October 27, 2025

AMD Powers U.S. Sovereign AI Factory Supercomputers, Accelerating an Open American AI Stack

News Highlights

- The Lux AI supercomputer, powered by AMD and housed at Oak Ridge National Labs, will be the first dedicated U.S. AI factory for science and will be deployed in early 2026.
- AMD will power the next-generation supercomputer at Oak Ridge National Labs, called Discovery – which will advance U.S. AI and scientific research at massive scale.
- Together, Lux and Discovery, directly support the U.S. AI Action Plan by accelerating AI-enabled science, strengthening national competitiveness, and enabling secure, sovereign AI infrastructure for the nation.

SANTA CLARA, Calif., Oct. 27, 2025 (GLOBE NEWSWIRE) -- AMD (NASDAQ: AMD) and the U.S. Department of Energy (DOE) today announced two next-generation systems at Oak Ridge National Laboratory (ORNL) designed to expand America's leadership in artificial intelligence (AI) and high-performance computing (HPC), the Lux AI supercomputer and the Discovery supercomputer. Discovery and Lux will be DOE flagship supercomputers designed to drive breakthroughs in science, energy, and national security. Both systems directly support the U.S. AI Action Plan by accelerating AI-enabled science, strengthening national competitiveness, and advancing secure, sovereign AI infrastructure for the nation.

Together with the U.S. government, ORNL and industry partners, AMD will deliver advanced computing platforms that help national researchers, agencies, and innovators tackle pressing challenges in energy, medicine, health, and national security. When fully deployed, the Lux and Discovery systems will represent a combined $1billion investment of private and public funding which will enable the DOE to build a secure, federated and standards-based infrastructure for sovereign U.S. AI and science.

"We are proud and honored to partner with the Department of Energy and Oak Ridge National Laboratory to accelerate America's foundation for science and innovation," said Dr. Lisa Su, chair and CEO, AMD. "Discovery and Lux will leverage AMD's high-performance and AI computing technologies to advance the most critical U.S. research priorities in science, energy, and medicine – demonstrating the power of public-private partnership at their best."

The Lux AI supercomputer – co-developed by ORNL, AMD, Oracle Cloud Infrastructure (OCI) and HPE and powered by AMD Instinct™ MI355X GPUs, AMD EPYC™ CPUs and AMD Pensando™ advanced networking technologies – will be deployed in early 2026, making it the first US AI Factory supercomputer setting the stage for novel public and private partnerships. The Lux system will address the immediate needs to expand the DOE's AI

leadership and accelerate progress and innovation in the areas of AI, energy research, materials, medicine and advanced manufacturing.

The Discovery supercomputer deepens the collaboration between the DOE, ORNL, HPE and AMD. At the heart of Discovery are next-gen AMD EPYC CPUs, codenamed “Venice,” and AMD Instinct™ MI430X GPUs— a new MI400 Series accelerator engineered specifically for sovereign AI and scientific computing. MI430X extends the leadership AMD architecture for AI and HPC, enabling the U.S. to train, simulate, and deploy AI models on domestically built systems that protect national data and scientific competitiveness.

“Winning the AI race requires new and creative partnerships that will bring together the brightest minds and industries American technology and science has to offer,” said U.S. Secretary of Energy Chris Wright. “That’s why the Trump administration is announcing the first example of a new commonsense approach to computing partnerships with Lux. We are also announcing, as part of a competitive procurement process, Discovery. Working with AMD and HPE, we’re bringing new capacity online faster than ever before, turning shared innovation into national strength, and proving that America leads when private-public partners build together.”

“The Discovery system will drive scientific innovation faster and farther than ever before,” said ORNL Director Stephen Streiffer. “ORNL’s leadership in supercomputing has dramatically shortened researchers’ time from problem to solution across a host of fields and industries. With Discovery, the integration of high-performance computing and AI promises breakthroughs at the accelerated speed and scale necessary for continued U.S. leadership in an increasingly competitive global environment.”

“We are proud that the partnership between HPE and AMD is at the forefront of advancing the next generation of supercomputing in the AI era for Oak Ridge National Laboratory,” said Antonio Neri, president and CEO at HPE. “HPE’s newest supercomputing solutions harness converged AI and HPC architectures, enabling the lab to achieve unprecedented productivity and scale. This collaboration also reinforces American leadership in applying AI to science, energy and national security.”

“Oracle is honored to work alongside the Department of Energy to help drive breakthroughs in science, energy, and national security,” said Mahesh Thiagarajan, executive vice president, Oracle Cloud Infrastructure. “Oracle will deliver sovereign, high-performance AI infrastructure that will support the co-development of the Lux AI cluster.”

Lux AI Supercomputer: Advancing U.S. AI Leadership

Lux at ORNL is the nation’s first dedicated AI Factory for science, energy, and national security—purpose-built to train, fine-tune, and deploy AI foundation models that will accelerate discovery and engineering innovation. Lux is designed to accelerate AI-driven science through its advanced architecture, optimized for data-intensive and model-centric workloads.

As a critical resource for the DOE and the National Laboratory system, Lux will give the United States an early and decisive advantage in deploying frontier AI capabilities that transform the pace and scale of scientific innovation. By dramatically expanding the nation’s capacity for scientific AI, Lux strengthens U.S. leadership, accelerates AI-for-Science, and helps ensure that the breakthroughs of the AI era are discovered, developed, and deployed in America.

Lux will unleash a new generation of scientific capabilities across the DOE's mission to, reveal nature's hidden structure, discover breakthrough materials, accelerate energy innovation, transform biology and biosecurity, and advance national security and resilience.

Discovery: America's Next Leap in Scientific Computing

Purpose-built for AI and science, Discovery features a "Bandwidth Everywhere" design that allows science and AI applications to run efficiently and deliver productive results. Importantly, Frontier laid the foundation for this work, and together with HPE Cray Supercomputing GX5000 – its next-generation supercomputing platform – Discovery builds on that success with a consistent programming environment to ensure an easy transition for users.

Discovery will be the DOE's next flagship supercomputer at ORNL, arriving in 2028, and will extend U.S. leadership in HPC and AI. Advancing America's innovation agenda, Discovery reflects a strong, public-private commitment to keep the United States at the forefront of AI and scientific innovation.

Building on the breakthroughs of Frontier—the world's first exascale system—Discovery is designed to deliver dramatically higher performance, energy efficiency, and AI capabilities. Discovery will drive breakthroughs in energy, biology, advanced materials, national security, and manufacturing innovation. It will help design next-generation reactors, batteries, catalysts, semiconductors, and critical materials.

Key Points of Discovery:

- **Bandwidth Everywhere:** Impressive memory capacity, node, and global network bandwidth compared to first-generation exascale machines, accelerating science and AI applications.
- **Seamless User Transition:** Applications built for Frontier transition seamlessly as Discovery continues the successful programming environment.
- **More Compute, Similar Energy:** More AI and scientific output at comparable power costs.
- **Open Standards:** Discovery is built on open-source software, open standards, and open science to provide the foundation for AI sovereignty.

ORNL is expected to take delivery of Discovery in 2028, with user operations expected in 2029. Once online, Discovery will serve as a cornerstone of the American AI Stack, advancing the science, security, and innovation goals that define the nation's next decade.

About AMD
For more than 55 years AMD has driven innovation in high-performance computing, graphics and visualization technologies. Billions of people, leading Fortune 500 businesses and cutting-edge scientific research institutions around the world rely on AMD technology daily to improve how they live, work and play. AMD employees are focused on building leadership high-performance and adaptive products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow, visit the AMD (NASDAQ: AMD) website, blog, LinkedIn, Facebook and X pages.

Cautionary Statement
This press release contains forward-looking statements concerning Advanced Micro

Devices, Inc. (AMD) such as features, functionality, performance, availability, timing and expected benefits of AMD products; deployment of AMD technology to power Lux AI supercomputer and timing and expected benefits thereof; and deployment of AMD Instinct™ MI430X accelerators to power the Discovery supercomputer and timing and expected benefits thereof, which are made pursuant to the Safe Harbor provisions of the Private Securities Litigation Reform Act of 1995. Forward-looking statements are commonly identified by words such as "would," "may," "expects," "believes," "plans," "intends," "projects" and other terms with similar meaning. Investors are cautioned that the forward-looking statements in this press release are based on current beliefs, assumptions and expectations, speak only as of the date of this press release and involve risks and uncertainties that could cause actual results to differ materially from current expectations. Such statements are subject to certain known and unknown risks and uncertainties, many of which are difficult to predict and generally beyond AMD's control, that could cause actual results and other future events to differ materially from those expressed in, or implied or projected by, the forward-looking information and statements. Material factors that could cause actual results to differ materially from current expectations include, without limitation, the following: competitive markets in which AMD's products are sold; the cyclical nature of the semiconductor industry; market conditions of the industries in which AMD products are sold; AMD's ability to introduce products on a timely basis with expected features and performance levels; loss of a significant customer; economic and market uncertainty; quarterly and seasonal sales patterns; AMD's ability to adequately protect its technology or other intellectual property; unfavorable currency exchange rate fluctuations; ability of third party manufacturers to manufacture AMD's products on a timely basis in sufficient quantities and using competitive technologies; availability of essential equipment, materials, substrates or manufacturing processes; ability to achieve expected manufacturing yields for AMD's products; AMD's ability to generate revenue from its semi-custom SoC products; potential security vulnerabilities; potential security incidents including IT outages, data loss, data breaches and cyberattacks; uncertainties involving the ordering and shipment of AMD's products; AMD's reliance on third-party intellectual property to design and introduce new products; AMD's reliance on third-party companies for design, manufacture and supply of motherboards, software, memory and other computer platform components; AMD's reliance on Microsoft and other software vendors' support to design and develop software to run on AMD's products; AMD's reliance on third-party distributors and add-in-board partners; impact of modification or interruption of AMD's internal business processes and information systems; compatibility of AMD's products with some or all industry-standard software and hardware; costs related to defective products; efficiency of AMD's supply chain; AMD's ability to rely on third party supply-chain logistics functions; AMD's ability to effectively control sales of its products on the gray market; impact of climate change on AMD's business; impact of government actions and regulations such as export regulations, tariffs and trade protection measures, and licensing requirements; AMD's ability to realize its deferred tax assets; potential tax liabilities; current and future claims and litigation; impact of environmental laws, conflict minerals related provisions and other laws or regulations; evolving expectations from governments, investors, customers and other stakeholders regarding corporate responsibility matters; issues related to the responsible use of AI; restrictions imposed by agreements governing AMD's notes, the guarantees of Xilinx's notes, the revolving credit agreement and the ZT Systems credit agreement; impact of acquisitions, joint ventures and/or strategic investments on AMD's business and AMD's ability to integrate acquired businesses, including ZT Systems; AMD's ability to complete the sale of ZT Systems' manufacturing business; impact of any impairment of the combined company's assets; political, legal and economic risks and natural disasters; future impairments of technology license purchases; AMD's ability to attract and retain qualified personnel; and AMD's stock price volatility. Investors are urged to review in detail the risks and uncertainties in AMD's

Securities and Exchange Commission filings, including but not limited to AMD's most recent reports on Forms 10-K and 10-Q.

Oracle Future Product Disclaimer
The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Oracle Forward-Looking Statements Disclaimer
Statements in this article relating to Oracle's future plans, expectations, beliefs, and intentions are "forward-looking statements" and are subject to material risks and uncertainties. Many factors could affect Oracle's current expectations and actual results, and could cause actual results to differ materially. A discussion of such factors and other risks that affect Oracle's business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including Oracle's most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at oracle.com/investor. All information in this article is current as of October 27, 2025 and Oracle undertakes no duty to update any statement in light of new information or future events.

Contact:
Aaron Grabein
AMD Communications
+1 512-602-8950
aaron.grabein@amd.com

Liz Stine
AMD Investor Relations
+1 720-652-3965
liz.stine@amd.com

Source: Advanced Micro Devices, Inc.