



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Stargate Project: OpenAI's \$500B AI Data Center Plan

Build practical AI for pharma and biotech with IntuitionLabs. We help life-science teams turn complex information and workflows into useful software, governed knowledge systems and AI tools.

Published by IntuitionLabs · 2025-10-08 · openai stargate · ai infrastructure · hyperscale datacenter · ai compute · nvidia gpus · oracle cloud · sovereign ai · sam altman · ai

Original article: <https://intuitionlabs.ai/articles/openai-stargate-datacenter-details>



In this article

1. [Stargate Project: OpenAI's \\$500B AI Data Center Plan](#)
 2. [U.S. Data Center Rollout](#)
 3. [Abilene, Texas: Stargate's Flagship Site](#)
 4. [Expanding Beyond Abilene](#)
 5. [International Expansion \("OpenAI for Countries"\)](#)
 6. [Technology and Capacity](#)
 7. [Economic Impact and Infrastructure](#)
 8. [Criticism and Challenges](#)
 9. [Status as of February 2026 and Outlook](#)
-

Stargate Project: OpenAI's \$500B AI Data Center Plan

In January 2025 [OpenAI](#) unveiled “**The Stargate Project**,” a massive AI infrastructure venture to build world-scale data centers. The new entity (funded by OpenAI, SoftBank, Oracle, and Abu Dhabi’s MGX) will invest an **initial \$100 billion, ramping up to \$500 billion** over the next four years to deploy **AI compute** in the U.S. ([openai.com](#)) ([www.datacenterdynamics.com](#)). SoftBank (CEO Masayoshi Son as chairman) handles the financing, while OpenAI manages operations ([openai.com](#)) ([techcrunch.com](#)). In December 2025, SoftBank completed its investment in OpenAI: \$22.5 billion on top of \$7.5 billion invested in April 2025, which together with \$11 billion from co-investors fully funded a **\$41 billion** commitment and left SoftBank with a stake of roughly 11% ([group.softbank](#)) ([cnbc.com](#)). This is framed as a public–private effort to “secure American leadership in AI,” creating jobs and strategic capability ([openai.com](#)) ([www.datacenterdynamics.com](#)). Key technology partners include Microsoft, NVIDIA, Oracle, Arm, and others ([www.datacenterdynamics.com](#)) ([www.datacenterdynamics.com](#)). (OpenAI will continue using Microsoft Azure alongside Stargate’s on-premises infrastructure ([www.datacenterdynamics.com](#)).) The project grew out of a White House event on January 21, 2025 ([techcrunch.com](#)) ([www.techtarget.com](#)), and aims to meet explosive AI compute demand with a **10-gigawatt (GW)** target (roughly 10 GW is enough to power ~7.5 million homes).

Update (October 2026): Stargate is still active, but its shape has changed since early 2026. In Abilene, Oracle and OpenAI dropped a planned expansion beyond the campus’s 1.2 GW in March 2026 ([tomshardware.com](#)); Epoch AI’s satellite tracker (updated October 5, 2026) estimates that four of the eight buildings were operating by mid-2026 and that the other four, already structurally complete, will come online around November 2026 ([epoch.ai](#)). The other six U.S. campuses (Shackelford County and Milam County in Texas, Doña Ana County in New Mexico, Port Washington in Wisconsin, Saline Township in Michigan, and Lordstown in Ohio) were in development but not yet operating at Epoch’s April 2026 review, which projects most of them to be finished by late 2028 ([epoch.ai](#)). In Europe, OpenAI paused Stargate UK in April 2026, citing energy costs and regulation ([cityam.com](#)), and Microsoft, not OpenAI, is now the customer for Nscale’s 230 MW Narvik campus in Norway ([nscale.com](#)) ([networkworld.com](#)). On funding, NVIDIA’s September 2025 letter of intent to invest up to \$100 billion gave way to a \$30 billion NVIDIA investment in OpenAI’s \$110 billion round of February 2026, which

also included \$30 billion from SoftBank and \$50 billion from Amazon ([openai.com](#)) ([business-standard.com](#)). OpenAI's first custom inference chip, Jalapeño, co-developed with Broadcom, published its first benchmark results in August 2026 ([openai.com](#)).

U.S. Data Center Rollout

Abilene, Texas: Stargate's Flagship Site

The first Stargate campus is in **Abilene, Texas**, a consortium-built facility on a vast site. Construction began almost immediately after the announcement. Oracle is overseeing Abilene, which is being developed on leased land with *Crusoe Energy* reportedly involved. Published reports say the Abilene buildout could reach nearly **1 GW** of power by mid-2026, costing around **\$3.4 billion** ([techcrunch.com](#)). At a White House briefing Oracle's Larry Ellison noted that "*10 [half-million-square-foot] buildings*" are under construction here ([www.datacenterdynamics.com](#)) ([techcrunch.com](#)). By September 2025 the Abilene campus had **6,400+ construction workers** on site ([tech.yahoo.com](#)). Sam Altman says portions of Abilene are "**up and running**", with the first two buildings operational since September 2025 and NVIDIA GB200 AI racks already delivered for early **model training** ([openai.com](#)). Oracle will ultimately deploy over **450,000 GB200 GPUs** at the Abilene campus, with the remaining six buildings expected to be completed by mid-2026 ([datacenterdynamics.com](#)), and leases the computing power to OpenAI under a 15-year agreement. The full Abilene campus will use **1.2 GW** of power, enough for roughly 900,000 homes ([networkworld.com](#)). As of October 2026, Epoch AI estimates that four buildings are operating and the remaining four will be online around November 2026 ([epoch.ai](#)).

Expanding Beyond Abilene

Stargate quickly expanded to multiple new U.S. campuses. By mid-2025 OpenAI announced a partnership with Oracle to add **4.5 GW more** of U.S. data center capacity ([openai.com](#)). Together with Abilene's 1 GW, this puts over **5 GW** under development, well on track toward the 10 GW goal. In September 2025, OpenAI, Oracle and SoftBank **announced five more U.S. sites**. Three are Oracle-led campuses (in Shackelford County, Texas; Doña Ana County, New Mexico; and a Midwest site later named as Port Washington, Wisconsin, developed with Vantage) and two are SoftBank-led (Lordstown, Ohio, and Milam County, Texas) ([openai.com](#)) ([techcrunch.com](#)). These five new campuses, together with Abilene and CoreWeave partnerships, push total planned capacity to about **7 GW** ([openai.com](#)) ([techcrunch.com](#)), which OpenAI said put it on a clear path to securing the full **10 GW, \$500B commitment by end-2025**, ahead of schedule ([openai.com](#)) ([www.theregister.com](#)).

- **Oracle sites (total ~5.5 GW):** The three Oracle-built campuses (initially announced July 2025) plus a potential 600 MW expansion near Abilene were expected to deliver about 5.5 GW combined ([openai.com](#)) ([www.datacenterdynamics.com](#)). The Abilene expansion was dropped in March 2026 ([tomshardware.com](#)). (The Shackelford County site, for example, is a ~1.4 GW hyperscale campus

being developed by Vantage and Oracle (www.datacenterdynamics.com.) Oracle is also contracted to supply OpenAI with cloud services (~\$300B over 5 years) (www.datacenterdynamics.com) and is delivering NVIDIA GPU racks to Abilene (openai.com) (openai.com).

- **SoftBank sites (total ~1.5 GW):** Two new campuses are via SoftBank: one in Lordstown, Ohio, and one in Milam County, Texas (openai.com). Together they can scale to ~1.5 GW in 18 months (openai.com) (www.datacenterdynamics.com). The Lordstown facility uses SoftBank's advanced data center design and was described in September 2025 as on track to be operational in 2026, while the Milam site is developed with SB Energy (SoftBank's power subsidiary) providing accelerated power infrastructure (openai.com) (www.datacenterdynamics.com). Epoch AI's April 2026 satellite review found no large-scale data center construction visible at Lordstown yet (epoch.ai).

In October 2025, OpenAI announced another major campus: **Stargate Michigan, a more than 1 GW, \$7 billion** facility in Saline Township developed by Related Digital and Oracle. Dubbed "The Barn" (for a historic red barn preserved at the entrance), it will consist of three 550,000-sq-ft buildings on 250 acres, with construction beginning in early 2026 and over 2,500 union construction jobs (openai.com) (techrepublic.com). Michigan Governor Whitmer called it "the largest economic project in the state's history."

In January 2026, OpenAI and SoftBank each invested **\$500 million in SB Energy** (SoftBank's power subsidiary), totaling a \$1 billion commitment to accelerate data center development. SB Energy was selected to build and operate the **1.2 GW Milam County, Texas** campus, with initial facilities expected to enter service in 2026 (cnbc.com) (openai.com).

In all, these U.S. Stargate centers are each designed for enormous scale (multi-hundred-megawatt buildings). For perspective, Larry Ellison said each building is about **500,000 square feet**, and ten such buildings are already being built (techcrunch.com). Adding the Michigan campus, the five September 2025 sites (Shackelford, Doña Ana, Lordstown, Milam, and Port Washington, Wisconsin), and Abilene brought planned capacity to over **8 GW** and more than \$450 billion in investment over three years, according to OpenAI (openai.com). OpenAI states this will create *tens of thousands of jobs* (25,000 onsite in the new sites alone) (openai.com) (www.theregister.com), though independent sources caution that long-term employment is much lower (Bloomberg reported the Abilene site might yield only ~57 permanent jobs (www.cnbc.com)).

International Expansion (“OpenAI for Countries”)

Alongside U.S. buildout, OpenAI has launched **Stargate projects abroad** under its “OpenAI for Countries” initiative (openai.com) (openai.com). The goal is to partner with friendly governments and firms to build local AI data centers that mirror Stargate's capabilities.

- **UAE (May 2025):** In late May 2025 OpenAI announced *Stargate UAE*, a 1 GW data center venture in the United Arab Emirates (www.capacitymedia.com). Partners include AI firm G42, Oracle, Nvidia, Cisco, and SoftBank, with UAE investors also funding the project. Initially 200 MW will be deployed by 2026 (eventually up to 1 GW) (www.capacitymedia.com). OpenAI said Stargate UAE has the potential to provide compute capacity within a 2,000-mile radius, “reaching up to half the world's population”

(www.capacitymedia.com). The deal benefited from U.S. export-rule changes to the UAE and was publicly lauded by officials (OpenAI thanked President Trump for support in the Middle East) (www.capacitymedia.com).

- **Norway (July 2025):** In July 2025 OpenAI announced *Stargate Norway*, its first European AI data center initiative (openai.com). The site is in Narvik, leveraging abundant hydropower. Initial capacity is **230 MW** (planned 100,000 NVIDIA GPUs by end-2026), with room to expand an extra 290 MW (openai.com). This facility will run entirely on renewables with state-of-the-art liquid cooling, and even supply waste heat to local industry (openai.com). The joint venture project (with local partners Nscale and Aker) aims to feed compute into Norway's and Northern Europe's AI ecosystem (openai.com) (openai.com). In April 2026, Nscale announced that Microsoft would take the Narvik campus's capacity, and Network World reported that OpenAI had pulled out of talks for the site (nscale.com) (networkworld.com).
- **United Kingdom (Sept 2025):** In September 2025 OpenAI unveiled *Stargate UK*, a partnership with NVIDIA and UK firm Nscale (openai.com). It will expand Nscale's planned sites (e.g. Cobalt Park) to provide "sovereign" UK compute. OpenAI planned to explore offtake of up to 8,000 GPUs in Q1 2026, with the potential to scale to 31,000 GPUs over time (openai.com). The aim is onshore specialist AI workloads (public services, finance, etc.) so that "world-leading AI models can run on local computing power" in the UK (openai.com) (openai.com). Sam Altman emphasized this builds on a July 2025 U.S.–UK collaboration pact and is part of broader efforts to share AI across partners (openai.com) (openai.com). In April 2026, OpenAI paused *Stargate UK*, saying it would move forward "when the right conditions such as regulation and the cost of energy enable long-term infrastructure investment" (cityam.com).
- **Argentina (Oct 2025):** OpenAI and Sur Energy signed a Letter of Intent for *Stargate Argentina*, a large-scale data center in Patagonia capable of reaching **500 MW**, representing an estimated investment of up to **\$25 billion**. Reports cited by DCD put the first 100 MW phase in 2027 (datacenterdynamics.com) (openai.com). In May 2026, the Buenos Aires Times reported that the project had not yet been presented to the Argentine government (batimes.com.ar).
- **South Korea (Oct 2025):** OpenAI struck deals with Samsung Electronics and SK Hynix for high-bandwidth memory (HBM) chip supply, while SK Telecom signed an MoU to develop, build, and operate AI data centers in South Korea as part of the *Stargate* ecosystem. SK Telecom's MoU covers an AI data center in the southwest region of Korea (theaiinsider.tech) (eng.sk.com).
- **Other Countries:** OpenAI has said it is exploring up to 10 international projects under the "OpenAI for Countries" initiative (openai.com). OpenAI has also submitted expressions of interest to join consortiums under the EU's "AI Gigafactories" initiative, and Fortune has described the program as a way for OpenAI to build global soft power (openai.com) (fortune.com).

Overall, *Stargate*'s global network is designed to deliver "**democratic AI**" infrastructure: data-sovereign clouds in allied nations that follow shared principles, rather than authoritarian alternatives (openai.com) (openai.com).

Technology and Capacity

The Stargate centers will be stuffed with cutting-edge AI hardware. OpenAI and partners are locking in massive GPU supply. **NVIDIA** signed on heavily: in September 2025 the two companies signed a letter of intent for NVIDIA to invest up to **\$100 billion** in OpenAI progressively as at least 10 GW of NVIDIA systems are deployed ([techcrunch.com](https://www.techcrunch.com)) (www.theregister.com). That letter of intent never progressed to a final agreement; in February 2026 NVIDIA instead invested \$30 billion in OpenAI's equity round, and OpenAI said it would use 3 GW of dedicated inference capacity and 2 GW of training on NVIDIA Vera Rubin systems (openai.com) ([business-standard.com](https://www.business-standard.com)). Likewise, AMD agreed in Oct 2025 to provide GPUs (announced at 6 GW of capacity for future deployment). Beyond third-party GPUs, OpenAI has developed its **own custom AI inference chip** with Broadcom, manufactured by TSMC, an effort TechCrunch flagged when Stargate launched ([techcrunch.com](https://www.techcrunch.com)). In October 2024, DCD reported that OpenAI's chip team, led by former Google TPU head Richard Ho, had about 20 people and hoped to have its first chips manufactured by 2026 (datacenterdynamics.com). The chip, named **Jalapeño**, was unveiled in June 2026, and in August 2026 OpenAI reported that it delivered 1.5 to 1.9 times more AI work per watt at peak throughput than leading commercially available systems on three tested models (openai.com) ([tomshardware.com](https://www.tomshardware.com)). OpenAI and Broadcom plan to deploy 10 GW of these custom accelerators, reducing OpenAI's long-term reliance on NVIDIA (investors.broadcom.com). Under the September 2025 NVIDIA letter of intent, the first gigawatt of NVIDIA systems was planned for the second half of 2026 on NVIDIA's next-generation **Vera Rubin** platform, succeeding the Blackwell GB200 racks (nvidianews.nvidia.com).

Each Stargate campus is envisioned as a **hyperscale AI factory**. OpenAI's NVIDIA partnership alone was described as deploying *millions* of GPUs (www.theregister.com). The combined sites will require immense power delivery. SoftBank's SB Energy and local utilities are bringing on accelerated power (e.g. new substation in Milam TX) to meet the gigawatt-scale loads. Energy sources vary: Epoch AI found that at least three of the seven U.S. sites will use on-site natural gas plants (epoch.ai). The Norway project was designed to run on renewable power (openai.com), while the UAE announcement did not specify energy sources (www.capacitymedia.com). Cooling will use techniques like closed-loop liquid cooling to keep GPUs at optimal temperatures (openai.com).

Partnerships extend beyond chips. For instance, Arm (UK chip designer owned by SoftBank) will provide CPU designs in NVIDIA's next-gen chips (openai.com). CoreWeave (a GPU-focused cloud provider) has a multi-billion-dollar deal with OpenAI to supply additional cloud-based GPU capacity to complement the physical datacenters. Microsoft remains a cloud partner as well, with Azure usage continuing alongside Stargate (www.datacenterdynamics.com) (openai.com).

Economic Impact and Infrastructure

OpenAI and its partners tout “**hundreds of thousands of jobs**” and massive economic benefit from the build-out (openai.com) (openai.com). OpenAI estimates that the 4.5 GW Oracle expansion alone will create **over 100,000 construction and operations jobs** (openai.com) (www.theregister.com). (For instance, OpenAI says

the five new U.S. sites alone will generate ~25,000 onsite jobs (www.theregister.com.) Work includes electricians, technicians, manufacturing and service roles across all 50 states. Significant investments in power plants, transmission lines, and water systems are also anticipated to support these centers.

However, some experts caution that long-term job numbers may be much lower. Large automated data centers often need only dozens of permanent staff once built. Bloomberg reported that Abilene's first facility may only employ ~57 ongoing workers, despite promises of thousands (www.cnbc.com). Critics also note the **enormous energy and water demands** of AI facilities (www.cnbc.com). Data centers consume ~2% of U.S. electricity, and cooling can use hundreds of thousands to millions of gallons of water per day (www.cnbc.com). Local communities will need upgrades to power grids and possibly new generation capacity. Regulators and utilities are closely involved to manage this strain.

Nonetheless, many analysts argue the U.S. needs such "AI megafactories" to stay competitive. Past projects (like nationwide broadband and highway construction) took decades and similar sums, and yielded vast productivity gains. OpenAI frames the buildout as a way to "accelerate America's reindustrialization" and advance U.S. AI leadership (openai.com) (www.datacenterdynamics.com).

Criticism and Challenges

Despite broad support from tech and government leaders, Stargate has faced skeptics and hurdles. Most notably, Elon Musk (who has his own AI company xAI and was then advising the U.S. administration) publicly **lambasted the plan as "fake"**, claiming SoftBank hadn't secured the money to fund it (www.forbes.com) (www.techtarget.com). Musk tweeted that "they don't actually have the money" and that SoftBank had "well under \$10B secured" (www.forbes.com) (www.techtarget.com). OpenAI's Altman quickly refuted this on social media, acknowledging that "what is great for the country isn't always what's optimal for your companies" and urging Musk to "put 🇺🇸 first" (www.forbes.com). The feud highlighted tensions (Musk and Altman have been public adversaries) but did not derail the project.

Another issue is **cost escalation**. In mid-2025 trade tariffs began hiking the price of imported servers, chips, and cooling equipment. Bloomberg reporting (cited by TechCrunch) warned that new U.S. tariffs could add **5–15%** to data center build costs (techcrunch.com). This, combined with tech market volatility and AI "bubbles," made some investors wary. By May 2025, reports emerged that SoftBank had not finalized majority funding and was still courting banks and asset managers (techcrunch.com). OpenAI has stated it will cover the initial \$100B through equity and debt. The full funding of SoftBank's \$41 billion OpenAI investment commitment in December 2025, and the joint \$1 billion SB Energy commitment in January 2026, have substantially addressed earlier funding skepticism (cnbc.com) (cnbc.com). However, the full \$500 billion commitment still requires ongoing fundraising from banks and asset managers over the coming years (www.datacenterdynamics.com).

While critics initially noted a **"slow start"** in 2025, the project has gained significant momentum: two Abilene buildings went operational in September 2025, and multiple new campuses are now under active construction (www.theregister.com). Some observers still wonder whether the full \$500B will be deployed or if projections will grow further. Altman himself has acknowledged the difficulty of building U.S. infrastructure, and has called for easing bureaucratic hurdles (techcrunch.com).

Finally, **national-security and economic concerns** are debated. An analyst noted that the underlying training data and control of AI models matter: if government-sponsored servers run U.S.-centric AI, that is seen as a benefit; but if foreign governments could leverage this infrastructure, questions arise (www.techtarget.com). OpenAI stresses sovereign compute (especially in “OpenAI for Countries” projects) and has engaged allied governments directly to assuage these worries (openai.com) (openai.com).

Status as of February 2026 and Outlook

As of early 2026, Stargate has moved from vision to reality at remarkable scale. The **Abilene** flagship campus has two buildings operational since September 2025, with the remaining six expected to complete by mid-2026, ultimately housing over 450,000 NVIDIA GB200 GPUs (datacenterdynamics.com). Update (October 2026): Epoch AI estimates four Abilene buildings are operating, with the other four expected online around November 2026 (epoch.ai). Six additional U.S. campuses are in various stages of development (Shackelford County TX, Doña Ana County NM, Lordstown OH, Milam County TX, a Midwest site, and Saline Township MI), with land secured and permits progressing (openai.com) (openai.com). SoftBank's **\$41 billion** OpenAI investment commitment (including \$11 billion from co-investors) is fully funded, and the two companies jointly invested **\$1 billion in SB Energy** to accelerate power infrastructure for the buildout (cnbc.com) (cnbc.com). OpenAI said in September 2025 that planned investment had passed **\$400 billion** over three years, and in October 2025 that it exceeded \$450 billion (openai.com) (openai.com).

Internationally, the picture has expanded significantly. Four overseas Stargate projects were announced in 2025 (UAE, Norway, UK, and Argentina), with South Korea joining as a hardware supply and data center partner. By April 2026, OpenAI had paused Stargate UK and stepped back from the Norway site, which Microsoft now leases (cityam.com) (networkworld.com). OpenAI's stated goal remains 10 projects with countries or regions in the first phase of OpenAI for Countries (openai.com) (fortune.com).

On the technology front, Broadcom has said it is targeting the start of deployments of racks containing OpenAI-designed AI accelerators and Broadcom networking systems for the second half of 2026, with completion targeted by the end of 2029 (investors.broadcom.com). OpenAI published the first benchmark results for its Broadcom co-developed Jalapeño chip in August 2026 (openai.com).

Going forward, OpenAI will continue selecting further U.S. site locations from hundreds of proposals (openai.com), and its global “10-country” rollout continues, despite the setbacks in the UK and Norway. Supply-chain deals (with NVIDIA, AMD, Samsung, SK Hynix, and others) continue to expand. Analysts will watch project financing, regulatory approvals, and power delivery: any delays in energy permits or grid upgrades could slow progress.

In summary, as of October 2026, **Stargate** is a sprawling, multi-continental initiative advancing OpenAI's original four-year commitment to invest **\$500 billion** in **10 GW** of U.S. AI infrastructure. It is led by OpenAI with SoftBank and Oracle, leveraging a wide circle of technology partners from chipmakers to energy developers. The technical details (sites, power, GPUs, custom silicon) are being fleshed out in real time, as are the economics (jobs, local benefits vs resource costs). With Abilene partly running and six more U.S. campuses in

development, Stargate is set to be one of the largest infrastructure projects undertaken in the tech sector ([openai.com](#)) ([openai.com](#)), positioning its backers, and arguably the U.S., at the forefront of next-generation AI.

Sources: Official OpenAI announcements and leading news reports (OpenAI blog, TechCrunch, CNBC, DCD, The Register, Broadcom, etc.) provide the above details ([openai.com](#)) ([openai.com](#)) ([openai.com](#)) ([openai.com](#)) ([openai.com](#)) ([www.theregister.com](#)) ([www.capacitymedia.com](#)) ([openai.com](#)) ([openai.com](#)) ([investors.broadcom.com](#)). These cover project scope, funding, partners, locations, technology, and critical commentary.

Source links

Links cited in this article, in order of first appearance. Full addresses are provided for readers using extracted text.

1. <https://intuitionlabs.ai/articles/openai-devday-2025-announcements>
2. <https://intuitionlabs.ai/articles/openai-amd-ai-hardware-partnership>
3. <https://openai.com/index/announcing-the-stargate-project/>
4. <https://www.datacenterdynamics.com/en/news/openai-announces-the-stargate-project-500bn-over-four-years-on-ai-infrastructure/?form=MG0AV3>
5. <https://techcrunch.com/2025/01/21/openai-teams-up-with-softbank-and-oracle-on-50b-data-center-project/>
6. <https://group.softbank/en/news/press/20251231>
7. <https://www.cnbc.com/2025/12/30/softbank-openai-investment.html>
8. <https://www.techtarget.com/searchenterpriseai/news/366618318/Mixed-views-on-Stargate-AI-infrastructure-plan-with-OpenAI>
9. <https://www.tomshardware.com/tech-industry/oracle-and-openai-scrap-planned-600mw-abilene-expansion>
10. <https://epoch.ai/data/ai-data-centers/directory/openai-stargate-abilene>
11. <https://epoch.ai/publications/openai-stargate-where-the-us-sites-stand>
12. <https://www.cityam.com/openai-pauses-stargate-uk-project-over-energy-costs-and-regulation/>
13. <https://www.nscale.com/press-releases/nscale-microsoft-norway>
14. <https://www.networkworld.com/article/4159269/openai-pulls-out-of-a-second-stargate-data-center-deal.html>
15. <https://openai.com/index/scaling-ai-for-everyone/>
16. https://www.business-standard.com/technology/tech-news/nvidia-30bn-openai-investment-after-shelving-100bn-deal-126022000503_1.html
17. <https://openai.com/index/jalapeno-first-results/>
18. <https://tech.yahoo.com/ai/articles/openai-ceo-sam-altman-officials-173850853.html>
19. <https://intuitionlabs.ai/articles/key-innovations-behind-chatgpt>
20. <https://openai.com/index/stargate-advances-with-partnership-with-oracle/>
21. <https://www.datacenterdynamics.com/en/news/openai-and-oracle-to-deploy-450000-gb200-gpus-at-stargate-abilene-data-center/>
22. <https://www.networkworld.com/article/3995015/oracle-to-spend-40b-on-nvidia-chips-for-openai-data-center-in-texas.html>
23. <https://openai.com/index/five-new-stargate-sites/>
24. <https://techcrunch.com/2025/09/23/openai-is-building-five-new-stargate-data-centers-with-oracle-and-softbank/>

25. https://www.theregister.com/2025/09/24/openai_oracle_softbank_datacenters/
26. <https://www.datacenterdynamics.com/en/news/openai-announces-five-more-us-stargate-data-centers-with-oracle-and-softbank/>
27. <https://openai.com/index/expanding-stargate-to-michigan/>
28. <https://www.techrepublic.com/article/news-openai-michigan-stargate-data-center-7b/>
29. <https://www.cnbc.com/2026/01/09/openai-and-softbank-group-announce-1-billion-investment-in-sb-energy-.html>
30. <https://openai.com/index/stargate-sb-energy-partnership/>
31. <https://www.cnbc.com/2025/02/06/openai-looking-at-16-states-for-data-center-campuses-tied-to-stargate.html>
32. <https://openai.com/global-affairs/openai-for-countries/>
33. <https://openai.com/index/introducing-stargate-norway/>
34. <https://www.capacitymedia.com/article/openai-to-open-first-international-stargate-data-centre-in-uae>
35. <https://openai.com/index/introducing-stargate-uk/>
36. <https://www.datacenterdynamics.com/en/news/openai-plans-500mw-data-center-in-argentina/>
37. <https://openai.com/global-affairs/argentinai-ai-opportunity/>
38. <https://www.batimes.com.ar/news/economy/chau-gpt-no-trace-of-mega-billion-dollar-ai-project-announced-before-elections.phtml>
39. <https://theaiinsider.tech/2025/10/01/openai-strikes-landmark-chip-and-data-center-deals-in-south-korea-to-power-stargate-project/>
40. <https://eng.sk.com/news/sk-group-partners-with-openai-to-advance-global-ai-infrastructure>
41. <https://fortune.com/article/openai-stargate-for-countries-soft-global-power/>
42. <https://www.datacenterdynamics.com/en/news/openai-building-first-custom-ai-inference-chip-with-tsmc-and-broadcom-report/>
43. <https://www.tomshardware.com/tech-industry/semiconductors/openai-says-its-jalapeno-chip-beats-nvidias-gb300-in-first-published-benchmarks>
44. <https://investors.broadcom.com/news-releases/news-release-details/openai-and-broadcom-announce-strategic-collaboration-deploy-10>
45. <https://nvidianews.nvidia.com/news/openai-and-nvidia-announce-strategic-partnership-to-deploy-10gw-of-nvidia-systems>
46. <https://www.forbes.com/sites/dereksaul/2025/01/22/musk-is-bashing-trumps-new-stargate-ai-initiative-with-openai-and-oracle-heres-why/>
47. <https://techcrunch.com/2025/05/12/openais-stargate-project-reportedly-struggling-to-get-off-the-ground-thanks-to-tariffs/>

THE PUBLISHER

About IntuitionLabs

IntuitionLabs is an AI consulting, custom software development and data engineering firm serving pharmaceutical, biotechnology, medical-device and other life-science organizations. We work with clinical, regulatory, medical-affairs, commercial, quality and IT teams to connect technology decisions with the work people need to accomplish.

AI consulting and adoption

Our AI enablement services cover readiness assessments, use-case selection, governance and policies, team workshops, adoption measurement and ongoing advisory support. We help organizations structure the information layer behind AI: source material, context, permissions and maintained knowledge that make generated answers useful and reviewable. Private LLM inference and hosted AI options support teams evaluating how to operate AI with appropriate control over their data and infrastructure.

Software, data and life-science workflows

IntuitionLabs develops custom software for pharma and biotech, integrates enterprise systems, and builds data engineering and business intelligence solutions. Areas of focus include AI agents, regulatory research, medical writing, medical affairs, CMC information, competitive intelligence and clinical-document workflows. Our eTMF intelligence work includes cross-system reconciliation and inspection-readiness support.

Enterprise platforms and regulated delivery

We provide Veeva services, application support, managed services, integrations and custom applications, alongside enterprise content work involving platforms such as Egnyte. For regulated workflows, our services include GxP enablement, computer-system validation and software development addressing 21 CFR Part 11 requirements. The applicable controls, validation responsibilities and acceptance criteria are defined for each engagement.

Work with IntuitionLabs

Explore [AI enablement](#), [pharma and biotech software development](#), [data engineering and BI](#), and [Veeva services](#). [Contact IntuitionLabs](#) to discuss your workflow, information sources and implementation needs.

IntuitionLabs publishes educational research to help life-science teams make informed technology decisions. Coverage of a product or organization does not imply a client relationship, endorsement or partnership.

Website: <https://intuitionlabs.ai>

This article is educational. Verify consequential medical, legal, financial, regulatory and technical decisions with appropriate professionals and the cited primary sources. Company services are described separately from the article's research findings.