

The Long Read

12 October 2026

10

STORIES

3


SECTIONS

2255

WORDS

2

HIGH IMPACT



● 2 high impact ● 8 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

10 stories · 2255 words · 3 sections · 2 charted



AI Models & Research

- 01 ● OpenAI Quietly Lists a Second GPT-Rosalind Model for Life Sciences
- 02 ● AI Models Can Provably Hide Arbitrary Capabilities — and Evals May Never Find Them
- 03 ● X-AgentRed: Red-Teaming Personal Agents Across Multi-Service Ecosystems

AI Tools & Ecosystem

- 04 ● RATED: a New Benchmark Judges Agent Tool Retrieval by Execution, Not Semantics
- 05 ● OpenAI Adds Composer Predictions to Codex — the Agent Now Anticipates Your Next

Change

AI Applications & Industry

- 06 ● Voice AI Hasn't Had Its ChatGPT Moment Yet — so Say the Execs at Cerebral Valley
- 07 ● A Week After Amazon's NDA Reversal, AI Data-Center Builders Bet Transparency Can Soften Opposition
- 08 ● New Zealand's Greens Run on AI Regulation: Moratoriums, Labeling and Data Sovereignty
- 09 ● Salesforce Renames AIForce to SIForce — the Super-Intelligence Rebrand Reaches the Fortune 500
- 10 ● Will the HomeView Be Apple's First Smart Home Success?

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 ● 1 high ● 2 medium

01 **MEDIUM IMPACT** TestingCatalog AI News

OpenAI Quietly Lists a Second GPT-Rosalind Model for Life Sciences

OpenAI has quietly added a second GPT-Rosalind model, gpt-rosalind-discovery, to its API pricing page without a blog post, system card, or help article.

The listing appears under a new "Life sciences models" section, where gpt-rosalind-discovery sits beside gpt-rosalind-research. Both carry identical pricing: \$5 per million input tokens, \$0.50 per million cached input tokens, and \$25 per million output tokens. Billing for gpt-rosalind-research began October 5, 2026, and cache-write pricing does not apply. Access remains restricted to approved internal research through OpenAI's trusted-access program, with data residency endpoints adding a 10% uplift for eligible models released on or after March 5, 2026.

OpenAI has not described what gpt-rosalind-discovery does. The name points toward drug discovery, and the company's April introduction of GPT-Rosalind framed it as a frontier reasoning model for biology, drug discovery, and translational medicine, with Amgen, Moderna, the Allen Institute, and Thermo Fisher Scientific among early partners. The help center describes the current model as most useful for early discovery and hypothesis generation—target biology, literature synthesis, omics interpretation. A variant tuned for target identification or molecular design would fit, but that remains inference.

OpenAI has also stated that "all eligible organizations will continue to get access to the latest GPT-Rosalind models as they're released," confirming a planned family rather than a single model. Context: GPT-Rosalind left research preview on September 11 and opened to eligible organizations worldwide through trusted access; published pricing took effect October 5. The latest version, GPT-Rosalind-5.5, is trained on GPT-5.5 and, per its system card, trained not to refuse sophisticated biology questions, with gated access as the primary safeguard—the same approach OpenAI used for GPT-5.4-Cyber.

OpenAI often lists models on the pricing page shortly before announcing them, so a formal announcement may be imminent. The open questions are whether this is a separately trained model or a tuned variant of the research model, and whether trusted-access requirements change.

KEY FACTS

INPUT PRICE

\$5 per million tokens

OUTPUT PRICE

\$25 per million tokens

DATA RESIDENCY UPLIFT

10%

CACHED INPUT PRICE

\$0.50 per million tokens

BILLING START (RESEARCH MODEL)

October 5, 2026

ACCESS

**Trusted-access program, approved
internal research**

WHY IT MATTERS

Teams already in the trusted-access program should watch for a formal announcement and prepare to evaluate whether a discovery-focused variant changes their biology workflows; those outside the program get no immediate access but a signal of where OpenAI is investing.

[Read the original at TestingCatalog AI News →](#)

02

HIGH IMPACT

arXiv.org

AI Models Can Provably Hide Arbitrary Capabilities — and Evals May Never Find Them

Researchers have constructed backdoors that embed encrypted neural networks inside host model weights, decrypting and running only on a secret trigger — and they prove, under cryptographic assumptions, that elicitation-based evals cannot surface them.

The paper shows that capability evaluations rest on an assumption that fails under adversarial conditions: that hidden capabilities can be elicited. The authors build backdoor attacks that embed encrypted neural networks within host model weights using digital lockers. The hidden circuits decrypt and execute only upon receiving a secret trigger, and remain provably hard to elicit or interpret under cryptographic assumptions even with full access to model weights. This extends prior work from hiding fixed strings to hiding arbitrary computations.

The proof-of-concept constructions are implemented in PyTorch and evaluated on a small chemical reaction prediction task, used as a proxy for CBRN-relevant capabilities. The authors test elicitation through fine-tuning, weight noising, gradient-based trigger search, and latent adversarial training. The hidden circuits resist all of these methods. The cryptographic guarantees, combined with the empirical resistance, point to a hard limit on elicitation-based capability audits.

The models are released to support work on defenses. The submission is dated 9 Oct 2026 on arXiv.

What is new here is the jump from hiding static data to hiding arbitrary computation, with formal hardness rather than empirical obscurity. Prior backdoor work generally hid fixed strings or simple behaviors; this construction makes the hidden capability itself a full neural network that is cryptographically sealed until triggered.

KEY FACTS

SUBMISSION DATE

9 Oct 2026

HIDDEN MECHANISM

Digital lockers

ELICITATION METHODS TESTED

Fine-tuning, weight noising, gradient-based trigger search, latent adversarial training

IMPLEMENTATION

PyTorch

EVALUATION TASK

Small chemical reaction prediction

MODELS RELEASED

Yes

WHY IT MATTERS

Teams relying on red-teaming, fine-tuning, or gradient-based probing to surface dangerous capabilities should treat a clean eval as weak evidence: a model can carry a fully functional hidden network that no current elicitation method will expose. Defenses need to move beyond elicitation toward weight-level or cryptographic inspection.

[Read the original at arXiv.org](#) →

03

MEDIUM IMPACT

[arXiv.org](#)

X-AgentRed: Red-Teaming Personal Agents Across Multi-Service Ecosystems

X-AgentRed red-teams personal agents by building a directed graph of tool dependencies and using execution feedback to discover harmful cross-service tool call paths that prompt-based methods miss.

The framework addresses a specific failure mode in existing red-teaming approaches: when LLMs are prompted to generate harmful paths from task descriptions, generation is biased toward tool combinations seen during training, leaving novel cross-service interactions underexplored. X-AgentRed instead constructs a directed graph by analyzing each tool's effects, potential harms, and execution preconditions. Edges in this graph capture how one tool's effects can satisfy another tool's execution preconditions. Recipes sampled from the graph are instantiated as realistic scenarios in which agents can produce harmful outcomes while fulfilling benign user requests.

Execution feedback drives self-improvement through two mechanisms. When a path produces harm, the framework diversifies it through tool substitution. When a path is blocked, the framework refines preconditions at the blocked step to enable further execution. This iterative loop expands the space of discovered harmful paths beyond what static generation can reach.

The paper reports results across four frontier agentic models. X-AgentRed achieves $1.77\times$ the harmful path rate of STAC and explores a more diverse set of tools, yielding $2.99\times$ the tool coverage on average. The comparison against STAC is notable because STAC represents the prompt-based approach the authors argue is limited by training-data bias.

The contribution is less about a new model and more about a search strategy: grounding red-teaming in an explicit dependency graph rather than in LLM generation alone. The execution-feedback loop is the genuinely

new element, turning failed attack paths into refinement signals rather than discarding them.



KEY FACTS

HARMFUL PATH RATE VS STAC

1.77x

FRONTIER AGENTIC MODELS EVALUATED

4

TOOL COVERAGE VS STAC

2.99x

SUBMITTED

9 Oct 2026

WHY IT MATTERS

Teams deploying personal agents across multiple services should treat cross-service tool interactions as an attack surface that prompt-based red-teaming will systematically under-sample. A dependency-graph approach with execution feedback surfaces compositional failures that would otherwise ship.

[Read the original at arXiv.org](#) →

AI Tools & Ecosystem

2 stories ● 1 high ● 1 medium

04 **HIGH IMPACT** [arXiv.org](#)

RATED: a New Benchmark Judges Agent Tool Retrieval by Execution, Not Semantics


RATED re-ranks tool candidates on live execution feedback, cutting fault selection from 86.2% to 2.9% during severe API outages.

RATED (Reliability-Aware Tool Retrieval via Execution Dynamics) pairs a frozen dual-encoder with contrastive fine-tuning for semantic matching against a lightweight contextual bandit module that re-ranks candidates using real-time execution feedback. Because the bandit adapts to runtime tool failures without retraining the encoder, the framework responds to API outages as they happen rather than relying on a static relevance score.

The accompanying EcoAgent benchmark comprises 5,172 real-world tools and over 210,000 query-tool pairs, split between offline retrieval and dynamic reliability evaluation. On real-world queries, RATED posts a 52.2% relative MRR improvement over the ToolLLM API retriever. Under a severe outage scenario with a 30% failure rate, fault selection rate drops from 86.2% to 2.9%. Latency overhead is under 1 ms over encode-only inference, which matters for agents that chain multiple tool calls per turn.

What is new here is the explicit modeling of tool availability as a first-class retrieval signal. Prior work optimized semantic relevance alone; RATED treats reliability as a dynamic property learned from execution outcomes. The bandit is deliberately lightweight, so the adaptation loop does not require expensive retraining or a separate serving stack.

Code and dataset are public, which makes the benchmark immediately usable for teams evaluating their own retriever-plus-tool-execution pipelines.



KEY FACTS

RELATIVE MRR IMPROVEMENT OVER TOOLLLM API RETRIEVER

52.2%

BASELINE FAULT SELECTION RATE UNDER OUTAGE

86.2%

ECOAGENT TOOLS

5,172

FAULT SELECTION RATE UNDER 30% API OUTAGE

2.9%

LATENCY OVERHEAD OVER ENCODE-ONLY INFERENCE

<1 ms

ECOAGENT QUERY-TOOL PAIRS

210,000+

WHY IT MATTERS

Anyone building agents that call external APIs should treat tool availability as a ranking signal, not a post-hoc error handler. The public EcoAgent benchmark gives a concrete way to measure that behavior under failure.

[Read the original at arXiv.org](#) →

05 **MEDIUM IMPACT** help.openai.com

OpenAI Adds Composer Predictions to Codex — the Agent Now Anticipates Your Next Change

OpenAI has added Composer predictions to Codex, letting the coding agent anticipate and pre-draft a user's next change.

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 October 9 release note reportedly ships Composer predictions to Codex. The coding agent now anticipates your next change and pre-drafts it. The editor notes this is small on its face but marks a moment when coding tools shift from responding to predicting. No further details are available from the source.

WHY IT MATTERS

If accurate, predictive pre-drafting could change how developers interact with AI coding agents, though the practical impact remains unverified.

[Read the original at help.openai.com](https://help.openai.com) →

AI Applications & Industry

5 stories ● 5 medium

06 MEDIUM IMPACT TechCrunch

Voice AI Hasn't Had Its ChatGPT Moment Yet — so Say the Execs at Cerebral Valley

Voice AI has full-duplex models but still lacks the reasoning speed and user trust needed for a ChatGPT-scale breakthrough, according to executives at PolyAI and Otter.

PolyAI CTO Shawn Wen, speaking at the HumanX conference, said the field has reached the milestone of full-duplex models — systems that can speak while listening — but that the next challenge is making reasoning fast enough that answers arrive quickly and conversation feels natural. Without that, he argued, voice AI has not had its ChatGPT moment. Wen framed the progression in customer service terms: once voice quality is good enough for callers to engage for the first two or three turns, they build confidence and may conclude they do not need a human if the agent can solve the problem.

Otter CMO Alex Gay identified speaker identification, intent capture, and connecting that to organizational knowledge as the key step toward automation. Otter is working on digital twins that could represent people in meetings, and Gay said the output voice must carry the same emotive expressions as a human participant. Without the ability to debate, hold strategic discussions, and feel a relationship underpinning the conversation, he said, an avatar is just a Q&A chatbot.

Both executives pointed to transcription accuracy as a foundational weakness. Wen said ASR models often miss important keywords, which breaks context capture. Gay agreed, noting that for Otter, transcription was never the endpoint but the layer on which productivity gains are built. If the original transcription lacks accuracy, all downstream actions become flawed, and the moment the platform takes a wrong action, users lose trust. He called continued ASR improvement critical because the downstream impacts are significant.

On transparency, both argued that tools should declare when people are being recorded or talking to AI. Otter wants to instill trust even in meetings where its bot is not present, exploring methods like notifying everyone in the chat that a meeting is being recorded. PolyAI's Wen said establishing that people are talking to an AI in enterprise calls is important.

KEY FACTS

FULL-DUPLEX MILESTONE

Reached, per PolyAI CTO Shawn Wen

OTTER FOCUS AREAS

Speaker identification, intent capture, organizational knowledge

ASR WEAKNESS

Models often miss important keywords, breaking context capture

NEXT CHALLENGE

Making reasoning very fast for natural conversation

OTTER EMERGING TECH

Digital twins representing people in meetings

TRANSPARENCY POSITION

Tools should declare recording or AI interaction to users

WHY IT MATTERS

Builders should treat ASR accuracy and reasoning latency as the binding constraints on voice agent adoption, not voice naturalness alone. Trust mechanisms — disclosure and reliable transcription — are prerequisites before downstream automation can be deployed without eroding user confidence.

[Read the original at TechCrunch →](#)

07

MEDIUM IMPACT

TechCrunch

A Week After Amazon's NDA Reversal, AI Data-Center Builders Bet Transparency Can Soften Opposition

Amazon has joined Microsoft in abandoning nondisclosure agreements during negotiations with local governments over data centers, a move aimed at the secrecy that has fueled community backlash.

The announcement came buried in a longer Amazon blog post arguing that data centers benefit communities, rather than as a standalone policy statement. Sean O'Kane noted on TechCrunch's Equity podcast that this placement is telling: the NDA reversal sits at the root of many problems the industry faces, because projects are often negotiated essentially in secret — sometimes with the public knowing a data center is coming but not who will occupy it.

Microsoft made a similar announcement earlier this year. The podcast panel was measured about the effect. Rebecca Bellan pointed out that many moratoriums passed by communities last only a year or two, while data centers are not coming online for two to five years anyway, so the political pushback may not yet be biting. Anthony Ha cited New York's moratorium, which is technically in place until the state finalizes an environmental review process expected to take a year — framing that leaves room for approval later. Still, he argued the builders are clearly worried, pointing to what he called unhinged tweets from figures like Trump and David Sacks.

O'Kane framed the deeper issue as a communication failure across the tech industry, comparing it to Uber's mid-2010s secret negotiations for tax subsidies outside Orlando, where secrecy let the company ask for more than it otherwise could. He also cited Tesla's 'amazing abundance' mission, announced a year ago with a multi-page PDF that even fans criticized as unspecific and 'LLM nonsense.'

The panel's consensus: dropping NDAs is a good first step, but it will not automatically convert opponents. It does, however, open the conversation by removing the baseline objection that communities do not know what

is being negotiated.

KEY FACTS

AMAZON NDA POLICY

Will no longer use NDAs with local governments over data centers

TYPICAL MORATORIUM LENGTH

One to two years

NEW YORK MORATORIUM

Until state finalizes environmental review, expected to take a year

MICROSOFT NDA POLICY

Made similar announcement earlier this year

DATA CENTER LEAD TIME

Two to five years before coming online

WHY IT MATTERS

For teams planning data-center capacity, local permitting is becoming a contested, public-facing process rather than a private negotiation. Expect longer timelines and the need to make an explicit economic and environmental case to communities, not just to local officials.

[Read the original at TechCrunch →](#)

08

MEDIUM IMPACT

interest.co.nz

New Zealand's Greens Run on AI Regulation: Moratoriums, Labeling and Data Sovereignty

New Zealand's Green Party has released an AI policy platform that would impose consent requirements on training data, environmental assessments on data centres, and labelling mandates on AI-generated content.

The policy, announced by co-leader Chlöe Swarbrick on Monday, frames AI as a regulatory problem rather than an inevitability. The document argues that trillions of dollars invested in AI tools constitute a growing bubble presenting risks to jobs and economic stability, and that government should regulate to ensure benefits are shared and harms minimised.

On data governance, the Greens want consent standards for AI training data, remuneration for intellectual property used in model training with fines for breaches, and specific regulation requiring co-governance and co-development of Māori data and IP. Government AI systems would need transparent datasets, and onshoring public data should be prioritised. The policy also proposes privacy codes covering biometric processing and facial recognition technology.

For public-sector deployment, the plan would prohibit fully automated decisions on matters that could make a substantive difference to people's lives, require public registers of government AI systems, and establish a right to challenge AI-assisted decisions or opt out of them. A new AI office would sit within the Department of Prime Minister and Cabinet, with algorithm audits for high-risk systems.

Environmental provisions target hyperscale data centres specifically: mandatory environmental impact assessments, requirements to provide new renewable generation covering full energy use, water-use limits, and a pause on consenting new hyperscale facilities until a regulatory framework exists. The policy also calls for withholding tax rules to prevent tech companies offshoring profits, payments to local media outlets for data

used in AI summaries, mandates for multiple interchangeable AI providers when using third-party models, and research support for New Zealand-based LLMs to reduce dependence on foreign vendors.

KEY FACTS

POLICY ANNOUNCED

Monday, by Green Party co-leader Chlöe Swarbrick

PROPOSED AI OFFICE

Within Department of Prime Minister and Cabinet

PUBLIC SECTOR AUTOMATION

No fully automated decisions on substantive matters

AI INVESTMENT CITED

Trillions of dollars

DATA CENTRE CONSENTING

Pause on new hyperscale facilities until regulatory framework exists

WHY IT MATTERS

If enacted, these rules would reshape how AI systems are procured, trained, and deployed in New Zealand — particularly the consent and remuneration requirements for training data, the pause on hyperscale data centre consenting, and the public-sector ban on fully automated decisions with substantive impact.

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

09

MEDIUM IMPACT

TNW | Artificial-intelligence

Salesforce Renames AIForce to SIForce — the Super-Intelligence Rebrand Reaches the Fortune 500

Salesforce is renaming its AIForce agent platform to SIForce, following a Trump executive order that tells federal agencies to replace 'artificial intelligence' with 'super intelligence.'

Chief executive Marc Benioff announced the change on Friday, writing on X: "The era of Super Intelligence is here. AIForce is officially SIForce." The rename covers the agent platform only, not Salesforce itself. Business Insider's Katherine Li first reported the move.

Benioff is the second major tech executive to make such a change this week. Last Sunday, Elon Musk said SpaceXAI would become SpaceXSI. Both follow a 29 September executive order directing federal agencies to switch terminology in official documents and websites. The order does not bind private companies, but Trump escalated on Thursday, writing on Truth Social that anyone who still says "artificial intelligence" is "THE ENEMY."

The rebrand is purely nominal. Superintelligence conventionally denotes systems far more capable than humans across many tasks, and renaming a product does not demonstrate that threshold has been met, as Business Insider noted. The push also puts companies with AI in their names, such as OpenAI and Scale AI, in an awkward position; neither responded to requests for comment.

For practitioners, nothing about the underlying platform has changed. The announcement contains no new capabilities, benchmarks, model releases, or technical specifications. It is a terminology shift driven by political pressure rather than engineering progress.

KEY FACTS

ANNOUNCED

Friday

SCOPE

Agent platform only, not company name

EXECUTIVE ORDER DATE

29 September

SECOND RENAME THIS WEEK

SpaceXAI to SpaceXSI

WHY IT MATTERS

The rename signals that terminology pressure from Washington is reaching enterprise AI vendors, but it changes nothing about what SIForce can actually do. Teams evaluating the platform should treat this as branding, not a capability upgrade.

[Read the original at TNW | Artificial-intelligence →](#)

10

MEDIUM IMPACT

TechCrunch

Will the HomeView Be Apple's First Smart Home Success?

Apple is making its most serious smart home bet to date with the HomeView, a Siri-powered display hub built on the A18 chip.

Bloomberg's Mark Gurman reports that Apple will announce the HomeView on Tuesday, a smart home hub with a display and an A18 chip that lets customers control other smart home devices through Siri. It will ship in two form factors: one for countertops and one for wall mounting.

The launch is paired with refreshes to two existing home products. Apple will also unveil updated versions of the Apple TV set-top box and the HomePod mini on the same day. Earlier in the week, Gurman said Apple plans to announce a partnership with LG covering a smart lock, a thermostat, and a doorbell.

The HomeView arrives after Apple's prior smart home efforts failed to gain traction. HomeKit, its software framework, and the HomePod smart speaker both underperformed. The HomeView represents a hardware-led reset: an on-device display with Apple silicon, positioned as a control surface rather than a passive voice endpoint.

For developers and integrators, the notable shift is architectural. The A18 chip puts the device in the same silicon generation as current iPhones, which suggests on-device Siri processing and local execution of home automations rather than cloud round-trips. The LG partnership also signals a move toward first-party accessory hardware, expanding the Matter-compatible surface area Apple can control directly.

KEY FACTS

ANNOUNCEMENT DATE

Tuesday

MODELS

Countertop and wall-mounted

ALSO UNVEILING

Updated Apple TV and HomePod mini

CHIP

A18

LG PARTNERSHIP PRODUCTS

Smart lock, thermostat, doorbell

A display hub with an A18 chip implies local inference and on-device automation execution, which changes latency and privacy assumptions for anyone building HomeKit or Matter integrations. The LG accessory line also gives Apple a direct hardware path into locks, climate, and entry control.

[Read the original at TechCrunch](#) →

Sources

- 01 OpenAI Quietly Lists a Second GPT-Rosalind Model for Life Sciences
<https://www.testingcatalog.com/openai-lists-new-gpt-rosalind-discovery-model-in-docs>
- 02 AI Models Can Provably Hide Arbitrary Capabilities — and Evals May Never Find Them
<https://arxiv.org/abs/2610.13928>
- 03 X-AgentRed: Red-Teaming Personal Agents Across Multi-Service Ecosystems
<https://arxiv.org/abs/2610.13886>
- 04 RATED: a New Benchmark Judges Agent Tool Retrieval by Execution, Not Semantics
<https://arxiv.org/abs/2610.13889>
- 05 OpenAI Adds Composer Predictions to Codex — the Agent Now Anticipates Your Next Change
<https://help.openai.com/en/articles/6825453-chatgpt-release-notes>
- 06 Voice AI Hasn't Had Its ChatGPT Moment Yet — so Say the Execs at Cerebral Valley
<https://techcrunch.com/2026/10/11/these-execs-think-voice-ai-hasnt-reached-its-chatgpt-moment-yet/>
- 07 A Week After Amazon's NDA Reversal, AI Data-Center Builders Bet Transparency Can Soften Opposition
<https://techcrunch.com/2026/10/11/can-the-ai-industry-persuade-data-center-opponents-by-getting-rid-of-ndas/>
- 08 New Zealand's Greens Run on AI Regulation: Moratoriums, Labeling and Data Sovereignty
<https://www.interest.co.nz/technology/140641/greens-ai-plan-targets-data-sovereignty-environmental-rules-and-labelling>
- 09 Salesforce Renames AIForce to SIForce — the Super-Intelligence Rebrand Reaches the Fortune 500
<https://thenextweb.com/news/salesforce-aiforce-siforce-trump-super-intelligence>
- 10 Will the HomeView Be Apple's First Smart Home Success?
<https://techcrunch.com/2026/10/11/will-the-homeview-be-apples-first-smart-home-success/>

About this document. Every story in the 12 October 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 2026-10-12T05:34:29Z · model deepseek-v4-pro:cloud · New Horizon, new-horizon.tech · To unsubscribe or manage your subscription, see the footer of the daily email.