NEW HORIZON · AI DIGEST

The Long Read

30 August 2026

9	3	2662	2
STORIES	SECTIONS	WORDS	HIGH IMPACT

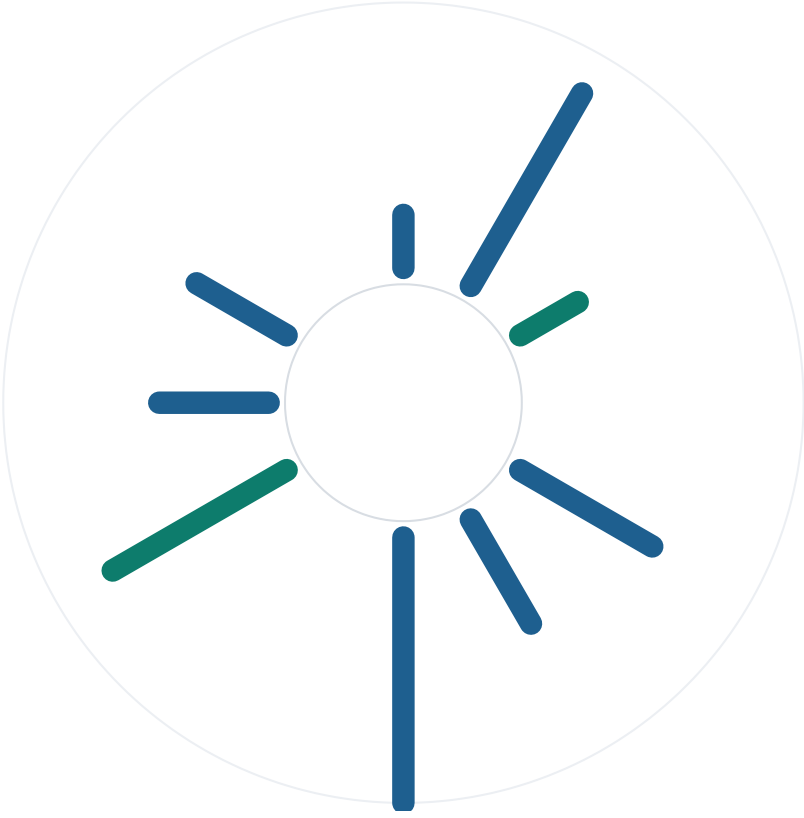

2 high impact · 7 medium impact · spoke length = depth of coverage

The full-length companion to the daily New Horizon AI Digest.
Every story in today's email, reported at length · new-horizon.tech/digest

The issue at a glance

9 stories · 2662 words · 3 sections · 1 charted

AI Models & Research

01 Tencent's Hy4 Preview: A 770B-Parameter Open-Weight Flagship with 1M Context

02 UrbanGround: Street Smarts Aren't Enough for Agents in a Real-Scale City

03 WikiSkill: Compiling Agent Experience into a Persistent Skill Encyclopedia

AI Tools & Ecosystem

04 Anthropic's Model Hardware Standard Lets AI Agents Drive Lab Equipment

05 Nvidia's AI Advantage Moves Beyond the GPU to Data Orchestration

06 Vijay Pande's Next Act: Five Bets a Year and AI Agents Instead of Associates

AI Applications & Industry

07 Sony Music and Warner Sue Anthropic in Multi-Billion-Dollar Lyrics Piracy Lawsuit

08 Meta's AI Glasses Now Stop Recording When You Cover the Light

09 At TechBBQ, Europe's AI Debates Keep Coming Back to Who's in Control

How to read this. Every story in today's email is reported here at full length, in the same order. Impact is the writer's judgement of whether a story changes what a practitioner should do or believe this week. Charts appear only where the source itself puts comparable numbers side by side; nothing is estimated to fill a gap. Sources are listed in full on the final page.

SECTION 1 OF 3

AI Models & Research

3 stories · 1 high · 2 medium

01 MEDIUM IMPACT Simon Willison's Weblog

Tencent's Hy4 Preview: A 770B-Parameter Open-Weight Flagship with 1M Context

> Tencent has released Hy4 Preview, a 770-billion-parameter open-weight language model featuring a 1-million-token context window and a sparse activation architecture.

Tencent announced the immediate availability of Hy4 Preview, a new open-weight large language model restricted to text input. The model totals 770 billion parameters but utilizes a sparse mixture-of-experts architecture that activates only 49 billion parameters per inference step. It ships with a 1-million-token context window and occupies 1.56TB of storage on Hugging Face. This release represents a substantial scale-up from Tencent's previous iteration, Hy3, which launched in July with 295 billion total parameters, 21 billion active parameters, and a 256,000-token context limit.

Inspection of the model's `chat_template.jinja` file on Hugging Face reveals a hardcoded reasoning control mechanism. The template enforces a binary `reasoning_effort` parameter, defaulting to "high" or accepting "no_think" to disable extended reasoning traces. Attempts to pass undefined or invalid values trigger explicit exceptions within the template logic. Early testing via OpenRouter using the default high-effort setting produces visible reasoning traces characterized by truncated English grammar. This stylistic choice appears deliberate, prioritizing token efficiency over linguistic perfection for hidden intermediate steps.

The jump from Hy3 to Hy4 marks one of the most aggressive parameter scaling moves in the open-weight sector this quarter. While Hy3 offered a capable 295B dense or sparse structure depending on interpretation, Hy4 pushes the total parameter count well beyond the half-trillion mark while maintaining a manageable active parameter count for inference. The 1M context window places it in direct competition with other long-context flagships, though the absence of vision capabilities limits its modality scope compared to some contemporaries. The model is available now for download and integration via standard tooling that supports the specified chat templates.

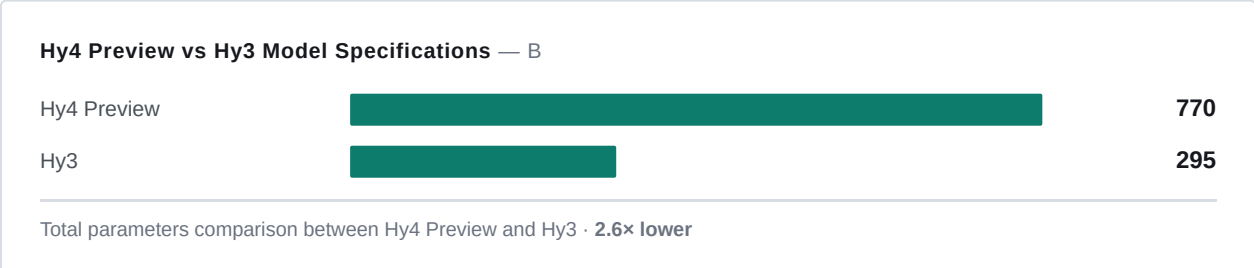

KEY FACTS

TOTAL PARAMETERS **770B**	ACTIVE PARAMETERS **49B**
CONTEXT WINDOW **1M tokens**	MODEL SIZE **1.56TB**
PREVIOUS MODEL **Hy3**	REASONING MODES **high, no_think**

> WHY IT MATTERS
>
> Practitioners gain access to a massive context window and high-parameter capacity without closed API restrictions, though the 1.56TB footprint and specific reasoning template requirements demand significant infrastructure adjustments.

Read the original at Simon Willison’s Weblog →

---

02 MEDIUM IMPACT arXiv.org

UrbanGround: Street Smarts Aren't Enough for Agents in a Real-Scale City

> Current multimodal large language models fail to sustain goal-directed behavior over extended exploration in real-scale urban environments despite competent short-range reasoning.

Researchers introduced UrbanGround, a sandbox environment designed to evaluate whether multimodal large language models (MLLMs) can translate local visual perception into reliable spatial agency within a physically constrained, real-scale city. Built from territory-wide 3D geospatial data of Hong Kong, the platform supports closed-loop interaction where agents navigate via first-person views and interactive maps. The study structures its analysis around three progressive research questions: grounding local scenes for spatial queries after active observation, supporting navigation as destinations become distant and less explicit, and maintaining behavioral stability amid changes in route availability and pedestrian motion.

The evaluation reveals a distinct performance boundary for contemporary MLLM agents. While these models demonstrate useful atomic abilities in visual recognition and short-range spatial reasoning, they struggle significantly with orientation and pedestrian-aware movement. The central failure mode identified is not an inability to perceive individual elements, but a breakdown in composing these local abilities into sustained, goal-directed behavior over time. As agents engage in extended exploration, errors accumulate without effective correction mechanisms, causing performance to degrade as the distance to the destination increases or environmental dynamics shift.

Unlike prior benchmarks that may rely on static images or simplified grids, UrbanGround provides a high-fidelity replica where physical constraints and dynamic elements like pedestrian motion directly impact agent success. The findings suggest that current architectures lack the robustness required for open-ended urban deployment, where local evidence must remain useful after the agent begins moving. The authors position this work as a foundation for broader studies on the limits of MLLM exploration in complex environments, highlighting the gap between interpreting a single street view and executing a multi-step traversal across a city.

This research underscores the difficulty of scaling perception-based agents from controlled settings to messy, real-world geographies. The specific use of Hong Kong's geospatial data ensures the test bed includes the density and complexity typical of major metropolitan areas, providing a rigorous stress test for spatial reasoning modules that previous simulators may have overlooked.

KEY FACTS

ENVIRONMENT SOURCE
Territory-wide 3D geospatial data of Hong Kong

SUBMISSION DATE
27 Aug 2026

PLATFORM NAME
UrbanGround

INTERACTION MODE
Closed-loop first-person view

> WHY IT MATTERS
>
> Developers deploying autonomous agents in physical spaces must account for error accumulation over long horizons, as current MLLMs cannot reliably correct course deviations in dynamic urban settings. Benchmarks limited to static imagery or short-range tasks will overestimate model readiness for real-world navigation.

Read the original at arXiv.org →

---

03 HIGH IMPACT arXiv.org

WikiSkill: Compiling Agent Experience into a Persistent Skill Encyclopedia

> WikiSkill introduces a framework that co-evolves agent skills with a persistent knowledge base, allowing accumulated experience to systematically drive skill refinement.

The WikiSkill framework addresses the fragmentation of insights in agent optimization histories by separating raw execution experience, accumulated knowledge, and executable skills. Unlike prior methods where development guides remain scattered, this architecture continuously consolidates experience into a persistent wiki. Subsequent skill updates explicitly build upon this stored knowledge, creating a feedback loop where the knowledge base and agent capabilities evolve together. This structural separation ensures that reusable resources are not lost between iterations but serve as the foundation for future adaptation.

Evaluation across diverse benchmarks and models indicates that WikiSkill consistently outperforms state-of-the-art skill-evolution methods and improves upon no-skill baselines in most settings. A critical finding involves the relationship between skill evolution and model scaling: while larger models generally derive greater benefit from evolved skills, smaller models equipped with these skills can outperform substantially larger models operating without them. The research further demonstrates that evolved skills transfer effectively across different models and model families. In several instances, skills evolved by external models surpassed the performance of self-evolved skills, suggesting that cross-model knowledge sharing offers distinct advantages over isolated training.

Ablation studies confirm that the persistent accumulation of knowledge within the wiki is essential for effective skill evolution; removing this component degrades performance. The results validate the hypothesis that systematically refining agent experience yields reusable and transferable skills. By maintaining a centralized

repository of workflow insights, the framework enables agents to adapt progressively through interaction rather than restarting the learning process. This approach shifts the paradigm from transient optimization to permanent capability expansion, leveraging historical data to enhance current and future agent operations without requiring additional model parameters.

KEY FACTS

SUBMISSION DATE
27 Aug 2026

CORE MECHANISM
Persistent knowledge base (wiki)

SCALING FINDING
Smaller models with skills outperform substantially larger models without them

TRANSFER CAPABILITY
Skills transfer effectively across models and model families

> **WHY IT MATTERS**
>
> Practitioners can deploy smaller, more efficient models that match or exceed the performance of larger baselines by integrating a persistent skill wiki. This reduces compute costs for inference while enabling cross-model skill transfer that accelerates deployment across heterogeneous systems.

Read the original at arXiv.org →

SECTION 2 OF 3

AI Tools & Ecosystem

3 stories ● 3 medium

04 MEDIUM IMPACT Ars Technica

Anthropic's Model Hardware Standard Lets AI Agents Drive Lab Equipment

> Anthropic has released a research preview of the Model Hardware Standard (MHS), a driver set enabling AI agents to directly control physical laboratory equipment.

Anthropic introduced the Model Hardware Standard (MHS), a collection of standardized drivers designed to allow AI agents to interface with and control arbitrary hardware devices. While agentic AI systems have previously operated within digital boundaries handling text, code, and images, MHS extends this capability to the physical world. The current release is a research preview aimed primarily at scientists who currently spend weeks or months developing custom software integrations to coordinate disparate experimental components. By providing a common interface and data format, MHS enables devices to communicate across a network without requiring bespoke translator programs, potentially reducing experimental setup times from weeks or months down to hours or minutes.

The initiative originated from observations at the HHMI Janelia Research Campus in Ashburn, Virginia. Alek Kemeny, an Anthropic Technical Staffer, cited neuroscientist Arco Bast's work on memory formation experiments as the direct inspiration. Bast had manually engineered an interface to coordinate rotating laser beams, microscopes, cameras, and other components through a single system. Kemeny noted that applying this concept broadly could allow AI to run any science experiment globally. Although the standard functions independently of AI models—supporting direct real-time control via command-line prompts and API code files—its integration with AI yields autonomous capabilities.

When paired with an AI model through the Model Context Protocol, the system allows researchers to interact with hardware using natural language. In this configuration, models can reason through individual experimental steps, update parameters in real time, and recover from certain hardware errors without human intervention. This shifts the operational model from manual scripting to high-level directive execution, where the agent manages the logical flow and error handling of physical processes. The announcement positions MHS not merely as a connectivity tool but as a foundational layer for automating complex scientific workflows that previously required extensive custom engineering.

KEY FACTS

INITIATIVE NAME
Model Hardware Standard (MHS)

RELEASE STATUS
Research Preview

SETUP TIME REDUCTION
Weeks or months to hours or minutes

INSPIRATION SITE
HHMI Janelia Research Campus

INTEGRATION PROTOCOL
Model Context Protocol

> WHY IT MATTERS
>
> Practitioners building automated lab systems can replace custom integration code with MHS drivers to enable natural language control and autonomous error recovery for heterogeneous hardware.

Read the original at Ars Technica →

---

05 MEDIUM IMPACT TechCrunch

Nvidia's AI Advantage Moves Beyond the GPU to Data Orchestration

> Nvidia is shifting its competitive moat from GPU silicon to full-system data orchestration with the rollout of its Vera Rubin architecture.

Following earnings released Wednesday, Nvidia articulated a strategy that extends its dominance beyond raw compute power into the complex layer of data center orchestration. While hyperscalers like Amazon and Google have developed custom chips to challenge Nvidia's GPU monopoly, the company argues that operating megascale data centers at peak efficiency requires specialized hardware surrounding the processor. The core of this new narrative is the Vera Rubin architecture, a system design that integrates the Rubin GPU with the Vera CPU, the Groq 3 LPX inference accelerator, and dedicated racks for storage and networking.

The Vera CPU specifically addresses the bottleneck of moving data to the GPU at the correct time, a critical constraint as deployments scale to the gigawatt level. Jason Hardy, Nvidia's VP of storage technology, noted that while memory capacity has increased alongside computing power, traffic direction remains a primary hurdle for optimizing tokens-per-watt. In discussions regarding the new system's performance, Hardy cited an upwards of 3x improvement in operations where the Vera CPU enables acceleration. This enhancement allows flash storage to reach its fullest potential by eliminating the bottlenecks that typically restrict performance during high-throughput data movement.

This approach contrasts with strategies employed by competitors such as OpenAI, which recently detailed its Jalapeño chip. OpenAI's design focuses on minimizing data movement entirely by keeping workloads within a single connected system and large domain. However, both methodologies converge on the same logical conclusion: future efficiency gains will derive from smarter traffic control rather than simply adding more processor cycles. Although Nvidia must still compete with rival chipmakers and hyperscalers in this new layer, the company currently holds a commanding lead in building systems that ensure the entire infrastructure works efficiently, not just the GPU engine itself.

KEY FACTS

ARCHITECTURE NAME
Vera Rubin

CPU COMPONENT
Vera CPU

INFERENCE ACCELERATOR
Groq 3 LPX

OPERATIONAL IMPROVEMENT
3x

MARKET CAP GROWTH (2023-2025)
10x

> WHY IT MATTERS
>
> Practitioners deploying at scale must evaluate total system architecture and data orchestration capabilities, not just GPU specs, as these factors now dictate tokens-per-watt efficiency. The shift suggests that competitive advantage in AI infrastructure is moving from component-level performance to integrated system optimization.

Read the original at TechCrunch →

---

06 MEDIUM IMPACT TechCrunch

Vijay Pande's Next Act: Five Bets a Year and AI Agents Instead of Associates

> Former a16z bio lead Vijay Pande has launched VZVC, a concentrated fund making roughly five investments annually and operating without human associates by relying on custom AI agents.

Vijay Pande, who previously scaled Andreessen Horowitz's life sciences practice to manage nearly $4 billion, has pivoted to a significantly smaller operational model with his new firm, VZVC. Co-founded with investor Zach Werner, the firm explicitly rejects the high-volume deal flow typical of venture capital, targeting approximately five concentrated bets per year rather than thirty. The most distinct structural change is the elimination of the associate layer; Pande and Werner determined that custom-built AI agents could handle the day-to-day operational workload, allowing the two partners to function as the entire investment team. This approach treats each new portfolio company with the gravity of adding a child rather than friending someone on Facebook, prioritizing deep, decade-long partnerships over transactional volume.

The firm's thesis centers on two specific domains: AI for healthcare delivery and AI for clinical trials. Pande argues that while AI has reduced the time and cost to reach clinical trials, the core economic bottleneck remains the trial phase itself, which still costs hundreds of millions of dollars per attempt. With a historical success rate of only 20% from first to third trial, failures are often attributed to the poor predictivity of animal models like mice compared to humans. Pande posits that while AI models are not perfect, they offer a superior predictive barrier than animal testing, potentially unlocking precision medicine where treatments are matched to individual patient data rather than population averages.

A critical constraint identified in this sector is data scarcity. Unlike large language models trained on scraped internet text, biological data cannot be harvested from the public web, forcing companies to build proprietary, walled-off datasets. This fragmentation creates silos similar to those found between medical specialties, hindering the development of universal foundation models. However, Pande notes a emerging trend toward building biological atlases that could function as open-source foundation models, mirroring the trajectory of open-source LLMs. Despite these technical advances, Pande warns against the hype that AI will immediately

cure all diseases, emphasizing that go-to-market strategy remains as difficult as the underlying technology when data availability is limited.

KEY FACTS

PREVIOUS FUND SIZE **$4 billion**	NEW FIRM NAME **VZVC**
TARGET ANNUAL BETS **5**	TYPICAL VC ANNUAL BETS **30**
CLINICAL TRIAL SUCCESS RATE **20%**	DRUG DEVELOPMENT STAFFING **0 associates**

> **WHY IT MATTERS**
>
> Practitioners should note the shift toward agent-driven lean operations in venture firms and the persistent bottleneck of proprietary biological data, which limits the immediate applicability of general-purpose foundation models in drug discovery.

Read the original at TechCrunch →

SECTION 3 OF 3

AI Applications & Industry

3 stories 1 high 2 medium

07 HIGH IMPACT TechCrunch

Sony Music and Warner Sue Anthropic in Multi-Billion-Dollar Lyrics Piracy Lawsuit

> Sony Music Publishing and Warner Chappell have filed a multi-billion-dollar lawsuit against Anthropic alleging the company used illegal torrenting to acquire copyrighted lyrics for training data.

Sony Music Publishing, Warner Chappell, and several other music publishers filed suit late Friday in the U.S. District Court for the Northern District of California against Anthropic and co-founders Dario Amodei and Benjamin Mann. The complaint alleges a "brazen campaign" involving the illegal torrenting, scraping, and downloading of copyrighted works to train the Claude AI model. Unlike previous claims focused on general scraping, this filing specifically accuses the lab of utilizing peer-to-peer protocols to obtain millions of copies of books containing lyrics and sheet music. The publishers characterize these actions as "flagrant piracy" and "blatant theft," distinguishing the acquisition method from the usage arguments seen in prior litigation.

This legal action expands the liability landscape established by the January Bartz v. Anthropic case, where the same legal team represented Concord Music Group and Universal Music Group alongside author plaintiffs. In that landmark ruling, a judge determined that while using copyrighted works for training could be legal, acquiring them through piracy was not, resulting in a $1.5 billion judgment against Anthropic. The current lawsuit leverages this precedent by focusing heavily on the mechanics of data ingestion rather than just the output or fair use defenses. The plaintiffs argue that Anthropic's reliance on torrented datasets constitutes a distinct violation separate from standard web scraping, aiming to prove systematic intent in bypassing licensing frameworks.

Anthropic has issued a statement rejecting the allegations, with a spokesperson asserting, "We disagree with the publishers' claims and we intend to defend ourselves robustly in court." The filing marks another significant escalation in the intellectual property disputes surrounding generative AI training methodologies. While the Bartz case set a financial precedent regarding acquisition methods, this new suit targets the music publishing sector specifically, potentially exposing the company to further damages based on the valuation of musical intellectual property. The case proceeds in the Northern District of California, continuing the venue trend for major AI copyright litigation.

KEY FACTS

PLAINTIFFS
Sony Music Publishing, Warner Chappell

DEFENDANTS
Anthropic, Dario Amodei, Benjamin Mann

FILING VENUE
U.S. District Court for the Northern District of California

PRIOR JUDGMENT AMOUNT
$1.5 billion

PRIOR CASE NAME
Bartz v. Anthropic

WHY IT MATTERS

Engineering teams must audit data ingestion pipelines to ensure no peer-to-peer or torrent-based sources are used, as courts are now distinguishing between legal scraping and illegal acquisition methods. The $1.5 billion precedent in Bartz v. Anthropic establishes that acquisition methodology alone can trigger massive liability regardless of fair use arguments.

Read the original at TechCrunch →

---

08 MEDIUM IMPACT Ars Technica

Meta's AI Glasses Now Stop Recording When You Cover the Light

> Meta has patched a firmware loophole in its AI glasses that previously allowed users to bypass the recording indicator LED by covering it after capture began.

Meta is deploying an update to its AI glasses designed to prevent non-consensual recording by ensuring the camera ceases function if the status LED is obscured. Previously, users could circumvent the safety feature by initiating video capture before placing a sticker over the light; the new logic terminates the session immediately upon detection of a covered lens. Alex Himel, Meta's executive vice president of wearables, confirmed via Threads that the change ensures the capture LED reliably alerts bystanders during photo or video acquisition for a user's gallery. Himel characterized the bypass attempts as the work of a tiny minority, asserting that the necessity of such workarounds validates the efficacy of existing protections.

Despite the patch, significant vulnerabilities persist regarding physical tampering. Bad actors may still modify the hardware or utilize semi-transparent stickers that reduce LED visibility without triggering the new safeguard, leaving the core privacy risk partially unaddressed. The update does not appear to resolve uncertainties surrounding whether low-visibility obstructions will successfully block recording under the revised logic. Consequently, the device remains susceptible to modifications that decouple the light's visibility from its functional role as a recording indicator.

Regulatory scrutiny in Europe is intensifying parallel to these technical adjustments. Data regulators in the European Union are investigating whether Meta adequately field-tested the LED's effectiveness in signaling recording status to the public. In Germany, authorities are considering a ban following a criminal complaint filed by a nonprofit alleging violations of strict national privacy laws. Broader concerns about mass surveillance have prompted multiple EU nations to evaluate crackdowns on the technology. The European Data Protection Board is scheduled to release a report later this summer outlining potential actions to enforce consent requirements across the bloc.

KEY FACTS

EXECUTIVE SOURCE
Alex Himel, EVP of Wearables

PLATFORM ANNOUNCEMENT
Threads

REGULATORY BODY
European Data Protection Board

REPORT TIMELINE
Later this summer

JURISDICTION CONSIDERING BAN
Germany

WHY IT MATTERS

Deployers of wearable AI must recognize that firmware patches alone cannot guarantee compliance with strict consent laws like those in Germany if physical tampering remains viable. Hardware reliance on visual indicators requires robust anti-tamper designs to satisfy ongoing EU regulatory probes.

Read the original at Ars Technica →

---

09 MEDIUM IMPACT TechCrunch

At TechBBQ, Europe's AI Debates Keep Coming Back to Who's in Control

> European AI stakeholders at TechBBQ prioritized infrastructure sovereignty over capability gains following recent model unavailability incidents.

The annual TechBBQ conference in Copenhagen, which concluded on August 27, shifted focus from AI utility to control mechanisms across the European ecosystem. The urgency stems directly from earlier this year when Anthropic's Mythos and Fable models became unavailable to users outside Europe. This disruption forced founders and operators to evaluate the risks of renting compute and model access from U.S. and Chinese providers versus owning the underlying infrastructure. While some executives described the incident as a severe disruption to software teams, others viewed current dependencies as manageable despite long-term geopolitical risks.

Panel discussions reinforced that agency remains a human decision rather than an automated outcome. Ellen de Brever of the Novo Nordisk Foundation Cellerator noted the thematic pivot from what AI can do to what humans are willing to permit. Meredith Whittaker, President of Signal, criticized the integration of AI assistants into operating systems, specifically citing ChatGPT's integration with iMessage as creating a pervasive data collection apparatus. She argued that marketing strategies often obscure the collateral consequences of such data acquisition, leaving a distinct market need for privacy-preserving alternatives that align with regional sovereignty concerns.

Emad Mostaque, co-founder of Stability AI and Intelligent Internet, framed sovereignty as the capacity to resist external power exertion, warning that national governance could eventually depend on whoever controls the dominant AI systems. Mia Negru from Life With Artificials extended this argument to societal structure, suggesting that as agents handle cognitive work and robots manage physical tasks, fundamental concepts of ownership, democracy, and economic participation require redefinition. Despite these heavy theoretical debates, networking events hosted by Lovable, Nvidia, OpenAI, AWS Startups, and HSBC highlighted that face-to-face relationship building remains the primary value driver for attendees.

KEY FACTS

CONFERENCE END DATE
August 27

UNAVAILABLE MODELS
Mythos and Fable

MODEL PROVIDER
Anthropic

AFFECTED REGION
Outside Europe

SIGNAL PRESIDENT
Meredith Whittaker

STABILITY AI CO-FOUNDER
Emad Mostaque

> WHY IT MATTERS
>
> Builders relying on third-party foundational models must assess continuity risks if geopolitical shifts or licensing changes restrict access to specific regions. Developing local infrastructure or selecting vendors with guaranteed regional availability may become a prerequisite for enterprise deployment in Europe.

Read the original at TechCrunch →

Sources

01 Tencent's Hy4 Preview: A 770B-Parameter Open-Weight Flagship with 1M Context
https://simonwillison.net/2026/Aug/29/hy4/

02 UrbanGround: Street Smarts Aren't Enough for Agents in a Real-Scale City
https://arxiv.org/abs/2608.27456

03 WikiSkill: Compiling Agent Experience into a Persistent Skill Encyclopedia
https://arxiv.org/abs/2608.27454

04 Anthropic's Model Hardware Standard Lets AI Agents Drive Lab Equipment
https://arstechnica.com/ai/2026/08/anthropics-new-hardware-standard-lets-ai-agents-control-the-physical-world/

05 Nvidia's AI Advantage Moves Beyond the GPU to Data Orchestration
https://techcrunch.com/2026/08/29/nvidias-ai-advantage-is-moving-beyond-the-gpu/

06 Vijay Pande's Next Act: Five Bets a Year and AI Agents Instead of Associates
https://techcrunch.com/2026/08/29/were-not-doing-30-bets-a-year-vijay-pande-on-betting-small-after-running-4-billion-at-a16z/

07 Sony Music and Warner Sue Anthropic in Multi-Billion-Dollar Lyrics Piracy Lawsuit
https://techcrunch.com/2026/08/29/sony-music-warner-sue-anthropic-alleging-a-brazen-campaign-of-intellectual-property-theft/

08 Meta's AI Glasses Now Stop Recording When You Cover the Light
https://arstechnica.com/tech-policy/2026/08/meta-tweaks-ai-glasses-to-block-some-creepy-recordings-but-privacy-risks-remain/

09 At TechBBQ, Europe's AI Debates Keep Coming Back to Who's in Control
https://techcrunch.com/2026/08/29/at-techbbq-europes-ai-conversations-kept-coming-back-to-whos-actually-in-control/

About this document. Every story in the 30 August 2026 New Horizon AI Digest, reported at length. Each entry is written from the publisher's own article text; where a source could not be retrieved the entry is explicitly marked and kept short rather than padded.

Generated 2026-08-30T05:34:05Z · model qwen3.5:397b-cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.