

OpenAI's Stargate Project: A Guide to the AI Infrastructure

By IntuitionLabs.ai • 10/8/2025 • 15 min read

[openai stargate](#)[ai infrastructure](#)[hyperscale datacenter](#)[ai compute](#)[nvidia gpus](#)[oracle cloud](#)[sovereign ai](#)[sam altman](#)



The Stargate Project: An Overview

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](#)). 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).

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 on the order of **\$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 late 2025 the Abilene campus already employed **6,400+ construction workers** ([tech.yahoo.com](#)). Sam Altman says portions of Abilene are "**up and running**", with NVIDIA GB200 AI racks already delivered and being used for early **model training** ([openai.com](#)).

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 site in the U.S. Midwest) 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), on course to hit 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 600 MW expansion at Abilene are expected to deliver about 5.5 GW combined (openai.com) (www.datacenterdynamics.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 should be online soon, while the Milam site is developed with SB Energy (SoftBank's power subsidiary) providing accelerated power infrastructure (openai.com) (www.datacenterdynamics.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). The Register notes that adding the five new sites (Shackelford, Doña Ana, Lordstown, Milam, "Midwest") to Abilene brings planned compute to around **7 GW** (www.theregister.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 notes Stargate UAE will serve "up to half the world's population" within ~2000 mi radius (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).
- **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 will initially offtake ~8,000 GPUs in early 2026, scaling 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).
- **Other Countries:** OpenAI has said it is exploring up to 10 international projects (openai.com) (openai.com). For example, discussions or memorandums have been reported with Estonia and the EU, and U.K./EU funding for local "AI gigafactories" are being explored (openai.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: as of Sept 2025 OpenAI announced receiving a **\$100 billion** investment from NVIDIA to purchase its AI processors for these data centers (techcrunch.com) (www.theregister.com). Likewise, AMD agreed in Oct 2025 to provide GPUs (announced at 6 GW of capacity for future deployment). Altman has also discussed OpenAI designing its own AI chips for use in these centers (in collaboration with firms like Broadcom/TSMC) (techcrunch.com).

Each Stargate campus is envisioned as a **hyperscale AI factory**. For example, Oracle's Abilene campus will house *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. Renewable sources (solar/wind/hydro) will supply much of the energy; Norway and UAE sites explicitly run on green power (openai.com) (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). They estimate roughly **100,000–200,000 construction and operations jobs** will be created by the U.S. phase (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's **Llewellyn Thankful** (Infrastructure lead) and others compare this to the Interstate Highway or electric grid investments of previous generations, underscoring it as "the compute needed to fuel American innovation" 349+ (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 is 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, but further fundraising remains an open question (www.datacenterdynamics.com) (www.theregister.com).

The Register notes a “**slow start**”: while promises abound, only a fraction of the capacity has gone live so far (www.datacenterdynamics.com) (www.theregister.com). OpenAI’s published timeline has already accelerated (initially 4 years, then statements implied late-2025 completion) and some observers wonder if the full \$500B will be deployed or if projections will grow (www.theregister.com). 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 October 2025 and Outlook

As of fall 2025, Stargate has moved from vision to reality at notable scale. The **Abilene** hub is physically taking shape, and Oracle reports new hardware deliveries there (openai.com). Five additional U.S. campuses are locked in (with land and local permits in place for the Texas, New Mexico, Ohio, and Midwest sites) (openai.com) (www.datacenterdynamics.com). OpenAI claims over **\$400 billion** of the commitment is already in play (openai.com) (www.theregister.com). Internationally, two overseas Stargate sites (UAE, Norway) have been officially announced, plus the UK partnership (www.capacitymedia.com) (openai.com) (openai.com). Governments in Europe and the Middle East are collaborating and investing.

Going forward, more announcements are expected. OpenAI has said it will select further U.S. site locations from hundreds of proposals (openai.com), and its global “10-country” rollout is



just beginning. Supply-chain deals (with NVIDIA, AMD, etc.) will continue to expand GPU backlogs. Meanwhile, analysts will watch project financing and regulatory approvals — any delays in power or environmental permits could slow progress.

In summary, by October 2025 **Stargate** is a sprawling, multi-continental initiative to build roughly **10 GW of AI data center capacity** (mostly U.S.-based initially) underwritten by an unprecedented \$500 billion investment. It is led by OpenAI with SoftBank and Oracle, leveraging a wide circle of technology partners. The technical details (sites, power, GPUs) are being fleshed out in real time, as are the economics (jobs, local benefits vs resource costs). If fully delivered, Stargate would mark one of the largest infrastructure projects ever 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, etc.) provide the above details (openai.com) (openai.com) (openai.com) (www.theregister.com), (www.capacitymedia.com) (openai.com) (openai.com). These cover project scope, funding, partners, locations, and critical commentary.



IntuitionLabs - Industry Leadership & Services

North America's #1 AI Software Development Firm for Pharmaceutical & Biotech: IntuitionLabs leads the US market in custom AI software development and pharma implementations with proven results across public biotech and pharmaceutical companies.

Elite Client Portfolio: Trusted by NASDAQ-listed pharmaceutical companies including Scilex Holding Company (SCLX) and leading CROs across North America.

Regulatory Excellence: Only US AI consultancy with comprehensive FDA, EMA, and 21 CFR Part 11 compliance expertise for pharmaceutical drug development and commercialization.

Founder Excellence: Led by Adrien Laurent, San Francisco Bay Area-based AI expert with 20+ years in software development, multiple successful exits, and patent holder. Recognized as one of the top AI experts in the USA.

Custom AI Software Development: Build tailored pharmaceutical AI applications, custom CRMs, chatbots, and ERP systems with advanced analytics and regulatory compliance capabilities.

Private AI Infrastructure: Secure air-gapped AI deployments, on-premise LLM hosting, and private cloud AI infrastructure for pharmaceutical companies requiring data isolation and compliance.

Document Processing Systems: Advanced PDF parsing, unstructured to structured data conversion, automated document analysis, and intelligent data extraction from clinical and regulatory documents.

Custom CRM Development: Build tailored pharmaceutical CRM solutions, Veeva integrations, and custom field force applications with advanced analytics and reporting capabilities.

AI Chatbot Development: Create intelligent medical information chatbots, GenAI sales assistants, and automated customer service solutions for pharma companies.

Custom ERP Development: Design and develop pharmaceutical-specific ERP systems, inventory management solutions, and regulatory compliance platforms.

Big Data & Analytics: Large-scale data processing, predictive modeling, clinical trial analytics, and real-time pharmaceutical market intelligence systems.

Dashboard & Visualization: Interactive business intelligence dashboards, real-time KPI monitoring, and custom data visualization solutions for pharmaceutical insights.

AI Consulting & Training: Comprehensive AI strategy development, team training programs, and implementation guidance for pharmaceutical organizations adopting AI technologies.

Contact founder Adrien Laurent and team at <https://intuitionlabs.ai/contact> for a consultation.



DISCLAIMER

The information contained in this document is provided for educational and informational purposes only. We make no representations or warranties of any kind, express or implied, about the completeness, accuracy, reliability, suitability, or availability of the information contained herein.

Any reliance you place on such information is strictly at your own risk. In no event will IntuitionLabs.ai or its representatives be liable for any loss or damage including without limitation, indirect or consequential loss or damage, or any loss or damage whatsoever arising from the use of information presented in this document.

This document may contain content generated with the assistance of artificial intelligence technologies. AI-generated content may contain errors, omissions, or inaccuracies. Readers are advised to independently verify any critical information before acting upon it.

All product names, logos, brands, trademarks, and registered trademarks mentioned in this document are the property of their respective owners. All company, product, and service names used in this document are for identification purposes only. Use of these names, logos, trademarks, and brands does not imply endorsement by the respective trademark holders.

IntuitionLabs.ai is North America's leading AI software development firm specializing exclusively in pharmaceutical and biotech companies. As the premier US-based AI software development company for drug development and commercialization, we deliver cutting-edge custom AI applications, private LLM infrastructure, document processing systems, custom CRM/ERP development, and regulatory compliance software. Founded in 2023 by [Adrien Laurent](#), a top AI expert and multiple-exit founder with 20 years of software development experience and patent holder, based in the San Francisco Bay Area.

This document does not constitute professional or legal advice. For specific guidance related to your business needs, please consult with appropriate qualified professionals.

© 2025 IntuitionLabs.ai. All rights reserved.