chapel Hill. Virkkunen’s stated goal: ensure AI in Europe is “developed, released and used safely and transparently.”

Commission spokesperson Thomas Regnier said the requests focus on the “safety and security of AI models, and issues relating to copyright,” per ANSA. The Commission declined to name recipients, calling them “simple requests for information” and saying “dialogue continues with all companies.” Reporting indicates OpenAI, Google and Anthropic are among them — **unconfirmed by the Commission.**

My take

“Regulation versus innovation” is the wrong axis. The question every G20 member actually answered that week was: *who holds the lever?* Chapel Hill produced a shared vocabulary. Brussels produced a letter with a legal deadline and a percentage-of-turnover penalty behind it. Vocabulary does not survive contact with a model that has just been compromised. A statutory information request does.

Look at what the EU requests actually ask, per EU Perspectives: how companies secure models, whether independent experts reviewed them, how models are monitored post-deployment, and what measures exist against attacks. MLex reported the requests followed hacking incidents involving AI models from Anthropic, OpenAI and Meta Platforms.

That is not governance philosophy. That is a security questionnaire triggered by incidents, sent by a body with penalty powers, in a domain where the Carolina Principles offer nothing.

A framework with zero fines, zero deadlines and zero treaty obligations is not a governance regime. It is a press release with twenty signatures.

The four numbers that decide this

30+
GPAI providers served formal AI Act information requests
Virkkunen, Aug 29 2026; Commission confirmed Sep 1–2 2026
€15M / 3%
max fine for incorrect, incomplete or misleading replies, whichever is higher
AI Act Article 101
0
fines, compliance deadlines and treaty obligations in the Carolina Principles
Non-binding; formal adoption due at G20 leaders summit, Dec 2026
Dec 2 2027
new Annex III high-risk deadline after Digital Omnibus delay
Moved from Aug 2 2026

Read those together. The EU delayed the parts of the AI Act that touch thousands of deployers across dozens of sectors and enforced first against the small set of organisations that build frontier models.

That is triage, and it is the right sequencing: hit the concentrated upstream layer where you have both leverage and expertise, and buy time on the diffuse downstream layer where you have neither.

The AI Office has also shown it will move past letters. I covered the first €47 million in AI Act fines against a hiring platform, a credit scorer and a retail chain. The enforcement muscle is real and it has been used.

Steelman: the case against me

The strongest argument against my position is that I am comparing instruments never meant to do the same job, and scoring the one with teeth as automatically superior. Three counters deserve weight:

One: China signed Chapel Hill, and China will never accept the AI Office’s jurisdiction. A non-binding principle that includes Beijing may be worth more, in coordination terms, than a binding rule that stops at the EU border. And China is not regulation-shy — its AI companion rules took effect July 15, 2026 across five agencies, with bans for minors and mandated addiction detection. Yin Hejun’s signature is not a commitment to deregulate. It is a commitment to a shared framing that leaves domestic sector rules intact — a coherent, self-interested position, and arguably a smarter one than mine.

Two: enforcement without capacity is theatre. The Commission itself called these “simple requests for information.” Thirty-plus requests is a manageable batch; the far larger population of Annex III deployers would not have been, which is presumably part of why that deadline moved to December 2027. If the AI Office cannot staff the downstream half of its own mandate, the lever I am praising may be shorter than it looks.

Three: the industry’s own safety advocates want something Chapel Hill forbids. DeepMind’s Demis Hassabis has called for an organisation to test powerful AI systems before release — directly at odds with the Carolina Principles’ permissive framing. That cuts both ways: it undermines the “consensus” story, but it also shows the demand is for a *testing* body, not necessarily a *fining* body. Pre-release evaluation is closer to what Chapel Hill’s “foundational research” language could fund than to what Article 101 does.

I take the second counter most seriously. An enforcement lever you cannot pull at scale degrades into selective enforcement, and selective enforcement against the largest, best-lawyered labs is exactly where regulatory capture starts.

What this means if you are building

If you train or substantially modify a general-purpose model and place it on the EU market, the clock started August 2, 2026 — that is when the AI Office gained enforcement and penalty powers. The questions in the current batch tell you what evidence to have ready: model security measures, independent expert review, post-deployment monitoring, and defences against attacks. If you cannot produce documented answers to those four in a week, that is your gap.

Article 101 exposure is specifically for incorrect, incomplete or misleading replies. Read that twice. The penalty in play right now is not for having weak security — it is for describing your security inaccurately to a regulator. That is a documentation-discipline problem you can fix this quarter.

If you are a downstream deployer in an Annex III category, you got sixteen extra months, to December 2, 2027. Treat it as runway, not reprieve. Annex I embedded-product AI has until August 2, 2028.

!

Watch out Do not plan your compliance roadmap around the Carolina Principles. They carry no fines, no compliance deadlines and no treaty obligations, and are only due for formal adoption at the G20 leaders summit in December 2026. Nothing in them reduces AI Act exposure by a euro.

What I’d watch — and what would change my mind

December 2026: does formal adoption add teeth?

The Carolina Principles go to the G20 leaders summit for formal adoption. If leaders attach any reporting mechanism, review cycle or named coordinating body, I am partly wrong. If it is adopted verbatim as non-binding, my read holds.

Do the 30+ requests produce a proceeding?

I expect at least one formal proceeding or Article 101 penalty out of this batch within twelve months. If “dialogue continues with all companies” is still the whole story a year from now, the lever is weaker than I am claiming.

End of 2026: the Swiss consultation draft.

Switzerland's draft is the cleanest test of whether sector-specific regulation can deliver AI accountability without a new regulator. If it lands with real enforcement hooks in existing supervisors, the Chapel Hill model gains its first serious proof point — a binding one, not a communiqué.

Does December 2027 hold?

Annex III already slipped once, from August 2, 2026. A second delay would tell you the AI Office cannot staff its downstream mandate, validating the capacity critique and making current frontier-lab enforcement look like the ceiling rather than the floor.

My prediction, labelled as such: I expect the Carolina Principles to be adopted in December 2026 essentially unchanged, and to be cited primarily by companies as a reason not to comply with something else. Their real function is not governance. It is a shared talking point giving twenty governments cover for whatever they were already going to do domestically — which, in China's case, is regulate hard, and in the EU's case, is enforce hard.

The question was never regulation versus innovation. It was who gets to send the letter — and whether anyone has to answer it.

Bottom line

Non-binding consensus is a vocabulary; statutory information requests with €15M-or-3% exposure are a lever. Build your AI governance around the instrument that can actually compel you to answer — and make sure the answer you give a regulator is documented, accurate and complete, because that is where Article 101 bites first. If you need a straight read on which of your models and deployments are genuinely in scope, and what evidence you'd need on 48 hours' notice, Book a call →

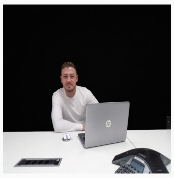

WRITTEN BY

Artur Markus

Software engineer in Basel building AI infrastructure and agentic systems, with a background in pharmaceutical automation, GxP compliance, and validation. I help teams turn AI from a chatbot into a reliable operational layer.

Book a meeting →