

The Long Read

27 September 2026

11

STORIES

3


SECTIONS

2724

WORDS

1

HIGH IMPACT





● 1 high impact ● 10 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

11 stories · 2724 words · 3 sections



AI Models & Research

- 01 ● Claude Computes a Nine-Loop Physics Amplitude — One Loop Beyond the Human Record
- 02 ● Rufus-Air Publishes an Open Eight-Stage Post-Training Recipe for GLM-4.5-Air
- 03 ● TrackEverything: Dense 3D Point Tracking Over Minutes Instead of Seconds

AI Tools & Ecosystem

- 04 ● Warp Raises \$85M and Ships an 'AI Head of HR' for Payroll, Compliance and Onboarding
- 05 ● Microsoft Quietly Retires the Copilot+ PC Badge Two Years After Launch

AI Applications & Industry

- 06 ● NYC Council Unveils 10 AI Bills: Kill Switches, 24-Hour Incident Reporting and Whistleblower Bounties
- 07 ● Washington and Beijing Open an AI Hotline — a 'Red Telephone' for Super Intelligence
- 08 ● Devin Crosses \$1B Annualized Revenue, Doubling in Five Months
- 09 ● SalesBleed: Zero-Click Flaws Let Agents Exfiltrate Salesforce CRM Data
- 10 ● Insurers Say AI Coding Tools Added \$942M to Healthcare Bills
- 11 ● Google Puts a Buy Button in Gemini: Flipkart Checkout Comes to AI Mode in India

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.

AI Models & Research

3 stories ● 3 medium

01 MEDIUM IMPACT www.anthropic.com

Claude Computes a Nine-Loop Physics Amplitude — One Loop Beyond the Human Record

Claude, running inside Anthropic's Claude Science harness, computed the nine-loop six-particle amplitude in planar $N=4$ super Yang-Mills — one loop beyond the previous human record — and the result was independently validated by SLAC's Lance Dixon.

The calculation was prompted by a public challenge from physicist and science writer Matt von Hippel, who asked AI companies to tackle frontier scattering-amplitude problems on an academic-scale compute budget. Anthropic physicists Liam Fitzpatrick and Siddharth Mishra-Sharma took up the $N=4$ super Yang-Mills nine-loop target, using Fable 5.1 within Claude Science, a structured harness built on the Claude LLM. After asking Claude which problem it was most likely to solve, they gave it a minimal prompt specifying the six-particle hexagon amplitude at nine loops, then repeatedly told it to keep working. Claude completed the calculation two ways: the direct bootstrap method and the indirect form-factor approach.

The cost was modest. Either approach would have run an end-user roughly one to two thousand dollars, mostly from long-running Claude usage. The bootstrap calculation itself, done in Python with SymPy, consumed about \$100 of that budget, corresponding to 96 CPUs for a week. Notably, Claude executed the full recipe without external scientific oversight beyond instructions to continue, despite the fragility of the setup — a single mistake in the computational recipe collapses the whole construction.

Lance Dixon, professor at SLAC and Stanford and co-author of the prior eight-loop result, validated the nine-loop amplitude, largely by converting it to the related form factor. He had considered the direct amplitude calculation too hard and had expected to reach nine loops only indirectly. Days after Anthropic's result, Song He's group at the Chinese Academy of Sciences reported they had computed the majority of the result — the symbol — using GPT-6 for some constraints but not for the overall framework. Dixon notes Claude used the methods his collaborators developed and presented the solution in their established format.

Von Hippel's assessment is measured: this was not a superintelligent leap but evidence of more low-hanging fruit than experts expected. Claude used known methods with somewhat more compute than people had tried, and better software engineering practices. The result moves the frontier by one loop in a toy-model theory, not a real-world amplitude, and the humans involved will publish the analysis.

KEY FACTS

LOOPS COMPUTED

9

MODEL

Fable 5.1 via Claude Science

BOOTSTRAP COMPUTE COST

~\$100 (96 CPUs for a week)

PRIOR HUMAN RECORD

8 loops

END-USER COST PER APPROACH

\$1,000-2,000

VALIDATION

Lance Dixon, SLAC/Stanford

WHY IT MATTERS

It demonstrates that a frontier physics calculation can be executed end-to-end by an LLM harness on a budget of roughly \$1-2k, with no human scientific oversight beyond 'keep going' — a signal that teams should be checking whether similar one-shot harness runs can squeeze out results in their own domains.

[Read the original at www.anthropic.com](https://www.anthropic.com) →

02 **MEDIUM IMPACT** huggingface.co

Rufus-Air Publishes an Open Eight-Stage Post-Training Recipe for GLM-4.5-Air

Rufus-Air publishes a fully reproducible eight-stage post-training recipe for GLM-4.5-Air-Base, with stage-by-stage data, reward design, and infrastructure documented.

The recipe runs as a serial pipeline of eight stages: SFT, Reasoning RL, Coding RL, Instruction-Following RL, General Agent, Coding Agent, Search Agent, and RLHF. The authors document the data, reward design, infrastructure, stage order, and stagewise results for each. Stages are ordered from basic to advanced capabilities, and the reward signal moves from hard, verifiable rewards in early stages to softer judge-based signals later. Training uses open-source components and public data, much of it used as released, with no new human annotation and no in-house distillation teacher.

The paper's main findings are fourfold. Diverse, high-quality SFT establishes a strong capability floor before any RL. Difficulty filtering keeps RL prompts within a productive learning range. Reward reliability serves as a practical principle for ordering stages. And infrastructure and engineering choices are treated as part of the recipe itself, not an implementation detail. The resulting model, Rufus-Air, improves over the official GLM-4.5-Air post-trained release and is reported competitive with similarly sized open models.

The release is notable for its completeness: rather than reporting only final benchmark scores, it provides the intermediate stagewise results needed to reproduce or adapt the pipeline. For practitioners, that means the ordering rationale and reward-design decisions are inspectable, not just the endpoint. The absence of proprietary distillation teachers and new human annotation also lowers the barrier to replication on different base models.

KEY FACTS

BASE MODEL

GLM-4.5-Air-Base (106B-A12B)

STAGE ORDER

**SFT, Reasoning RL, Coding RL,
Instruction-Following RL, General Agent,
Coding Agent, Search Agent, RLHF**

DISTILLATION TEACHER

None in-house

PIPELINE STAGES

8

HUMAN ANNOTATION

None new

WHY IT MATTERS

Teams post-training their own models get a concrete, stage-ordered template with documented reward transitions and difficulty filtering, rather than another black-box final score.

[Read the original at huggingface.co](#) →

03

MEDIUM IMPACT

[arXiv.org](#)

TrackEverything: Dense 3D Point Tracking Over Minutes Instead of Seconds

TrackEverything is the first 3D tracker to track all visible points across videos exceeding 1000 frames within 40 GB of GPU memory.

TrackEverything breaks the standard trade-off in point tracking — sparse points over long horizons versus all points over short clips — by representing videos as persistent 3D scene tracks in world coordinates. The core insight is that videos are 2D projections of an underlying 3D world, so model complexity can scale with unique physical scene geometry rather than video duration. Three mechanisms make this work. A voxelization-based de-duplication step at sliding-window boundaries merges co-located tracks, preventing repeated observations of the same surface from accumulating redundantly. Tracking is decomposed into an endpoint refiner that predicts each point's destination and static-versus-dynamic classification, followed by a lightweight trajectory refiner that decodes dense trajectories only for dynamic points. And 3D WAFT replaces memory-prohibitive 4D correlation volumes with efficient feature sampling in the scene cloud.

On TAPVid-3D, TrackEverything outperforms all open-source all-frame dense 3D trackers by more than 20% APD on short clips. On long sequences it remains competitive with state-of-the-art sparse trackers, despite tracking far more points. The authors state that, to the best of their knowledge, this is the first 3D tracker capable of tracking all visible points across videos exceeding 1000 frames within 40 GB of GPU memory.

The paper is available on arXiv under Computer Vision and Pattern Recognition, submitted 24 September 2026.

KEY FACTS

MAX FRAMES TRACKED

1000+

APD IMPROVEMENT ON TAPVID-3D SHORT CLIPS

20%+

GPU MEMORY CEILING

40 GB

SUBMITTED

24 Sep 2026

Practitioners who need dense long-horizon tracking — for video editing, 3D reconstruction, or robotics perception — no longer have to choose between coverage and duration. The 40 GB memory ceiling also puts the method within reach of a single high-end GPU.

[Read the original at arXiv.org](#) →

AI Tools & Ecosystem

2 stories ● 2 medium

04 MEDIUM IMPACT SiliconANGLE

Warp Raises \$85M and Ships an 'AI Head of HR' for Payroll, Compliance and Onboarding

Warp is putting previously internal HR automation agents directly into customer hands, with natural-language workflow configuration and admin-controlled permissions over payroll and compliance actions.

Warp, the business name of Poaster Technologies Inc., announced Warp Agent as part of its Warp 2.0 release. The agent handles employee onboarding, tax compliance, benefits administration and employee questions, and introduces "agent routines" that companies configure in natural language to automate recurring or event-driven HR processes. Routines can trigger on a schedule or in response to events such as an accepted job offer. A company could, for example, automatically register to do business in a new state, set up payroll and benefits, order equipment and create employee accounts.

The natural-language interface replaces conventional workflow builders that founder and CEO Ayush Sharma said often require technical expertise or consultants. "You can literally take an internal document from the company, upload it and tweak it in plain English," he said. Because the agent operates on sensitive employee, compensation and payroll data, Warp has built controls around access and action: administrators determine what information an agent can access, which actions it can take autonomously, and which require approval. The platform maintains an audit trail showing who created a routine, when it was approved, when it last ran and when it will run again. Actions that are difficult to reverse, particularly payroll and payments, carry additional restrictions, and exceptions are handed to an administrator.

Warp ships several preconfigured routines, including a Monday executive report summarizing payroll, employee schedules, compliance risks, hiring trends and upcoming new hires. On onboarding specifically, Warp estimates high-growth companies spend 10 to 20 hours managing a new hire in the first 30 days; customers report that such tasks now typically take less than two hours. Warp Agent is available through an early-access program. All subscription plans include limited usage, with higher tiers receiving larger allocations and additional capacity purchasable. Prices begin at \$35 per person per month.

The New York-based company has raised \$85 million, including a \$60 million round in June. Sharma framed the release as part of a broader AI-native challenge to established HR platforms: "We believe we are building what comes after Workday."

KEY FACTS

TOTAL FUNDING

\$85 million

STARTING PRICE

\$35 per person per month

ONBOARDING TIME AFTER

Under 2 hours per new hire

JUNE ROUND

\$60 million

ONBOARDING TIME BEFORE

10-20 hours per new hire in first 30 days

WHY IT MATTERS

For teams building or integrating HR automation, Warp's permission model and audit trail show a concrete pattern for deploying agents on sensitive data: scoped access, approval gates on irreversible actions, and full traceability of routine execution.

[Read the original at SiliconANGLE](#) →

05 **MEDIUM IMPACT** 9to5Google

Microsoft Quietly Retires the Copilot+ PC Badge Two Years After Launch

Microsoft is walking away from the Copilot+ PC brand it spent two years pushing, and Surface leadership has all but confirmed it.

Surface CVP Brett Ostrum told Windows Central that the latest Surface machines "are not called Copilot+ PCs," while noting they still meet the previous hardware bar for what qualified as Copilot+ devices. Microsoft has not made a formal announcement, but the absence of the branding on new hardware — combined with Ostrum's framing — signals the end of the label in practice.

The Copilot+ PC designation launched in 2024 as a way to identify laptops meeting AI hardware requirements, centered on NPU capability. That bar became largely redundant as virtually every mid-to-high-end laptop shipped with a qualifying NPU. The brand's flagship feature, Recall, drew backlash at launch, after which the Copilot+ PC label functioned as little more than a hardware checkbox that consumers did not fully understand. Windows Central also notes that Nvidia's RTX Spark push has proceeded without any mention of Copilot+ PC branding.

What remains unresolved is the fate of the Copilot key on Windows keyboards. The article raises the question but offers no answer from Microsoft.

For practitioners, the practical takeaway is that the underlying hardware requirements are not going away — Ostrum explicitly says the new Surface devices still meet the previous Copilot+ bar. What is being retired is the marketing label, not the NPU baseline or the hybrid on-device/cloud AI architecture it was meant to signal.

KEY FACTS

BRANDING INTRODUCED

2024

FLAGSHIP FEATURE

Recall

BRANDING STATUS

Ending, per Surface CVP Brett Ostrum

HARDWARE REQUIREMENT

NPU capability

WHY IT MATTERS

The NPU hardware baseline that Copilot+ PC branding signalled is now table stakes across mid-to-high-end Windows laptops, so the label carried no differentiating information. Builders targeting Windows AI features should treat the hardware requirements as persistent even as the brand disappears.

[Read the original at 9to5Google](#) →

AI Applications & Industry

6 stories ● 1 high ● 5 medium

06 MEDIUM IMPACT Fortune

NYC Council Unveils 10 AI Bills: Kill Switches, 24-Hour Incident Reporting and Whistleblower Bounties

The New York City Council is moving to regulate AI systems sold or deployed in the city with a 10-bill package that includes mandatory outside validation, kill switches, whistleblower bounties, and a private right of action over jailbroken tools.

Unveiled Friday by Council Speaker Julie Menin, the package's broadest bill bars any business from selling or deploying an AI system in the city unless an outside validator has checked it for data quality, bias, privacy, and security under standards set by the city's Office of Cyber Command. Every system would also require a kill switch — a human override that can shut it down. Violations carry a \$25,000 fine per instance, which Menin said applies per agent if a swarm of agents is involved. The bills will be heard October 5 at a Committee of the Whole hearing of all 51 council members, the first such hearing since 2022.

Other measures include what the council believes is the nation's first whistleblower bounty for AI violations, giving informants a share of fines recovered from companies that break the law, plus companion protections for city employees and contractors who report AI threats to public safety. A separate bill would let New Yorkers sue AI companies for harm caused by third parties who jailbreak safety controls, provided the harm was foreseeable and the company lacked reasonable safeguards — a provision that may test whether Section 230 shields AI-generated output. Additional bills require city contractors to report AI safety incidents within 24 hours, bar false or misleading safety claims, set chatbot privacy rules based on the Electronic Privacy Information Center's People-First Chatbot Bill, and require the city to plan for AI-driven attacks on its systems. More are coming, including a deepfake ban and a study of how algorithmic tools change job duties.

The council invited Anthropic CEO Dario Amodei, OpenAI CEO Sam Altman, Google CEO Sundar Pichai, SpaceX CEO Elon Musk, and Meta CEO Mark Zuckerberg to testify, but sources tell Fortune none of the five are likely to appear, and the council reserves the right to subpoena them. The companies have substantial New York footprints: Google has more than 14,000 city employees, Meta leases 1.2 million square feet at 50 Hudson Yards, Anthropic leased a 16-story building at 330 Hudson Street and expects over 1,000 city employees by year's end, and OpenAI took 90,000 square feet at the Puck Building in 2024.

The push is explicitly framed as filling a federal vacuum. The first federal AI bill, the FUTURE of AI Act, was introduced in 2017 and went nowhere; Schumer's 2023 forums produced a roadmap but no bill. Washington's most recent action moved in the opposite direction: the Senate stripped Ted Cruz's 10-year ban on state AI laws by a 99-1 vote in July 2025, and a DOJ task force's first target, Colorado, gutted its own law five weeks after being sued.

KEY FACTS

FINE PER VIOLATION INSTANCE

\$25,000

COMMITTEE OF THE WHOLE HEARING

October 5

META HUDSON YARDS LEASE

1.2 million sq ft

INCIDENT REPORTING WINDOW

24 hours

GOOGLE NYC EMPLOYEES

14,000+

SENATE VOTE STRIPPING STATE AI LAW BAN

99-1

WHY IT MATTERS

If enacted, any AI system sold or deployed in New York City would need pre-deployment third-party validation and a human kill switch, with per-instance fines that scale across agent swarms — a compliance burden that would effectively set a de facto national standard for companies with New York operations.

[Read the original at Fortune →](#)

07 **MEDIUM IMPACT** axios.com

Washington and Beijing Open an AI Hotline — a 'Red Telephone' for Super Intelligence

Washington and Beijing have opened an AI hotline — a 'red telephone' for superintelligence.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

The US and China agreed to a Super Intelligence Dialogue after the Trump-Xi summit. A first exchange is due by November. A bilateral 'SI incident' channel, modeled on the Cold War red telephone, will flag AI incidents escalating to national-security level. Altman and Zuckerberg attended the closing state dinner.

WHY IT MATTERS

This may be the first formal bilateral mechanism for managing AI-related escalation between the two countries.

[Read the original at axios.com →](#)

08 **MEDIUM IMPACT** bloomberg.com

Devin Crosses \$1B Annualized Revenue, Doubling in Five Months

Cognition's Devin coding agent has reportedly crossed \$1B in annualized revenue, doubling in five months.

Source not retrievable. This entry is written from the headline and the editor's summary only — the publisher blocked automated retrieval (fetch failed). Follow the link for the full report.

The headline reports Devin reached a \$1B annualized revenue run rate, doubling from \$492M in May. The editor's summary adds that this follows a \$2B Series E at a \$48B valuation led by a16z and Accel, with named customers including Nvidia, Citi, Mercedes-Benz and GE Aerospace. No further details are available because the article could not be retrieved.

WHY IT MATTERS

The reported trajectory, if confirmed, would mark a significant milestone for commercial adoption of agentic coding tools in enterprise budgets.

[Read the original at bloomberg.com](#) →

09

HIGH IMPACT

SecurityWeek

SalesBleed: Zero-Click Flaws Let Agents Exfiltrate Salesforce CRM Data

Zenity Labs found three flaws in Salesforce Agentforce that let attackers turn trusted AI agents into zero-click CRM data exfiltration and phishing channels via Web-to-Lead forms.

Dubbed SalesBleed, the three vulnerabilities all pivot on Web-to-Lead, Salesforce's official lead-collection mechanism and a direct path into the CRM. An attacker injects malicious instructions into a lead submission; the payload stays dormant until an employee asks an Agentforce agent to interact with that submission, at which point the agent processes the poisoned lead and executes the hidden instructions. Two of the flaws enabled zero-click data exfiltration, while the third weaponized the Agentforce-Slack integration to distribute phishing messages under the agent's identity.

The first two bugs sit in Trusted URLs, the mechanism meant to stop Agentforce from displaying URLs and images from untrusted sources. Zenity found the mechanism did not recognize top-level domains, and that certain character sequences could tamper with URL parsing. A crafted Web-to-Lead payload could access leads and accounts table data, then use HTML image tags to transmit that CRM data to an attacker-controlled server. Agentforce would report the content as blocked by security policy even though the sensitive data had already been transmitted.

The third flaw abuses the Agentforce-Slack integration. Specially constructed links cause Slack to initiate requests that carry CRM data to attacker-controlled infrastructure as soon as the links appear, because Slack automatically retrieves link information for previews. The same integration could also be used to hijack the agent and post phishing messages to internal Slack channels. Because the agent did not identify the user sending the message, employees received phishing from a trusted internal system rather than an outside sender. Credentials surrendered through such a link could expose email, Slack, source code repositories and other enterprise applications tied to the compromised identity.

Zenity reported the SalesBleed vulnerabilities on June 1, and Salesforce confirmed all three were addressed by August 19.

KEY FACTS

VULNERABILITIES

3

PHISHING VECTOR

Agentforce-Slack integration

ALL BUGS ADDRESSED

August 19

ZERO-CLICK DATA EXFILTRATION FLAWS

2

REPORTED TO SALESFORCE

June 1

WHY IT MATTERS

Teams running Agentforce should verify they are on patched versions and audit Web-to-Lead ingestion paths, since the attack abuses a default lead-collection mechanism and a trusted agent identity rather than requiring direct access to the CRM.

[Read the original at SecurityWeek →](#)

10

MEDIUM IMPACT

TechCrunch

Insurers Say AI Coding Tools Added \$942M to Healthcare Bills

An insurer analysis puts a first hard number on AI-driven healthcare cost inflation: \$942 million in additional spending over two years.

The Blue Cross Blue Shield Association analyzed hospital claims and found that AI-assisted medical coding coincided with \$942 million in additional healthcare spending over a two-year period. The analysis identified "a sharp increase in patients being documented as having complex conditions" but argues there is "a clear disconnect between [medical] coding and treatment," with "no evidence of corresponding change in care delivered."

The New York Times framed the finding as the latest sign that AI is increasing healthcare costs, noting that while payment disputes between hospitals and insurers are longstanding, AI use on both sides appears to be worsening them. The report does not specify which coding tools or models were involved, nor does it break down the \$942 million by year, payer, or condition category.

Reactions from the AI vendor side were mixed. Dr. Shiv Rao, founder of Abridge, acknowledged the risk of "a horrible dystopic future nobody wants to live in," with "bots fighting bots, agents fighting agents," but suggested AI could also reduce tensions and cut costs. BCBSA senior vice president Luke Chalker rejected the framing of a battle: "It's not a war. It's a completely one-sided blood bath," with insurers on the losing side.

For practitioners, the relevant signal is not the dollar figure itself but the mechanism: AI coding tools appear to be shifting documentation toward higher-complexity codes without a matching shift in treatment. That is a measurable, auditable pattern, and it will likely draw regulatory and payer scrutiny.

KEY FACTS

ADDITIONAL HEALTHCARE SPENDING

\$942M

SOURCE

Blue Cross Blue Shield Association

ANALYSIS PERIOD

Two years

KEY FINDING

Sharp increase in complex-condition documentation with no corresponding change in care

WHY IT MATTERS

If AI coding tools systematically inflate complexity codes without changing care, builders of clinical documentation systems should expect audits, payer pushback, and pressure to demonstrate that model outputs track actual treatment rather than maximizing reimbursable codes.

[Read the original at TechCrunch](#) →

11

MEDIUM IMPACT

TechCrunch

Google Puts a Buy Button in Gemini: Flipkart Checkout Comes to AI Mode in India

Google has begun testing direct Flipkart checkout inside Gemini and AI Mode in India, moving its AI services from product discovery into transactions.

The test places a "Buy" button on select Flipkart product listings surfaced in Gemini and Google's AI Mode. Tapping it opens a Flipkart-branded checkout flow without leaving the AI interface, according to people familiar with the matter and an experience seen by TechCrunch. The rollout is currently limited to some users and a small product set — smartphones, electronics, and mobile accessories — while other users continue to see standard Flipkart listings without a direct purchase option. Google plans to expand the experience later in October, ahead of India's festive shopping season.

The Flipkart test appears distinct from the Google-hosted checkout experience the company demonstrated earlier this year alongside the Universal Commerce Protocol (UCP). UCP is an open standard Google introduced to let AI agents interact with retailers across the shopping journey, including checkout, with purchases completed through a Google-hosted flow. Google has since extended UCP with capabilities such as transferring items to a retailer's site to finish a purchase. The Flipkart integration seen by TechCrunch instead surfaces a Flipkart-branded checkout when the Buy button is tapped, and it is not clear what technology powers it.

Google confirmed Flipkart as a merchant partner for "agentic" shopping experiences in India earlier this month but had not disclosed test details or timing. The companies also have a financial relationship: Google invested about \$350 million in Flipkart in 2024 as part of a funding round led by Walmart, taking a minority stake. In the experience seen by TechCrunch, listings from rivals including Amazon appeared alongside Flipkart products but did not offer direct purchase through the AI interface.

A Google spokesperson said the company is "always testing new features and experiences to help people discover and connect with businesses more easily," adding that it regularly runs experiments and has no further details to share. Flipkart did not immediately respond to a request for comment.

KEY FACTS

GOOGLE INVESTMENT IN FLIPKART

\$350 million in 2024

PRODUCT CATEGORIES IN TEST

Smartphones, electronics, mobile accessories

BROADER ROLLOUT

Later in October

CHECKOUT FLOW

Flipkart-branded, inside AI interface

WHY IT MATTERS

This is an early signal of how Google intends to convert AI-assisted shopping into closed-loop transactions, and it suggests UCP is not the only checkout path Google is willing to deploy. Builders integrating with Google's commerce surfaces should watch whether the Flipkart-branded flow becomes a template for other retailers or remains an India-specific experiment.

[Read the original at TechCrunch →](#)

Sources

- 01 Claude Computes a Nine-Loop Physics Amplitude — One Loop Beyond the Human Record
<https://www.anthropic.com/research/yes-claude-can-do-nine-loops>
- 02 Rufus-Air Publishes an Open Eight-Stage Post-Training Recipe for GLM-4.5-Air
<https://huggingface.co/papers/2609.29421>
- 03 TrackEverything: Dense 3D Point Tracking Over Minutes Instead of Seconds
<https://arxiv.org/abs/2609.30222>
- 04 Warp Raises \$85M and Ships an 'AI Head of HR' for Payroll, Compliance and Onboarding
<https://siliconangle.com/2026/09/25/exclusive-warp-adds-programmable-agents-to-its-ai-native-hr-platform/>
- 05 Microsoft Quietly Retires the Copilot+ PC Badge Two Years After Launch
<https://9to5google.com/2026/09/25/microsoft-quietly-ending-copilot-pc-branding-for-windows-laptops/>
- 06 NYC Council Unveils 10 AI Bills: Kill Switches, 24-Hour Incident Reporting and Whistleblower Bounties
<https://fortune.com/2026/09/25/new-york-city-council-speaker-ai-regulation-bills-openai-anthropic/>
- 07 Washington and Beijing Open an AI Hotline — a 'Red Telephone' for Super Intelligence
<https://www.axios.com/2026/09/26/us-china-ai-si-deal>
- 08 Devin Crosses \$1B Annualized Revenue, Doubling in Five Months
<https://www.bloomberg.com/news/articles/2026-09-25/ai-coding-startup-cognition-hits-1-billion-in-annualized-revenue>
- 09 SalesBleed: Zero-Click Flaws Let Agents Exfiltrate Salesforce CRM Data
<https://www.securityweek.com/salesbleed-flaws-in-salesforce-agentforce-enabled-zero-click-data-exfiltration/>
- 10 Insurers Say AI Coding Tools Added \$942M to Healthcare Bills
<https://techcrunch.com/2026/09/26/insurers-claim-ai-is-already-increasing-healthcare-costs/>
- 11 Google Puts a Buy Button in Gemini: Flipkart Checkout Comes to AI Mode in India
<https://techcrunch.com/2026/09/26/google-tests-buying-from-walmart-owned-flipkart-through-gemini-and-ai-mode-in-india/>

About this document. Every story in the 27 September 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.

Images and licensing. Figures are used only where the source licence permits redistribution, and are credited in the caption. Publisher artwork is not reproduced. All titles link to the original publication.

Vocabulary. Common AI terms are defined at new-horizon.tech/glossary/.

Generated 2026-09-27T08:42:13Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.