

AI Unveiled: Deep Research on the Most Important Discoveries and News in the World of AI from the Past 7 Days

Date of Report: November 9, 2025

Author: Senior Technology Analyst

1.0 Introduction: Unveiling the New Foundations of AI

1.1 Stating the Theme: AI Unveiled

This report adheres to the "AI Unveiled" theme, moving analysis beyond the iterative updates to existing large language models that have dominated headlines. The period of November 2-9, 2025, was not defined by new models achieving incremental benchmark gains. Instead, it was defined by the public unveiling of the foundational pillars—the strategic infrastructure, novel computational architectures, government-led industrial standards, and critical-sector applications—that will support the next decade of artificial intelligence.

1.2 Executive Summary of Discoveries

The past seven days have seen a correlated series of major announcements that, when analyzed as a whole, form a coherent picture of an industry entering its "industrial build-out" phase.

- We witnessed the unveiling of the *new super-scale infrastructure* for artificial general intelligence (AGI) development, solidified by a landmark \$38 billion multi-cloud

partnership.¹

- We saw the *new hybrid architecture* required to bridge the gap between classical supercomputing and the mercurial power of quantum processors.⁵
- We observed the creation of *new standards and metrics* by government bodies, a move designed to transform "quantum" from a physics buzzword into a defined, auditable engineering discipline.⁶
- We identified a *new critical-sector application* of AI in managing national energy grids, marking a pivotal transition from "soft" digital problems to "hard" physical-industrial challenges.⁹
- Finally, we saw the emergence of a *new user interface paradigm* for mass adoption, as platforms emerged to turn non-programmers into "citizen developers" of AI.¹⁰

1.3 Why These Discoveries Matter

These events are not isolated. They are deeply interconnected, revealing the full "stack" of the next computational era. This report will analyze how these five discoveries function as the simultaneous unveiling of the **Infrastructure Layer**, the **Architecture Layer**, the **Standards Layer**, the **Application Layer**, and the **Interface Layer** of a post-LLM world. This is the strategic "plumbing" for the next generation of AI being laid in real-time.

1.4 Methodology Note

This report adheres to a strict methodology. All findings are sourced from the past seven days and have been corroborated by multiple, independent, credible global sources, including official announcements, respected technology publications, and government research agencies.

2.0 Key Discoveries of the Week (November 2-9, 2025)

2.1 Discovery: The \$38B Deal Unveiling a Multi-Cloud AGI Future

2.1.1 The Announcement

In the first week of November 2025, multiple credible financial and technology news outlets confirmed that OpenAI has signed a landmark \$38 billion, multi-year partnership with Amazon Web Services (AWS).¹

2.1.2 Description and Context

This agreement is not a standard enterprise cloud contract. The deal provides OpenAI with massive, priority access to AWS cloud computing infrastructure, including, critically, access to hundreds of thousands of state-of-the-art NVIDIA GPUs.³ This move is strategically significant as it breaks OpenAI's long-standing *perceived* exclusivity with Microsoft Azure, its primary partner and investor.¹²

For Amazon, this partnership is being reported as a strategic "revival" for its AWS division, which some investors had perceived as lagging behind Microsoft in securing high-profile generative AI partnerships.³ The announcement had an immediate market impact, with news of the deal helping to push Amazon's market capitalization past \$2 trillion for the first time.⁴

2.1.3 Potential Impact and Strategic Implications

The single most important implication of this deal is that AGI development is now officially a multi-cloud endeavor. This \$38 billion commitment unveils a strategic de-risking by OpenAI. It ensures that its development of next-generation models, such as GPT-5 and beyond, is not dependent on a single infrastructure provider, mitigating a critical supply-chain bottleneck.

This deal also confirms that the primary competitive battleground for AI has shifted, at least temporarily, from model performance benchmarks to raw infrastructure access. The limiting factor for AGI is no longer just algorithmic innovation, but secure, long-term access to hundreds of thousands of interconnected, state-of-the-art accelerators. This move by AWS and OpenAI will almost certainly trigger responsive, large-scale infrastructure deals from competitors like Google and Anthropic, who cannot afford to be locked out of future GPU

supply.

2.1.4 Unveiling the "Capital-for-Capacity" Model

A \$38 billion figure ¹ is far too large to be a standard, pay-as-you-go cloud services contract, even for a compute-intensive client like OpenAI. This suggests the deal is a "capital-for-capacity" agreement. OpenAI is effectively pre-paying \$38 billion to *reserve* a significant portion of AWS's future GPU-based infrastructure, securing its supply chain for computation for the rest of the decade.

This action unveils a massive new barrier to entry for AGI development. If building a frontier model now requires \$30B+ *per provider* in infrastructure pre-commitments, the field is now definitively restricted to a handful of trillion-dollar corporations and their well-funded proxies. This commoditizes the cloud providers, who become "arms dealers" of computation, and solidifies the entrenched power of the handful of AI labs capable of making such investments.

2.2 Discovery: The Quantum-Classical Hybrid Architecture at Oak Ridge

2.2.1 The Announcement

Multiple technology sources reported this week on a new collaboration at Oak Ridge National Laboratory (ORNL).⁵ The project brings together ORNL, NVIDIA, and HPE to build a new breed of quantum-classical hybrid computing system. The system will feature the NVIDIA NVQLink system, built by HPE, and is set for installation in early 2026.⁵

2.2.2 Description and Context: The "Mercurial Qubit" Problem

"Pure" quantum computing has been perpetually hamstrung by the fundamental physical limitations of its core components. Qubits, the quantum equivalent of bits, are described as

"mercurial" and "exquisitely sensitive to every micro-disturbance".⁵ The slightest noise, flicker of magnetism, or sliver of heat can cause them to decohere and "collapse" their delicate quantum state, creating high error rates that render meaningful computation impossible.

2.2.3 Emerging Technology Deep-Dive: The Hybrid Solution

This new system *unveils* a new, pragmatic architectural paradigm. Instead of waiting to perfect the qubit (a physics problem), this architecture uses classical supercomputing (a solvable engineering problem) to manage the quantum system in real-time.

It leverages the "brute force of GPU and CPU clusters to supplement and manage quantum error correction and control".⁵ The lynchpin of this new technology is the software layer. The NVIDIA CUDA-Q programming environment is described as a "shared musical score" that allows researchers, for the first time, to "orchestrate both quantum and GPU simulations simultaneously".⁵

2.2.4 Potential Impact and Strategic Implications

This hybrid approach unveils the *first credible, engineering-led path* to commercially viable quantum computing. It reframes quantum from a theoretical physics experiment into a practical High-Performance Computing (HPC) accelerator, managed by a classical system.

Crucially, the system is designed to be "not locked to one specific protocol or quantum hardware vendor".⁵ This positions NVIDIA's CUDA-Q as the *essential software gateway* to the entire quantum industry. This strategy mirrors how NVIDIA's original CUDA platform became the indispensable standard for classical AI, locking developers into its ecosystem. This new architecture is designed to train AI models for error correction, emulate quantum noise, and solve complex optimization and simulation problems that neither paradigm can solve alone.⁵

2.2.5 NVIDIA's Pragmatic Pivot

This announcement represents a pragmatic pivot by the HPC industry, effectively acknowledging that the "pure quantum" dream is, for the near term, commercially unviable. The future is hybrid. This hybrid architecture can be seen as a "Trojan horse" for NVIDIA. It

ensures that even in a quantum-powered future, the most critical (and profitable) component will still be an NVIDIA-controlled GPU cluster running NVIDIA's proprietary software stack. This strategy allows NVIDIA to own the most valuable real estate—the software integration layer—while the quantum hardware (the QPUs) from various vendors risk becoming commoditized peripherals managed by the classical GPU.

2.3 Discovery: DARPA Unveils the Roadmap for "Utility-Scale" Quantum

2.3.1 The Announcement

On November 6, 2025, the U.S. Defense Advanced Research Projects Agency (DARPA) announced it had selected several companies, including IonQ and QuEra Computing, to advance to Stage B of its Quantum Benchmarking Initiative (QBI).⁶

2.3.2 Description and Context: A Market Without Metrics

To date, the quantum industry has operated as a "Wild West" of competing, unverified claims. Companies use different and often incompatible metrics (qubit count, "quantum volume," etc.), making it impossible for government, scientific, or commercial customers to assess performance or value.

2.3.3 The New "Utility-Scale" Paradigm

DARPA's QBI is the first serious, government-backed effort to *unveil a standard definition* of quantum utility. The initiative's goal is to "rigorously verify and validate" if any quantum approach can achieve "utility-scale operation." DARPA provides a clear and commercially relevant definition for this: "meaning its computational value exceeds its cost".⁶ Stage B, which has just been announced, involves creating a detailed R&D roadmap *to the year 2033* to

achieve this goal.⁶

2.3.4 Potential Impact and Strategic Implications

This announcement signals the moment "quantum computing" begins its transition from a marketing buzzword to a structured, auditable engineering discipline. DARPA is, in effect, *creating the metrics* that will define the market for the next decade.

For the selected companies like IonQ and QuEra, this is a massive stamp of "institutional validation".¹⁵ It separates their roadmaps from the general hype and signals to the market that their technical approaches are considered plausible by U.S. government experts. This initiative, and IonQ's stated goal to "support both our national security and national economic security"⁶, unveil the true driver: quantum is now a matter of state-level strategic competition.

2.3.5 The Coordinated Quantum Blueprint

This week's two major quantum announcements are not a coincidence. They are two sides of the same coin, revealing the "how" and the "what" of the U.S. quantum strategy.

DARPA's QBI sets the *demand signal* (the "what"): achieve "utility-scale" computing, defined as "value > cost," by 2033.⁸ The ORNL/NVIDIA hybrid system provides the *supply-side answer* (the "how"): a specific architecture designed to extract value from today's noisy, imperfect quantum systems by pairing them with cost-effective classical GPUs.⁵ The hybrid path is the most likely, and perhaps only, path to meet DARPA's "value > cost" definition within that timeframe. We are witnessing the simultaneous, coordinated unveiling of the new industrial-quantum paradigm. DARPA is creating the benchmarks, while the national labs and their partners (ORNL, NVIDIA) are unveiling the hybrid architecture to meet them.

2.4 Discovery: Hitachi Unveils AI for Critical National Infrastructure

2.4.1 The Announcement

In the past week, Hitachi announced a new strategic partnership with the Southwest Power Pool (SPP) to use AI for modernizing the U.S. energy grid.⁹ This announcement complements a broader Memorandum of Understanding (MoU) Hitachi signed with the U.S. Department of Commerce to promote investment in energy infrastructure.⁹

2.4.2 Description and Context: A New Application Vertical

This is a novel and highly significant deployment of AI. The Southwest Power Pool manages the high-voltage electrical grid for a 14-state region in the central U.S. This is not an enterprise IT project; it is the application of advanced AI to *critical national infrastructure* (CNI).

2.4.3 Industry Application Deep-Dive: The "Green Grid" Problem

The partnership's stated goal is to assist with "grid upgrades and renewable energy integration".⁹ This addresses a critical, unsolved challenge for modern economies. Renewable energy sources like wind and solar are *variable* and *intermittent*, which creates massive instability for a power grid designed for the stable, on-demand power of fossil fuels.

This implies Hitachi's AI is a new *control-plane technology*—a sophisticated forecasting, load-balancing, and optimization system that can manage these rapid fluctuations in real-time and ensure grid stability.

2.4.4 Potential Impact and Strategic Implications

This application unveils AI's transition from "soft" digital problems (like search, ad targeting, and text generation) to "hard" physical-industrial problems (like managing power flow, energy physics, and grid stability).

It serves as a blueprint for AI in other "hard" industries like logistics, manufacturing, and public utilities, where safety, reliability, and security are paramount. Furthermore, this unveils AI as a critical *enabling technology* for the green energy transition. The stated goals of

decarbonization are likely unachievable without AI-based control systems to manage the inherent instability of a renewable-heavy grid.

2.4.5 A New Class of Physical-World Challenge

This development is the first concrete example of AI being deployed to solve the *second-order consequences* of another technology (renewable energy). However, this creates a new and massive attack surface. By embedding AI at the heart of the U.S. energy grid⁹, this partnership also unveils a new, high-stakes national security vulnerability. An AI-controlled grid is a high-value target for adversarial attacks. The challenges of AI safety, robustness, and security are no longer academic; they are matters of CNI resilience.

2.5 Discovery: Google Unveils the "Citizen Developer" AI Platform

2.5.1 The Announcement

On November 6, 2025, Google Labs announced the massive global expansion of its "Opal" platform, a no-code AI app builder. The platform's availability was expanded from an initial 15 to over 160 countries.¹⁰

2.5.2 Description and Context: Beyond the Chatbot

AI interaction has so far been dominated by the chat interface (prompting), while AI *development* has been restricted to engineers and data scientists. Opal seeks to bridge this gap.

2.5.3 Emerging Technology Deep-Dive: The No-Code Paradigm

Opal is a platform that allows "non-programmers" or "citizen developers" to build and deploy their own custom AI applications.¹⁰ The key "unveiled" concept is the *workflow*. Users can visually build complex, multi-step "mini-apps"¹⁰ without writing any code. Examples include custom marketing asset generators, research automation tools, or language learning apps.¹⁰

2.5.4 Potential Impact and Strategic Implications

This is a paradigm shift. Google is not just giving people access to AI (like its Gemini model); it is giving them the tools to *build with* AI. This "democratization of AI development"¹⁰ could unleash a new wave of innovation, similar to what the Apple App Store did for mobile applications.

This is a strategic move by Google to *own the platform layer* for the next generation of AI-native applications. By providing the tools, Google aims to create a massive, proprietary ecosystem of AI apps built on its platform. This also unveils the next evolution of the AI user interface—moving from the *conversational* (the chatbot) to the *visual and structural* (the workflow builder).

2.5.5 The War for the Ecosystem

This is Google's strategic counter-move to the "prompt engineering" phenomenon. It is a bet that the future of scalable AI value is not in clever *prompts* but in shareable, productized *workflows*. A simple chat interface is not "sticky." A visual, no-code platform where users build their own intellectual property (the workflow) *is*.

This is Google's attempt to build the "Android" for AI. While OpenAI and Microsoft may own the core "model" (akin to an OS kernel), Google is trying to build the dominant "app store" (Google Play) and "development platform" (Android Studio) for a new generation of "citizen developers." This is a battle for the ecosystem.

3.0 Analysis of Emerging Technologies: The Hybrid & No-Code Paradigms

3.1 Synthesis of the Hybrid Paradigm

The past week's announcements clarify the true future of high-performance computing. The ORNL/NVIDIA system ⁵ is the physical manifestation of this new hybrid paradigm. The "mercurial" nature of qubits ⁵ is the central, unsolved *problem*. The use of classical GPUs for real-time error correction and control ⁵ is the *solution*. The DARPA QBI ⁸ provides the *commercial framework* that will measure the success of this solution, with its pragmatic "value exceeds cost" metric.

This combination of events reveals that the future of HPC is not "pure quantum" replacing "classical." It is a *permanent, fused, classical-assisted architecture*. This is a fundamental re-definition of the quantum computing roadmap, moving it from physics to engineering.

3.2 Synthesis of the No-Code Paradigm

Google's global expansion of Opal ¹⁰ is the vanguard of a new AI interface paradigm. The *problem* it solves is that AI's immense power is currently locked behind two barriers: the "high wall" of programming expertise and the "narrow keyhole" of the chat prompt. The *solution* is a visual, workflow-based, no-code platform. This signals the start of a new war for platform dominance. The race is no longer just to build the biggest model, but to build the *stickiest ecosystem* and *most accessible interface* for the "citizen developer".¹⁰ This is the "App Store" moment for generative AI.

4.0 Survey of New Industry Applications: From Grids to the Cloud

4.1 AI in Critical, Regulated Sectors

The Hitachi/Southwest Power Pool partnership ⁹ serves as a powerful case study for AI's jump

from the digital world to the physical, industrial world. This is the "AI as Critical Enabler" thesis. The goals of the green energy transition—specifically, the integration of unstable renewable energy sources—are likely physically impossible to achieve at scale *without* an AI-based control plane to manage grid instability.⁹ This elevates AI from an IT tool to a foundational, C-suite, and national-level strategic asset.

4.2 Infrastructure as a Strategic Battleground

The cloud computing market itself is the first industry being fundamentally reshaped by AI. The \$38 billion AWS/OpenAI deal¹ demonstrates that the most significant "application" of AI right now is its own *supply chain*. The multi-billion dollar expenditures are not for a finished AI *product*, but for the *foundational capacity to build* future products. This reveals a market that is consolidating at the infrastructure level *before* it has even fully matured at the application level.

4.3 The Unveiled AI Ecosystem Stack

The discoveries of the past week are not disconnected. They form the five distinct layers of a new, emerging technology stack.

Table 1: The Unveiled AI Ecosystem: Connecting This Week's Foundational Discoveries

Discovery (Date)	Key Organizations	Technology Layer Unveiled	Key Technology/Paradigm	Core Strategic Implication
OpenAI/AWS Deal (Nov 2-7)	OpenAI, Amazon (AWS)	Infrastructure Layer	Multi-Cloud AGI	AGI development requires multi-vendor, \$30B+ infrastructure commitments, consolidating

				the market. ¹
Hybrid System (Nov 5)	Oak Ridge, NVIDIA, HPE	Architecture Layer	Quantum-Classical Hybrid (CUDA-Q)	The path to quantum value is not "pure" but a hybrid, GPU-controlled architecture. ⁵
DARPA QBI (Nov 6)	DARPA, IonQ, QuEra	Standards Layer	"Utility-Scale" Benchmarking	The "Wild West" era of quantum hype is over; a new, metrics-based industrial discipline is being created. ⁶
Hitachi/Grid (Nov 5-7)	Hitachi, Southwest Power Pool	Application Layer	Industrial AI Control Plane	AI moves from "soft" digital to "hard" physical (CNI) problems, enabling the green transition. ⁹
Google Opal (Nov 6)	Google	Interface Layer	No-Code "Citizen Developer"	The AI user interface is evolving from "prompt" to "visual workflow," creating a new app-ecosystem battleground. ¹⁰

5.0 Challenges and Strategic Considerations

5.1 The Physical Bottleneck (Quantum)

As the ORNL research implies, the "mercurial"⁵ and "exquisitely sensitive"⁵ nature of qubits remains the fundamental, unsolved physics problem. The hybrid system is an *engineering patch*, not a *physics solution*. The challenge is that the industry is now building a sophisticated, software-defined ecosystem (CUDA-Q) on top of an unstable, unsolved physics experiment.

5.2 The Security of AI in CNI (Industrial AI)

The Hitachi/grid partnership⁹ unveils a massive new challenge. By making AI the central control plane for the U.S. energy grid, it creates what may be the single most high-value cybersecurity target in the nation. The challenge is that AI "safety" and "robustness" are no longer academic. A single adversarial attack, a data-poisoning incident, or a critical model hallucination could have catastrophic, physical-world consequences.

5.3 The Consolidation of Power (Infrastructure)

The \$38 billion AWS/OpenAI deal¹ is not a story of innovation; it is a story of *consolidation*. The challenge is that AGI development is now clearly restricted to a handful of trillion-dollar tech titans (Amazon, Google, Microsoft) and their proxies (OpenAI, Anthropic). This level of capital requirement³ creates insurmountable barriers to entry, stifles open competition, and concentrates the power to build the world's most powerful technology in the hands of three or four corporations.

5.4 The Ecosystem "Lock-in" (Interface)

The Google Opal¹⁰ and NVIDIA CUDA-Q⁵ announcements reveal the true business model: *platform lock-in*. The challenge for customers and developers is that these "vendor-agnostic"⁵ or "no-code"¹⁰ platforms are gateways to highly proprietary ecosystems. By building on Opal, "citizen developers" are building on and for the Google Cloud ecosystem. By developing on CUDA-Q, researchers are locking themselves into the NVIDIA hardware and software stack.

6.0 Strategic Outlook and Near-Future Directions

6.1 Near-Term Trends (Next 3-6 Months)

- **Competitor Response to Infrastructure:** Expect a major counter-announcement from Microsoft, re-asserting its relationship with OpenAI and its own infrastructure prowess. Expect Google and/or Oracle to announce a new, massive partnership with a major AI lab (such as Anthropic) to counter the AWS deal.³ The "infrastructure wars" are now in full swing.
- **The "Citizen Developer" Race:** Expect Microsoft to launch a "PowerAutomate for AI" and Amazon a "SageMaker for no-code," directly competing with Google's Opal.¹⁰ The battle for the AI *interface layer* will be the most active new market of 2026.

6.2 Mid-Term Trends (Next 1-3 Years)

- **The Rise of the "CUDA-Q" Ecosystem:** The NVIDIA hybrid architecture⁵ will become the *de facto* standard for quantum-classical HPC. Universities and research labs will rapidly adopt it, creating a new generation of researchers fluent in CUDA-Q, cementing NVIDIA's dominance long-term.
- **AI in CNI Becomes a Vertical:** The Hitachi/grid deal⁹ will serve as the pathfinder. Expect to see "AI for Logistics," "AI for Water Utilities," and "AI for Smart Cities" emerge as high-value, defensible enterprise verticals.

6.3 Concluding Insight: The "Industrial Build-Out"

The "AI Unveiled" theme of the past week has been consistent. We are witnessing the end of the purely theoretical, "model-magic" phase of AI and the start of the "industrial build-out." The key discoveries were not new models, but the foundational infrastructure ⁴, architectures ⁵, standards ⁸, applications ⁹, and interfaces ¹⁰ that will define the market for the next decade. The "unveiling" shows a new, full-stack ecosystem being built in real-time, from the silicon and qubits up to the citizen developer's mouse-click. This is the strategic landscape on which the next battles for AI dominance will be fought.

Works cited

1. AI News Tracker - Spreaker, accessed November 9, 2025, <https://www.spreaker.com/podcast/ai-news-tracker--6204412>
2. Amazon Issues Cease-and-Desist to Perplexity AI Over Shopping Tool - Alpha Spread, accessed November 9, 2025, <https://www.alphaspread.com/market-news/regulatory-actions/amazon-issues-cease-and-desist-to-perplexity-ai-over-shopping-tool>
3. All News for AMZN : Amazon.com - Zacks.com, accessed November 9, 2025, <https://www.zacks.com/stock/research/AMZN/all-news>
4. The AI Industry's Critical Transition: Landmark Deals, Surging Valuations, and Enterprise Adoption - Spreaker, accessed November 9, 2025, <https://www.spreaker.com/episode/the-ai-industry-s-critical-transition-landmark-deals-surging-valuations-and-enterprise-adoption--68459655>
5. Quantum Computing 101 - Spreaker, accessed November 9, 2025, <https://www.spreaker.com/podcast/quantum-computing-101--6443895>
6. IonQ Advances to Stage B of DARPA's Quantum Benchmarking ..., accessed November 9, 2025, https://markets.ft.com/data/announce/detail?dockey=600-202511061745BIZWIRE_USPRX_20251106_BW234988-1
7. DARPA Selects QuEra for Stage B of Quantum Benchmarking Initiative (QBI) - ADVFN, accessed November 9, 2025, <https://www.advfn.com/stock-market/stock-news/97178383/darpa-selects-quera-for-stage-b-of-quantum-benchma>
8. QBI: Quantum Benchmarking Initiative - DARPA, accessed November 9, 2025, <https://www.darpa.mil/research/programs/quantum-benchmarking-initiative>
9. Hitachi (TSE:6501) Deepens AI and Energy Ties—Does This Signal a New Strategic Focus?, accessed November 9, 2025, <https://www.webull.com/news/13784370542601216>
10. Google's Opal AI App Builder Goes Global in 160 ... - The Tech Buzz, accessed November 9, 2025, <https://www.techbuzz.ai/articles/google-s-opal-ai-app-builder-goes-global-in-160-countries>

11. Artificial Intelligence Archives - TechRepublic, accessed November 9, 2025, <https://www.techrepublic.com/topic/artificial-intelligence/>
12. TECHi: The operating system for understanding technology., accessed November 9, 2025, <https://www.techi.com/>
13. Adam Spatacco Articles | The Motley Fool, accessed November 9, 2025, https://www.fool.com/author/20479/?utm_source=msnrss&utm_medium=feed&utm_campaign=article&referring_guid=5f91eece-65ac-486a-a101-f31e0d9fb819
14. IonQ Co-Founds Q-Alliance to Create Lombardy Quantum Hub | IONQ Stock News, accessed November 9, 2025, <https://www.stocktitan.net/news/IONQ/ion-q-partners-with-italy-s-national-strategy-for-quantum-technology-linrupf2qou3.html>
15. Final thoughts before earnings - IonQ (Nasdaq: IONQ) Live: Full Day Earnings Coverage | 24/7 Wall St., accessed November 9, 2025, https://247wallst.com/ionq-nasdaq-ionq-live-full-day-earnings-coverage/liveupdates/5/?tpid=liveblog-1499897-6&tv=link&tc=in_content
16. 3 things to watch after the closing bell - IonQ (Nasdaq: IONQ) Live: Full Day Earnings Coverage - 247WallSt, accessed November 9, 2025, https://247wallst.com/ionq-nasdaq-ionq-live-full-day-earnings-coverage/liveupdates/1/?tpid=liveblog-1499897-7&tv=link&tc=in_content
17. AI - PPC Land, accessed November 9, 2025, <https://ppc.land/tag/ai/>