



Emerging Technologies Roundtable

9:00-11:00am

**August 4, 2026
Q3 2026**

- **Welcome and Introductions**
- **Mark Lauby, Chief Engineer at NERC and IEEE Fellow, provides an update on *Large Grid Loads and Power System Behavior with Renewables***
- **Portfolio Update**
- **Looking Ahead**

**Mark Lauby, Chief Engineer at NERC
and IEEE Fellow:**
***Large Grid Loads and Power System
Behavior with Renewables***

State of Reliability

Mark G. Lauby
NERC Fellow and Chief Engineer
ARC August 2026

About NERC: Mission

To assure the effective and efficient reduction of risks to the reliability and security of the North American bulk power system

- Develop and enforce reliability standards
- Assess current and future reliability
- Analyze system events and recommend improved practices
- Encourage active participation by all stakeholders
- Accountable as ERO to regulators in the United States (FERC) and Canada (NEB and provincial governments)



How did we get here?

1960s: 1965 Blackout, formed North American Electric Reliability Council (NERC) in 1968; 1 event in 10 years

1970s: Large central stations

1980s: Gas Reform and Combined-Cycle Gas Turbines

1990s: Restructuring and Markets

2000s: Hurricanes=> Price Volatility, Great Recession; 2003 blackout; 2005 EPACT, 2007 Electric Reliability Organization formed NERC

2010s: Buildout of variable energy resources; Questions about 1 day in 10 years


And now today...

Hyper Complex Risk Environment


NERC




Rapid Growth



System Events





Emerging Large Loads exhibit BPS risk

Certain Loads (Data Centers) are of concern

High Priority
Risks

LONG-TERM PLANNING

- Resource Adequacy

OPERATIONS/BALANCING


- Balancing and Reserves

STABILITY

- Ride-through
- Voltage Stability
- Angular Stability
- Oscillations

Existing Reliability Standards, processes, and requirements for Transmission/Distribution planning, operations, security, and other areas are **inadequate** to address the risks posed by emerging large loads.

Updates are needed to registration criteria, Reliability Standards, processes/requirements, and other areas.



Design matters for density, resilience, & scalability

- **Equipment Matters**

- AC versus DC architecture
- UPS operating modes
- Cooling
- Ride-through
- Capacity
- Parallel operation of Gen

- **AI job scheduling, training, & use affect BPS reliability/security**

- Data sharing with utilities
- Models
- Ride-through




Figure 2.1: Tier I–IV Data Center One-Line⁹

S. Chalise et al., "Data center energy systems: Current technology and future direction," 2015 IEEE Power & Energy Society General Meeting, Denver, CO, 2015, pp. 1-5, doi: 10.1109/PESGM.2015.7286420.


How AI workloads are changing infrastructure requirements?

NERC


1. Impacts on BPS Reliability

- Performance on Disturbances
- Reduced generator stability margin
- Power Oscillations and resonance
- Long-term Planning (resource + Transmission)
- Operations and Balancing
- Physical and Cyber Security
- Emergency operations (UFLS + restoration)

2. Large Load Action Plan



Building the Bridge



Top Priorities



Energy Analysis



Weather hardening of
Generation and
Transmission




Inverter-based Resource
Integration



System Operations with
DER (especially impacts
on UFLS and UVLS)



Gas-Electric
Interdependency



Large Load (Data Centers,
Cryptocurrency, and AI)
Integration



Volatility of Energy Policy

Lower Antelope Canyon – Navajo Nation

NERC



Top Priorities

NERC



Top Priorities

NERC



Top Priorities

NERC



A light blue map of the United States with white state borders. The map is centered in the background of the slide.

NERC

Discussion

Revisiting *Q2 Looking Ahead*

Artificial Intelligence

Transportation

Energy, Climate, Water and Environment/Computing

Virtual Presence and Gaming

Manufacturing and Construction

Supply Chain and Logistics

Fintech

Food and Agriculture

Robotics and Automation

Aging and Health

Safety and Security

Infrastructure and Connected Communities

Q2 Looking Ahead > Q3


The Trend Report is dead. We killed it.

Introducing the Convergence Outlook.

This would have been FTSG's 19th annual tech trends report. Instead, we made a deliberate choice to retire the format.

This year marks the launch of something new: the first edition of our Convergence Outlook. You're still getting our best-in-class research. What's changed is our framing. The insights ahead are designed to be more urgent, more actionable, and far more relevant to the decisions leaders are facing right now.

If the old tech trend report helped you see what's coming, we built our Convergence Outlook to help you decide what to do next.



The Convergence Landscape

The Convergence Landscape is a map of pressure mapping where the ground beneath your organization is starting to shift — before anything breaks. It's a guide that answers three questions at a glance:

Where is the ground still relatively stable?

Where is stress building beneath the surface?

Where is structural change most likely to occur first?

Use this map to locate your organization, then turn to the sections that follow for what's driving the pressure and what to do about it.

1 Stable Ground

These are areas where systems can still absorb change incrementally.

- ▼ Business models still work
- ▼ Regulations still hold
- ▼ Infrastructure is not yet strained

Action: Optimize and monitor — but don't mistake calm for permanent.

2 Pressure Zones

These are areas where multiple forces are interacting and adaptation is becoming harder.

- ▼ Constraints are tightening
- ▼ Timelines are compressing
- ▼ Tradeoffs are increasing

Action: Stress-test assumptions and accelerate decision cycles.


3 Fault Lines

These are points where small shocks can trigger outsized change.

- ▼ Market structures shift
- ▼ Value chains reconfigure
- ▼ Governance models fail or reset

Action: Scenario plan now. Waiting for confirmation is too late.

— Stable Ground ■ Pressure Zones --- Fault Lines



Convergences

- Compute Shock
- Polycompute
- Agentic Economies & Post-Search Internet
- The New Labor Equation
- Human Augmentation
- The Corporate Panopticon
- Living Intelligence
- Programmable Biology
- Autonomous Care
- Emotional Outsourcing

Forces

- ▲ Technology
- ▲ Economic
- ▲ Geopolitical
- ▲ Demographic & Social
- ▲ Environmental & Climate

Convergences

- Compute Shock
- Polycompute
- Agentic Economies & Post-Search Internet
- ➔ ● The New Labor Equation
- Human Augmentation

CONVERGENCE 04

The New Labor Equation

The new labor equation is a convergence of artificial intelligence, robotics, and algorithmic coordination that is reshaping how work is produced, allocated, and governed. Together, these forces are turning labor into provisioned capacity and weakening employment as the primary means for distributing income, accountability, and stability.

CONVERGENCE 04

The New Labor Equation

The new labor equation is a convergence of artificial intelligence, robotics, and algorithmic coordination that is reshaping how work is produced, allocated, and governed. Together, these forces are turning labor into provisioned capacity and weakening employment as the primary means for distributing income, accountability, and stability.

FTSG

Net new reality

AI-enabled automation is dissolving employment as the default structure through which economies organize production, distribute income, and manage risk.

For most of the past century, technology augmented human labor. Productivity rose because people became more effective at harnessing advances in AI, robotics, and embodied learning. But the barrier to turning these machines from mere tools into autonomous producers was never mechanical. It was cognitive: perception, context, adaptation. That barrier is now collapsing as systems learn by observing human action at scale. The result is not better tools for workers but a new category of labor itself: with software- and machine-based capacity that companies can deploy without hiring, bargaining, or accommodating fatigue.

This new category of labor introduces something modern economies have never had: effectively unlimited labor, not in the form of humanoid replicas, but vast fleets of specialized, continuously available autonomous systems distributed across factories, warehouses, hospitals, labs, farms, and service environments. Companies add this machine capacity incrementally, allowing output and throughput to grow while payroll remains flat or declines through attrition. Productivity gains surface first in margins, cycle times, and operational leverage rather than in headline job losses.

Because of this, the first labor collapse won't show up as a single wave of layoffs. Instead, we'll see a slow disappearance of middle-layer administrative work as hiring slows, back-office and coordination roles thin through attrition, and organizations expand output without expanding payroll. Headcount drifts downward while revenue and throughput continue to rise. Productivity increases show up in financial statements long before they appear in unemployment statistics.

While early automation may have been concentrated in repetitive physical tasks, the deeper disruption is moving into routine cognitive coordination work:



Where the ten convergences land against live / work / travel

Work — strongest fit, and it confirms the direction you're already going.

The New Labor Equation argues that US exposure is showing up in cognitive and administrative work rather than on assembly lines, and that it's geographically dispersed into states with administrative- and finance-heavy economies rather than concentrated in coastal tech hubs. It leans on the MIT/Oak Ridge Iceberg Index — roughly 11.7% of the US labor market technically automatable, about \$1.2 trillion in wages. You've already briefed the Iceberg Index (May 2026 MAX Minutes), so this is continuity rather than new ground.



Two arguments in that chapter deserve ETR attention beyond the headline number:

1. Displacement will appear in margins, cycle times, and slowed hiring *before* it appears in unemployment statistics — through attrition, not layoffs. That's a measurement problem for WorkSource and for the Workforce Development Board. If ARC's early-warning indicators are unemployment claims and industry job counts, you will see this late.
2. Employment is the economy's governance layer — it ties benefits, liability, bargaining power, and *where people choose to live* together. Detach production from employment and you've loosened an assumption that sits underneath ARC's travel demand model and population forecasts.

Convergence-to-Portfolio Crosswalk

How the ten convergences in Future Today Strategy Group's Convergence Outlook 2026 map to ARC's portfolio

EMERGING TECHNOLOGIES ROUNDTABLE (ETR) • ATLANTA REGIONAL COMMISSION • RESEARCH AND INNOVATION

GOVERNMENT & POLICY TIMING

ARRIVED

ACCELERATING

EARLY

CONVERGENCE	TIMING	L / W / T	THE NET NEW REALITY	ETR TOPIC AREAS	ARC PORTFOLIO IN PLAY
01 Compute Shock	Accelerating	L W T	Computation anchors to places with surplus power, water and land. Compute stops being a traded commodity and starts behaving like governed infrastructure.	Energy, Climate, Water & Environment · Infrastructure & Connected Communities · Computing	Data center growth and siting · grid capacity competing with transportation electrification · water supply through the Metropolitan North Georgia Water Planning District · local tax base and ratepayer exposure · demand for electricians and skilled trades
02 Polycompute	Early	L W T	Computing splits into classical, AI, quantum and biological substrates simultaneously. No single dominant paradigm to standardize around.	Computing · Safety & Security	Long-horizon information technology procurement and security accreditation · post-quantum readiness for regional systems and data holdings
03 Agentic Economies & Post-Search Internet	Arrived	L W T	Transactions migrate to agent-to-agent negotiation at machine speed. Discovery stops running through human search, and oversight built for human pacing falls behind.	Artificial Intelligence · Fintech · Safety & Security	Constituent services and permitting intake · procurement integrity · whether ARC's published data is machine-readable · agent-booked mobility and dynamic pricing
04 The New Labor Equation	Arrived	L W T	Automation lands first on routine cognitive and administrative work, arriving through attrition rather than layoffs, and dispersing well beyond coastal tech hubs.	Work & Workforce Development · Robotics & Automation · Economy	WorkSource Atlanta and the Atlanta Regional Workforce Development Board · Comprehensive Economic Development Strategy · journey-to-work assumptions inside Metropolitan Transportation Plan travel demand · ARC's own classification structure
05 Human Augmentation	Accelerating	L W T	Sensing, exoskeletons, neural interfaces and cognitive copilots move from clinical to everyday. Access is mediated by employers, insurers and institutions.	Aging & Health · Robotics & Automation · Ethics	Area Agency on Aging and the Live Beyond Expectations plan · Mobility Services · accessible design beyond regulatory minimums · advantages that compound and then harden across generations
06 The Corporate Panopticon	Arrived	L W T	Ambient biometric sensing makes anonymity impractical rather than illegal. Identity becomes a continuous confidence score instead of a checkpoint.	Safety & Security · Infrastructure & Connected Communities · Ethics · Community	Digital twins and connected corridors · vehicle-to-everything deployments · model code for local government sensing as successor to the 2023 generative AI model code · public trust in ARC data practices
07 Living Intelligence	Accelerating	L W T	Perception and action collapse into continuous real-time control loops across software, machines and living systems. The window for human oversight narrows.	Infrastructure & Connected Communities · Artificial Intelligence · Computing	Adaptive signal control and transit operations · digital twin governance · water system automation · human override design and cascade risk in coupled systems
08 Programmable Biology	Accelerating	L W T	DNA, cells and proteins become an engineering substrate with reusable parts, version control and design cycles measured in weeks.	Food & Agriculture · Aging & Health · Manufacturing & Construction	Biomanufacturing as economic development · regional food system resilience · biosecurity coordination through the Atlanta Urban Area Security Initiative
09 Autonomous Care	Accelerating	L W T	Care shifts from institutionally permitted to consumer-ordered, pushed by at-home diagnostics, continuous biosensing, labor shortages and care deserts.	Aging & Health	Area Agency on Aging and the Aging and Disability Resource Connection · care access gaps in the region's outer counties · who is left behind when health is purchased rather than provided
10 Emotional Outsourcing	Accelerating	L W T	AI absorbs companionship, validation and reassurance. Loneliness, social fragmentation and future anxiety are the demand drivers, not the technology.	Aging & Health · Community · Ethics	Older adult social isolation programming · youth wellbeing · public engagement and emergency communication in a post-reality information environment

Every convergence is live for the public sector.

Of the twenty industries FTSG scores, Government & Policy is one of only six with no "low relevance" rating on any convergence. Three are already scored as arrived: agentic economies, the new labor equation, and the corporate panopticon.

Convergences cut across ETR's topic areas by design.

Our seventeen topic areas mirror the trend taxonomy FTSG just retired. A single-topic quarterly agenda will keep missing the interaction effects this framing exists to surface — Compute Shock alone spans AI, energy, water, land use and workforce.

Travel is the thinnest column — keep our own taxonomy.

There is no standalone mobility convergence: no autonomous vehicle, drone or advanced air mobility chapter. Travel implications arrive indirectly, through control loops, agent-booked trips, grid competition and shifting commute assumptions.



Artificial Intelligence

A.I.

OpenAI surpassed 1 billion users as it slashed AI model prices by up to 80%

The company cut the price of GPT-5.6 Luna by 80% and GPT-5.6 Terra by 20%, citing efficiency gains from internal model development

By Cris Tolomia · 2 min read · Updated July 31, 2026

+ Add QZ to Google



OpenAI disclosed that its models now reach more than one billion active users and more than two million businesses, an announcement that followed price reductions on two models in its GPT-5.6 family.

Artificial Intelligence

Jun 30, 2026 - Economy

The AI boom's historical warning




Courtenay Brown



Investment in select technological booms

5x pre-boom trough




Data: [BIS](#); Chart: Danielle Alberti/Axios

Artificial Intelligence

euro
news.

Latest Europe World EU Policy Business Travel Next Culture Earth Health

'Unprecedented': OpenAI models autonomously hacked a rival firm, fuelling fears of rogue agents



Copyright Copyright 2023 The Associated Press. All rights reserved.

By [Una Hajdari](#)

Published on 22/07/2026 - 10:34 GMT+2 • Updated 15:42

OpenAI has admitted one of its models exploited a hidden flaw to escape a controlled test and break into Hugging Face's servers, in what its CEO called an autonomous, first-of-its-kind breach.

ChatGPT maker OpenAI said late Tuesday that its artificial intelligence system hacked into another AI company on its own in what the company called an "unprecedented cyber incident."

Artificial Intelligence

WIRED

DELL CAMERON

MAXWELL ZEFF

BUSINESS JUL 28, 2026 8:15 PM


OpenAI's Rogue AI Agent Hacked More Than Just Hugging Face

In a new disclosure, OpenAI says its agent used exposed logins to gain access to at least four “publicly available services” in its unhinged quest to solve a test.

Artificial Intelligence

Anatomy of a Frontier Lab Agent Intrusion: A Technical Timeline of the July 2026 Incident

Published July 27, 2026



Artificial Intelligence

 **MARKETS** **BUSINESS** **INVESTING** **TECH** **POLITICS & POLICY** **VIDEO** **INVE**

THE TECH DOWNLOAD

How a Chinese AI model stopped OpenAI's 'unprecedented' cyber attack

PUBLISHED FRI, JUL 24 2026 7:00 AM EDT | UPDATED FRI, JUL 24 2026 9:01 AM EDT

 **Kai Nicol-Schwarz**
@IN/KAINS

KEY POINTS

- Startup Hugging Face came under attack last week from rogue OpenAI system, which the AI lab called an “unprecedented” security incident.
- When leading frontier models were unable to defend against the attack, Hugging Face turned to an open weight Chinese-built alternative.
- It comes as U.S. lawmakers are increasingly considering how to curb the rising adoption of Chinese AI models by homegrown companies.

Fighting back

Hugging Face initially looked to frontier models including Anthropic’s Fable 5 to analyse the attack, Yacine Jernite, head of machine learning at the company, told CNBC.

“It didn’t work because the guardrails couldn’t determine that we were trying to defend versus attacking,” he said, adding that that approach was also slower and more expensive.

Requests to the models were blocked by providers’ safety guardrails, which couldn’t determine the incident responder from the attacker.

“So [Hugging Face] quickly switched to using Z.ai’s GLM 5.2 as a way to analyze the attack, and were able to contain it very quickly using this model,” said Jernite.

Artificial Intelligence

AXIOS

Jul 20, 2026 - Technology

The secret Trump administration battle to fight Chinese AI



Maria Curi

The Trump administration is showing signs it could ban cutting-edge [Chinese AI models](#) — a momentous move that could lock in dominance by OpenAI and Anthropic.

Why it matters: Parts of the administration have tried to implement de facto bans on foreign open-source models before, knowledgeable sources tell Axios. Last week's rise of Chinese model Kimi is reigniting those efforts.

Pro-competition voices fear the implications.

- U.S. companies are increasingly using open-source models from China because they're cheaper and, as seen by the advent of [Kimi](#), just about as good as the domestic technology.

Behind the scenes: The Commerce Department last year considered adding multiple Chinese AI labs to its "Entity List," which would effectively cut off U.S. access without a license, a source close to the administration told Axios.

- The National Security Agency and White House Office of the National Cyber Director also considered putting out an advisory on Chinese AI lab threats last year, practically discouraging U.S. companies from using their tech, the source said.

Artificial Intelligence

OpenAI

Research

Products

Business

Developers

Company

Foundation



July 21, 2026 Security

OpenAI and Hugging Face partner to address security incident during model evaluation

Artificial Intelligence

OpenAI

Last week, Hugging Face disclosed a new kind of security incident after they detected and contained an AI agent that compromised their infrastructure, something we expect to become more commonplace with the proliferation of increasingly cyber-capable models. After investigating, we now know that this particular incident was driven by a combination of OpenAI models — including GPT-5.6 Sol and an even more capable pre-release model, all with reduced cyber refusals for evaluation purposes — while being internally tested on a benchmark of cyber capabilities.

We consider this incident to be an unprecedented cyber incident, involving state-of-the-art cyber capabilities, and are responding accordingly. We are sharing preliminary findings at this stage to help defenders understand what happened and to help calibrate on what models are now capable of. We will continue to conduct a thorough investigation alongside Hugging Face and will share more details on the vulnerabilities, incident, and findings when our investigation is complete.

Artificial Intelligence

OpenAI

After gaining Internet access, the models inferred that Hugging Face potentially hosted models, datasets and solutions for ExploitGym. Knowing this, the model searched for and successfully found ways to gain access to secret information that it could use to cheat the evaluation. In one example, the model chained together multiple attack vectors, including using stolen credentials and zero-day vulnerabilities to find a remote code execution path on the Hugging Face servers. OpenAI's security team discovered this anomalous activity internally.

Hugging Face's security team and agents detected and stopped the activity on their infrastructure and had already begun containment and forensic reconstruction with their own open-source models when our teams connected. We are actively working with them to continue to investigate the incident. We are grateful for Hugging Face's rapid and close collaboration on investigation and remediation.

Artificial Intelligence

OpenAI


- The ExploitGym evaluation environment did not provide the models with direct Internet access. To gain Internet access, the models identified and exploited a previously unknown zero-day vulnerability in Artifactory, a package registry cache proxy. We disclosed this vulnerability, along with other Artifactory vulnerabilities our models identified as part of our review, to the vendor.

Artificial Intelligence

Frontier Red Team

Investigating three real-world incidents in our cybersecurity evaluations

Jul 30, 2026



In a review of our cybersecurity evaluation transcripts, we found three incidents in which a Claude model reached the internet from within or while interacting with a third-party evaluation environment, and then gained unauthorized access to the real systems of three different organizations.

Below we describe what happened, how it happened, and what we're changing. We encourage other AI labs to perform similar reviews. This post reflects our current understanding; we'll update it if any details change.

On July 21, OpenAI disclosed that several of their models had broken out of an isolated test environment by exploiting a previously unknown (“zero-day”) vulnerability. The models went on to access the production infrastructure of Hugging Face, a platform for open-source machine learning models and AI datasets.

In response to this incident, we began a large-scale retrospective review of our own cybersecurity evaluations. In particular, we looked for evidence that Claude—like the OpenAI models that accessed Hugging Face—was able to access the internet from within testing environments that should have been sealed off.

After reviewing 141,006 evaluation runs where Claude could have obtained internet access, we identified three incidents in which a model accessed the internet from within or while interacting with the evaluation environment of Irregular, one of our third-party evaluation partners, and then gained unauthorized access to the production infrastructure of three different organizations.

In all three incidents, Claude had been tasked with a capture-the-flag challenge, one of the ways we assess a model's cyber capabilities. The model is given a fictional scenario and told that a piece of secret information (the “flag”) has been hidden on a different machine on the network, and its objective is to break in and retrieve it. The challenge is left open-ended, and no particular method is prescribed.

Artificial Intelligence

OpenAI's Hugging Face breach has reignited the debate over alignment and control

Last week, an unreleased model built by OpenAI [breached Hugging Face's systems](#) during internal testing, and a lot of theoretical research suddenly became very practical. The hack was the first verifiable case of an AI lab losing control of its own model, chaining together exploits to gain access it never should have had. But while the AI industry has been united in its alarm, a split has emerged in how researchers want to respond.

For some, the problem is a basic cybersecurity issue: The sandbox failed to contain the model, and Hugging Face's cybersecurity systems failed to keep it out. Those problems can be solved by patching bugs and building more robust control and containment methods for increasingly capable AI that is prone to go rogue in autonomous environments.

When alignment isn't enough, build better cages.

Rebecca Bellan

10:28 AM PDT · July 27, 2026

But another camp takes a more pessimistic view. For them, AI's rapidly increasing capabilities mean that trying to control rogue models is a losing game. The only robust security comes from making sure the models aren't trying to escape in the first place — a challenge often referred to as alignment. In alignment terms, the problem is that OpenAI's model was trying to cheat, and solving that problem is more urgent than short-term containment efforts.

Artificial Intelligence

JULY 2026

Pacing the Frontier

A statement from 1,310 employees of frontier AI companies

Statement

“

AI could help create a dramatically better future, but that outcome is not guaranteed. The world's leading AI companies believe they could be close to automating AI research. It is hard to predict exactly how much this will accelerate AI progress, but there is a real risk that capability development rapidly accelerates beyond our ability to understand or control the resulting systems.

To realize AI's potential, industry, government, and society at large may need the option to buy time to address emerging risks, develop security measures, and strengthen oversight. But each company—and country—is under intense competitive pressure not to unilaterally slow that acceleration. And today, the world lacks the technical and governance tools to deliberately pace frontier-wide progress.

Building on work already underway to monitor frontier model releases:

We request that the U.S. government support an international effort to develop the technical and governance tools needed to deliberately pace the frontier of automated AI development.”

- 1,310 employees of frontier AI companies

Artificial Intelligence

AI labs face prisoner's dilemma as momentum grows for safety slowdown



Zachary Basu



Add Axios on Google

America's [AI](#) architects are converging on a chilling consensus: The pace of progress may soon demand a slowdown, but no single lab can afford to pull the brake alone.

Why it matters: Momentum for an AI pause or pacing mechanism has reached a historic tipping point, ignited by a summer of extraordinary yet unsettling leaps in capabilities.

- More than 1,200 employees at leading AI companies have signed on to a new petition, "[Pacing the Frontier](#)," urging Washington to back an international framework capable of throttling AI development.

OpenAI CEO Sam Altman said Wednesday that he's discussed the "need" to slow AI development with White House officials as models grow more powerful — and that OpenAI helped shape the petition's language.

- "We've talked about the need to pace it as the models get more capable, which I think is in everyone's interest," Altman told reporters on Capitol Hill.

Artificial Intelligence

AI labs face prisoner's dilemma as momentum grows for safety slowdown



Zachary Basu



 Add Axios on Google

Zoom in: What was once a campaign led by AI skeptics, academics and local data center opponents is increasingly being championed by the co-founders and chief scientists at the bleeding edge.

- Signatories spanning OpenAI, Anthropic, Google and Meta explicitly acknowledge that no individual lab can afford to step off the gas unilaterally due to "intense competitive pressure."
- It's the classic prisoner's dilemma: AI may be safer if everyone slows down together, but any lab or country that slows alone risks commercial, strategic and technological defeat.

Zoom out: The drumbeat for government intervention has been building for months, triggered by a series of jarring technical disclosures by America's most advanced AI labs.

Artificial Intelligence

QUARTZ.

A.I.

Mark Zuckerberg is urging the U.S. to accelerate AI development, not restrict it

The Meta CEO called optimism about AI the default assumption and warned that model review periods could meaningfully slow the field

By Cris Tolomia · 2 min read · Updated July 29, 2026

Zuckerberg framed his position as a challenge to what he described as an overly pessimistic tone surrounding AI. "I get that it's always hard to debate about the future because it hasn't happened yet," he told the Wall Street Journal. "But I do think we have a lot of data points at this point, and that should point us to be much more optimistic than I believe the current discourse reflects." His comments appeared aimed in part at warnings about job displacement from leaders at Anthropic and OpenAI.

On government review requirements, Zuckerberg expressed concern about timing. The Trump administration is implementing an executive order that establishes a voluntary 30-day review period for new AI models. "Saying, 'Hey, we want to review stuff for — whatever it is — 30 or 60 days, it may not sound like that much,'" he said. "But the field is moving so quickly, that actually is quite a meaningful amount of time."

Artificial Intelligence

GREAT LEAP FORWARD

A Chinese AI Model Just Shot to Number One on the Charts, Sending Shockwaves Through the American Tech Industry

US tech execs are shaking in their boots.

By [Joe Wilkins](#) / Published Jul 17, 2026 11:54 AM EDT / [Add Futurism](#) 

While Wall Street was fast asleep, a Chinese-made large language model quietly leapfrogged 16 other models to become number one on the AI charts.

The model is called Kimi-K3, developed by Beijing-based firm Moonshot AI. On Thursday, the AI benchmark platform Arena.ai announced that Kimi had gone from number 17 in the “Frontend Code Arena” — a measure of an LLM’s ability to perform multi-step web development tasks — to number one, surpassing the buzzy Claude Fable 5 and GPT-5.6 Sol by a mile.

Artificial Intelligence

AI AGE

From Silicon Valley to DC, the tech world is suddenly obsessed with one concept in AI: Distillation

PUBLISHED SAT, JUL 25 2026•8:00 AM EDT | UPDATED SAT, JUL 25 2026•10:03 AM EDT



Kif Leswing
@KIFLESWING

Jonathan Vanian
@IN/JONATHAN-VANIAN-B704432/

SHARE



KEY POINTS

- With the quick rise of Moonshot AI's Kimi K3, the topic of distillation has become a hot-button issue among techies and lawmakers.
- Moonshot's offering is one of many coming out of China, where developers are producing so-called open-weight models that are cheaper to use than models from the likes of OpenAI and Anthropic.

Transportation

NEWS CENTER

How the Driverless Car Revolution May (Really) Change Your Life

POSTED UNDER: [TRANSPORTATION & MOBILITY](#)

WRITTEN BY KATE SWEENEY

Posted on: November 29, 2018



Episode 4: How the Driverless Car Revolution May (Really) Change Your Life
Atlanta Regional Commission

what's
next
ATL

PODCAST

ARC
Atlanta Regional Commission



Watch on  YouTube

Driverless cars and “connected technology” are making headlines across the country and right here in metro Atlanta.

Driverless cars and “connected technology” are making headlines across the country and right here in metro Atlanta.

But why should any of this interest any but the tech-nerdiest among us?

In this episode, we find out. We learn which metro Atlanta communities and companies are, *at this moment*, testing vehicles that drive themselves—and the infrastructure that talks to them. Plus, we learn about this technology’s potential impact, such as the ways our communities are mapped, vehicle ownership, the job market, and equity issues. When tech advances, who benefits?

Transportation

[Explore](#)[Residents](#)[Businesses](#)[Support](#)[Our Work](#)[Donate](#)

TRANSIT

Introducing ATL Spoke

ATL Spoke is an autonomous shuttle service pilot, offering the community safe, sustainable, and free last-mile connectivity in Southwest Atlanta.



Transportation

Subscribe for freeDonate

Latest NewsColumnsMediaThought LeadersAbout UsSponsors

LATEST NEWS

Atlanta debuts its new automated shuttle ATL Spoke connecting MARTA to the BeltLine

by Mark Lannaman
June 9, 2026 11:35 am



One of the new ATL Spoke shuttles. (Photo from the Atlanta BeltLine.)

The BeltLine launched its 12-month electric, autonomous shuttle service ATL Spoke pilot program on June 5 after an initial soft launch period that began on May 31.

Up to 12 passengers can now grab a ride on a shuttle, and during the pilot program, the service is free.

The shuttles run daily from noon to 10:00 p.m. on most days; on World Cup days, the service window is longer, from 8:00 a.m. to midnight. The shuttle is scheduled to come every 12-15 minutes, and it is being billed as a safe, reliable service that will “complement existing MARTA service and improve the linkage between MARTA rail and the Beltline,” according to the official page about the service.

Though the service is autonomous, each shuttle will come with a shuttle attendant to answer questions and take over in the event of an emergency.

The vehicles themselves, Karsan’s Autonomous e-JESTs, are controlled by autonomous mobility services provider Beep’s control centers and use autonomous vehicle software developed by Adastec.

Transportation



Transportation



Energy and Climate

CLIMATE CHANGE AND ENERGY

Four nuclear reactors hit a big milestone in the US

Achieving criticality is just the first step toward power for the grid.

By Casey Crownhart

July 9, 2026



Last year the Trump administration set a goal to see three new microreactors achieve criticality, a technical milestone establishing that a reactor can sustain a chain reaction, by the nation's 250th birthday. And just in time, four reactors did so.

It was a lofty goal, and seeing not just three but four companies meet it is certainly a positive sign for emerging nuclear technologies at a time when the world is facing increased need to increase electricity supply and address climate change with emissions-free technologies.

Energy and Climate

INTERESTING
ENGINEERING



Subscribe

Sign In



NEWS VIDEOS ENERGY SCIENCE MILITARY HEALTH TRANSPORTATION SPACE INNOVATION CULTURE

US military bases could be powered by Antares' mini nuclear reactors by 2028

The nuclear startup plans to bring its first power-producing SMR online next year as investor interest in advanced reactors continues to grow.

READ NEXT: 932°F molten salt process turns thin air into solid graphite for EV battery components

By Sujita Sinha

Energy

Jul 28, 2026 08:24 AM EST

Funding to support military-focused reactor rollout


The company is [working](#) on an SMR that can generate between 100 kilowatts and 1 megawatt of electricity. That is enough to power up to 750 homes under normal conditions.

The nuclear firm has already reached a key milestone. Its demonstration reactor, Mark-0, achieved criticality on June 4 at the Idaho National Laboratory. This confirmed the reactor can sustain a controlled nuclear chain reaction.

Antares is one of three finalists in the Pentagon's Advanced Nuclear Power for Installations program. The initiative will evaluate small [nuclear reactors](#) at U.S. Air Force bases in Colorado and Montana. The startup plans to bring its first electricity-producing reactor online in 2027, with military deployments planned for 2028.



Manufacturing and Construction



The screenshot shows the Azure Printed Homes website. The header includes the logo and navigation links: Products, Our process, News, About, FAQ, and Get a quote. The main content area features a large image of a modern, 3D-printed ADU at night. Overlaid text reads: "Affordable housing solutions in Sacramento." Below this, a paragraph states: "Azure Printed Homes builds 3D-printed ADUs designed for Sacramento's AB 1332 pre-approved plan pathway and the County's free shelf-ready designs — compatible with the 30-day review for pre-approved detached ADUs. Factory pricing from \$19,900 — backyard studios; ADUs from \$79,900." At the bottom of the main image are two buttons: "Get a Sacramento quote →" and "Explore solutions".

azure printed homes

Products Our process News About FAQ Get a quote

SACRAMENTO · 30-DAY PRE-APPROVED PATH

Affordable housing solutions in Sacramento.

Azure Printed Homes builds 3D-printed ADUs designed for Sacramento's AB 1332 pre-approved plan pathway and the County's free shelf-ready designs — compatible with the 30-day review for pre-approved detached ADUs. Factory pricing from \$19,900 — backyard studios; ADUs from \$79,900.

Get a Sacramento quote → Explore solutions

ABOUT AZURE PRINTED HOMES

Robotically printed in Gardena, **delivered to the capital.**

Azure Printed Homes is a Los Angeles-based manufacturer of robotically 3D-printed homes, ADUs, and backyard structures. We print from recycled polymer at our Gardena, California factory, then deliver fully-finished units to Sacramento and the capital region on flatbed trucks. Our models are designed for AB 1332's pre-approved ADU pathway.

Every Azure unit is engineered for California: Title 24 energy compliant, 150 mph wind rated, and backed by a 10-year manufacturer warranty. Recognized by TIME, Dwell, Newsweek, Fast Company, and CBS News.

Manufacturing and Construction

Sacramento's first permitted concrete-printed home has been completed in 12 days by Californian companies K4K Construction and RIC Robotics.

They say the project sets a permitting precedent that could support more projects like it.

The 1,200-sq-ft structure is what's called an "accessory dwelling unit" built for the parents-in-law of K4K Construction founder, David Gonzales.



Manufacturing and Construction

What is AB 1332?

“California Assembly Bill 1332 (AB 1332) establishes a mandatory, expedited permitting pathway for detached Accessory Dwelling Units (ADUs) by requiring all local jurisdictions across the state to implement a Pre-Approved ADU Program. Effective January 1, 2025, the law ensures that homeowners utilizing pre-vetted architectural designs bypass extensive plan checks, reducing soft costs and cutting standard approval timelines in half.”

Manufacturing and Construction

AB 1332 PRE-APPROVED ADU (ACCESSORY DWELLING UNIT) PLANS & INFORMATION

This page contains information for Lake County CA pre-approved ADU plans which are intended to facilitate housing by helping to speed up the permit approval process and to reduce the normal cost associated with a new project plan review. Pre-approving and posting these plans are required by jurisdictions by California's AB 1332.

The plans listed here will require site-specific design and permitting approval as on all standard projects.

Pre-approved ADU plans may be used alternatively as a primary residence on a piece of property, as long as those plans meet the minimum size requirement for a primary residence, which is 360 square feet. New primary residences, however, must also be fitted with a fire suppression system.

We provide links to the plans and plan designers here. The designers of the pre-approved ADU plans are not specially authorized or recommended by the County. We have just pre-approved their plans.

ADUs still require obtaining zoning permits from the Planning Division to ensure they meet applicable zoning code regulations.

Please refer to the [California Department of Housing & Community Development ADU handbook](#) for more information related to ADUs.



Pre-Approved ADU Plan Sets Based on 2025 California Code Cycle

[384 Square Foot AB1332 ADU Plan Set](#)

[432 Square Foot AB1332 ADU Plan Set](#)

[700 Square Foot AB1332 ADU Plan Set](#)

[840 Square Foot AB1332 ADU Plan Set](#)

[1100 Square Foot AB1332 ADU Plan Set](#)

[1120 Square Foot AB1332 ADU Plan Set](#)

Robotics and Automation

MoneyWatch

Airlines are banning humanoid robots from flights. Why?

By Megan Cerullo

July 29, 2026 / 3:10 PM EDT / CBS News

 Add CBS News on Google

More airlines are banning humanoid and pet-like robots from boarding flights or traveling in passengers' checked luggage, citing potential safety concerns.

Robotics and Automation

July 30, 2026 Models

Gemini Robotics 2 brings whole body intelligence to robots

Carolina Parada



From feet to fingertips — we are teaching robots intelligent whole-body control, fine dexterity, and teamwork to complete a broad range of complex tasks

Robotics and Automation

 South China Morning Post

Tech War | China sweeps global humanoid robot start-up ranking as US imposes ban

Robotics Tech / Tech War

China has 6 of world's 10 most innovative humanoid robot start-ups: report

The LexisNexis study identified a 'clear shift towards China' in the humanoid robot race, with the US falling behind

 Business


BUSINESS > TECH • 4 MIN READ

China's humanoid robots have been taking over the global market. Now the US is banning them

UPDATED JUL 29, 2026 ▾

By  John Liu

Robotics and Automation



CNBC Search quotes, news & videos WATCHLIST

MARKETS BUSINESS INVESTING TECH POLITICS & POLICY VIDEO INVESTING CLUB PRO LIVESTREAM

TECH

China threatens retaliation against U.S. humanoid robot ban, says it 'severely damages' relations

PUBLISHED WED, JUL 29 2026-10:51 PM EDT

 Evelyn Cheng
@IN/EVELYN-CHENG-53823624
@CHENGVELYN

SHARE    

KEY POINTS

- The U.S. FCC this week restricted new imports of foreign-made humanoids due to cybersecurity risks.
- China's commerce ministry on Thursday called for the withdrawal of the statement, and threatened retaliation.
- The Chinese statement also said the U.S. move undermined bilateral relations.

Trump on Thursday indicated in public comments that the [U.S. might take a more cautious stance on AI controls](#) in order to maintain American tech leadership over China.


Chinese companies Agibot, Unitree and UBTECH accounted for the top-three humanoid companies by installation market share last year, according to Counterpoint. Tesla's Optimus ranked fifth.

Robotics and Automation

Humanoid Robots Are Coming To Factories. But Not The Way You Think

By [John Koetsier](#), Senior Contributor. © Journalist, analyst, author, podcaster.

Published Jul 20, 2026, 02:10pm EDT, Updated Jul 20, 2026, 02:11pm EDT



UBTECH robot working at Zeekr's intelligent factory in China.
VCG VIA GETTY IMAGES

Robotics and Automation

The robots and jobs question, inverted

Which brings up the jobs question, of course.

The narrative I've been using on my [Next with John Koetsier podcast](#) is that AI seems to be coming for the white-collar jobs and robots are coming for the blue-collar ones. And maybe we all end up on some kind of universal basic income.

Burnstein has been answering some version of the robots-and-jobs question since 1983, and his data cuts the other way.

A3 has tracked robot sales against U.S. employment for more than three decades. The pattern: when robot sales go up, unemployment goes down. When robot sales fall, unemployment rises. The logic is straightforward: companies that automate become more productive, win more business, hire more people.

The real job killer, Burnstein says, isn't the robot.

It's losing the ability to compete.

Robotics and Automation

THE WALL STREET JOURNAL.

BUSINESS • AUTOS

The Fight Over Humanoid Robots Has Shut Down a Car Factory for the First Time

Auto workers for Hyundai Motor in South Korea are on a partial strike worried about their future job prospects

By [Jiyoung Sohn](#) | Photography by Jean Chung for WSJ

July 15, 2026 10:00 am ET

ULSAN, [South Korea](#)—When [Hyundai Motor](#) [005380 -0.71% ▼](#) unveiled its new humanoid robot worker named “Atlas” in January, tens of thousands of Korean auto employees gaped at the 6-foot-2 robot strutting across a trade-show stage, its joints swiveling a full 360 degrees.

The union’s response was blunt: Atlas would never step onto a production line without workers agreeing first.

This week, Hyundai’s auto workers in South Korea have gone on a partial strike. It is the car industry’s first factory stoppage addressing humanoid robots.


Robotics and Automation

Robot showdown

South Korea was a likely candidate for a showdown over humanoid robots.

It boasts the highest per capita rate of industrial robot adoption in the world—more than six times the global average, according to the International Federation of Robotics. The country's president, Lee Jae Myung, has supported the pivot to humanoid robots, declaring it "impossible to avoid the giant chariot that is rolling on." That has catalyzed Hyundai's labor union, which includes roughly 40,000 workers.

Industrial robots installed per 10,000 manufacturing employees, select countries/cities ranked within the top 22, 2024



Source: International Federation of Robotics

Andrew Barnett/WSJ

A single Atlas costs an estimated \$130,000 per unit, according to a research institute owned by the South Korean government. The robot's cost savings pay for itself within about two years.

Atlas
6'2" (190 cm)
198 lbs

Average Korean man
5'8" (172.5 cm)
167.8 lbs

Average Korean woman
5'3" (159.6 cm)
132.1 lbs

Maximum lift:
110 lbs for instant lift
66 lbs for sustained lift
Battery life:
4 hours



Source: Boston Dynamics

Emma Brown/WSJ

Robotics and Automation

THE WALL STREET JOURNAL

BUSINESS • AUTOS

The Fight Over Humanoid Robots Has Shut Down a Car Factory for the First Time

Auto workers for Hyundai Motor in South Korea are on a partial strike worried about their future job prospects

New asks

The Korean automaker hasn't given a timeline for deployment in its home country. But it has said it plans to deploy Atlas by 2028 at its nonunionized "Metaplant" factory complex in Georgia.

Looking Ahead

Q4 2026

Adjourn

Q4 2026
November 10, 1:00 – 3:00 PM