



INTUITIONLABS / RESEARCH & PRACTICAL GUIDANCE

Oracle-OpenAI \$300B Deal Explained: 2026 Update

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 · 2025-11-30 · openai · oracle cloud infrastructure · ai infrastructure · cloud computing · data centers · stargate project · generative ai · large language models · financial analysis · nvidia · ai

Original article: <https://intuitionlabs.ai/articles/oracle-openai-300b-deal-analysis>



In this article

1. Executive Summary
 2. Introduction and Background
 3. The AI Computing Arms Race
 4. OpenAI's Growth and Funding
 5. Oracle's AI Strategy
 6. Oracle–OpenAI \$300B Contract Details
 7. Strategic Rationale
 8. Financial Analysis and Implications
 9. OpenAI's Projected Costs vs. Revenues
 10. Oracle's Financial Impact
 11. Progress and Implementation (to Dec 2025)
 12. Data Center Construction
 13. Chip Supply and Technology Partnerships
 14. Financing and Investment Flows
 15. Preliminary Outcomes & Current Status
 16. Multi-Stakeholder Perspectives and Analysis
 17. OpenAI's Perspective
 18. Oracle's Perspective
 19. Microsoft and Other Cloud Competitors
 20. National and Economic Perspectives
 21. Case Studies and Comparisons
 22. Data Analysis and Evidence
 23. Discussion: Implications and Future Directions
 24. National Competitiveness and Security
 25. Cloud Industry Impact
 26. Economic and Environmental Considerations
 27. Future Outlook
 28. Conclusion
-

[Revised October 10, 2026]

Executive Summary

Update (October 2026): The \$300 billion contract is still scheduled to start in 2027, but much has changed around it since our April revision. Oracle's remaining performance obligations reached **\$664 billion** at the end of its fiscal Q1 2027 (August 31, 2026), when cloud infrastructure revenue grew 121% to \$7.4 billion and quarterly capex was \$28.5 billion, after \$55.7 billion of capex in fiscal 2026 (www.prnewswire.com). Oracle expects fiscal 2027 capex of up to \$95 billion, with customers repaying up to \$25 billion of it, and an S&P downgrade in July 2026 left Oracle on the cusp of a junk credit rating (www.reuters.com). In March 2026,

Oracle and OpenAI dropped a planned 600 MW expansion of the Abilene campus, while their plan to develop another 4.5 GW of capacity remained on track (www.reuters.com). In September 2026, Oracle issued a force majeure notice on Project Jupiter, the New Mexico campus that is part of its OpenAI capacity, citing potential power delays; Oracle said the project remained on schedule (www.reuters.com). On the OpenAI side, Sam Altman said the company will not go public in 2026 (fortune.com), and OpenAI told investors its annualized revenue for September 2026 was almost \$50 billion (www.reuters.com).

In September 2025, **Oracle and OpenAI** announced a landmark **\$300 billion, five-year cloud computing contract**, beginning in 2027, to supply OpenAI with vast amounts of computing power (www.datacenterdynamics.com) (www.techradar.com). This deal is part of a much larger AI infrastructure campaign (often referred to as the “**Stargate**” project) involving OpenAI, Oracle, and partners like SoftBank, which collectively aims to build up to ~30 gigawatts (GW) of AI computing capacity in the U.S. at a total investment on the order of half a trillion to a trillion dollars (www.reuters.com) (www.reuters.com). The Oracle–OpenAI agreement alone covers roughly **4.5 GW** of data center capacity per year, enough power to supply millions of homes (www.tomshardware.com) (www.techradar.com). It represents one of the largest commercial cloud contracts ever signed.

Despite the headline size, the deal carries **substantial risks and challenges**. OpenAI’s revenues (~\$10–12 B annual run-rate by mid-2025, growing to ~\$25 B ARR by February 2026 (www.cnbc.com) (openai.com)) are still dwarfed by its projected \$60 billion per year cloud-compute bill, implying OpenAI must aggressively raise funds, cut costs, or secure credit to fulfill its end of the deal. However, a massive \$122 billion funding round closed in March 2026 at an \$852 billion valuation (www.cnbc.com), significantly bolstering OpenAI’s war chest. Oracle, for its part, must invest heavily in new data centers and attract sufficient customers to cover its colossal infrastructure investments. Credit agencies have warned this deal could stress Oracle’s finances (Moody’s flagged counterparty risk and increased leverage (www.reuters.com)), and Oracle’s FY2026 capex guidance was raised to ~\$50 billion (investor.oracle.com); actual fiscal 2026 capex came to \$55.7 billion (www.prnewswire.com).

As of early 2026, **implementation has progressed significantly**. The flagship Stargate data center in Abilene, Texas went live on Oracle Cloud Infrastructure in September 2025 (www.crusoe.ai) and is being built out to 1.2 GW across eight buildings, two of which were running by March 2026 (www.reuters.com). Construction continues on additional sites across multiple U.S. states, with the Stargate project expanding to nearly **7 gigawatts** of planned capacity and over **\$400 billion** in investment (openai.com). Oracle is acquiring hundreds of thousands of Nvidia GB200 GPUs (~\$40B worth) to stock these facilities (www.reuters.com). In parallel, OpenAI is diversifying its supply of compute by designing its own chips with partners like Broadcom, with the first custom AI inference chips expected to deploy in H2 2026 (www.tomshardware.com). OpenAI has also completed its restructuring into a Public Benefit Corporation and closed a record \$122 billion funding round at an \$852 billion valuation (www.cnbc.com). An IPO once expected in late 2026 has been put off: Altman said in September 2026 that it will not happen this year (fortune.com).

This report provides a comprehensive deep dive into the Oracle–OpenAI partnership and “**300B**” project, originally written in December 2025 and updated through October 2026. We chronicle the origins and context of this deal, analyze its financial and strategic implications, review progress and bottlenecks in infrastructure

rollout, compare it to other mega-deals in cloud/AI, and consider future scenarios. All claims are backed by recent reporting and data sources including news agencies and technical studies.

Introduction and Background

The AI Computing Arms Race

The past three years have seen **explosive growth in generative AI**, led by OpenAI's ChatGPT. Surging demand for **large language models (LLMs)** like GPT-3 and GPT-4 has placed unprecedented pressure on data center infrastructure. Training GPT-3 in 2020 alone consumed an estimated **1,300 megawatt-hours (MWh)** of electricity ([arxiv.org](#)), and future models may require **gigawatt-scale** continuous power budgets ([arxiv.org](#)). Energy consumption and computational cost are now among the central constraints on AI advancement. Studies warn that AI **model training** and inference are driving “unprecedented increase(s) in the electricity demand of AI data centers” ([arxiv.org](#)), posing challenges for grid capacity and sustainability planning.

To put numbers into perspective, large language models are consuming power on the order of dozens of *nuclear reactors*. For example, Oracle, when announcing the OpenAI contract, noted it involves **4.5 GW** of data center capacity annually, “equivalent to what four million homes use” ([www.tomshardware.com](#)) ([www.techradar.com](#)). Such scales are comparable to the largest conventional power plants. Recent academic work underlines that AI's hunger for compute is **exponential**. One analysis projected that state-of-the-art LLMs may soon demand *city-scale* power (gigawatt levels) for training ([arxiv.org](#)). Another found that an 8-GPU system based on Nvidia's H100 accelerator still consumes tens of kilowatts under load and emphasized that current estimates of AI energy usage incorporate considerable uncertainty ([arxiv.org](#)). Taken together, experts emphasize that **AI's infrastructure needs are skyrocketing**, fueling demand for massive computing and data center investments.

OpenAI's Growth and Funding

Since releasing **ChatGPT** in late 2022, OpenAI has seen stunning growth in user uptake and revenue. By mid-2025, Reuters reported OpenAI had reached **\$10–12 billion** in annualized recurring revenue ([www.reuters.com](#)) ([www.reuters.com](#)). This includes both consumer subscriptions and commercial API sales. By July 2025, ChatGPT boasted around **700 million weekly active users**, doubling its user base earlier in the year ([www.reuters.com](#)). Despite this scale, less than 10% of its users were paid subscribers ([www.reuters.com](#)), leaving the company with significant free-usage burdens on its servers. In 2024, OpenAI reportedly lost \$5 billion on revenues of about \$6 billion ([www.reuters.com](#)), reflecting the high fixed costs of the GPU clusters it runs and the rapid scaling of its service. Revenue growth has continued at a staggering pace: OpenAI CFO Sarah Friar confirmed \$20 billion in 2025 revenue, and by February 2026, the company had reached **\$25 billion ARR**, generating approximately **\$2 billion per month** ([openai.com](#)). OpenAI now counts over **9 million paying business users** and has launched an advertising pilot that already topped \$100 million in annualized revenue within its first two months ([www.cnn.com](#)), with projections of \$2.5 billion in ad revenue for 2026 and a target of \$100 billion by 2030 ([www.investing.com](#)).

OpenAI's valuation has correspondingly soared. A SoftBank-led funding round in mid-2025 valued the company at ~\$300 billion (www.reuters.com), and by autumn, SoftBank had helped drive a ~20% increase to a \$500 billion valuation (www.reuters.com). Then in March 2026, OpenAI closed a record-shattering **\$122 billion funding round** at an **\$852 billion valuation** (www.cnbc.com) (techcrunch.com). Key investors included Amazon (\$50 billion, with \$35 billion contingent on IPO or AGI milestones), Nvidia (\$30 billion), SoftBank (\$30 billion), alongside Andreessen Horowitz, Microsoft, and others. The company also completed its long-anticipated **restructuring into a Public Benefit Corporation (PBC)**, with a nonprofit Foundation holding a \$130 billion stake and Microsoft retaining ~27% ownership (openai.com). An IPO was reportedly being targeted for Q4 2026, with a potential \$1 trillion valuation goal (www.indexbox.io). In September 2026, Altman said OpenAI will not go public in 2026 (fortune.com), and OpenAI was reported to be discussing a new funding round at a valuation of about \$1.2 trillion (fortune.com). However, OpenAI is still not profitable: the company is projected to lose approximately **\$14 billion in 2026** (finance.yahoo.com), with cumulative long-term spending projections of ~\$115 billion through 2029. OpenAI CEO Sam Altman has publicly warned that continuing this growth requires “trillions” in infrastructure spending (www.datacenterdynamics.com) (www.reuters.com).

Oracle's AI Strategy

Oracle Corporation historically built its business on database and enterprise software. In recent years it has tried to expand its cloud infrastructure presence (Oracle Cloud Infrastructure, or OCI) to compete with the likes of Amazon Web Services (AWS), Microsoft Azure, and Google Cloud. Oracle co-founder Larry Ellison, its chairman and chief technology officer, has been particularly bullish on AI, pledging that “AI is the next frontier” for Oracle (www.datacenterdynamics.com). The company's stock jumped 40+% on news of the OpenAI deal (www.techradar.com) (www.pcgamer.com), underscoring investor enthusiasm for Oracle's pivot to high-growth AI infrastructure (even as analysts warn of execution risk). Oracle's quarterly filings show remaining performance obligations (future revenue under signed contracts) reached **\$553 billion** as of Q3 FY2026, up 325% year-over-year (investor.oracle.com), and **\$664 billion** by Q1 FY2027, the quarter ended August 31, 2026 (www.prnewswire.com). Q3 FY2026 was Oracle's strongest organic quarter in over 15 years: total revenue rose 22% to \$17.2 billion, while cloud revenue jumped 44% to \$8.9 billion, with Cloud Infrastructure (IaaS) alone surging 84% (www.cnbc.com). However, Moody's maintains a **negative outlook** on Oracle's Baa2 rating, warning that debt could surge to 4× its earnings given these massive AI commitments (www.reuters.com). Oracle's FY2026 capex guidance was raised by \$15 billion to **\$50 billion total**, and the company raised \$30 billion in debt and equity financing in early 2026 to fund the buildout (investor.oracle.com). Barclays downgraded Oracle's debt to Underweight in late 2025, citing a 500% debt-to-equity ratio and warning it could approach BBB- (the lowest investment-grade rating). Despite these financial pressures, Oracle's stock rose 10% on Q3 FY2026 earnings, though shares were down roughly 24% year-to-date in early April 2026 amid broader AI bubble concerns (www.fool.com).

Between them, Microsoft (the longtime OpenAI cloud partner) and Oracle represent competing cloud ecosystems. Historically, OpenAI's compute workloads ran almost entirely on Azure. The Oracle deal is a strategic shift to “diversify its cloud platform partnerships” (www.tomshardware.com) (www.techradar.com). Microsoft relaxed exclusivity clauses in early 2025 to allow OpenAI to pursue new cloud compute sources (www.reuters.com). Now OpenAI is actively engaging both the former (Microsoft) and new (Oracle, AWS, Google, etc.) for compute capacity.

Table 1 summarizes key AI infrastructure agreements relevant to this report:

Announced Deal	Parties	Term	Value	Purpose
OpenAI–Oracle Cloud (Sep 2025)	OpenAI & Oracle	5 years	~\$300 billion	Supply ~4.5 GW/year of compute (cloud services) starting 2027 (www.datacenterdynamics.com) (www.tomshardware.com)
OpenAI–SoftBank–Oracle ‘Stargate’ JV (Jan 2025)	OpenAI, SoftBank, Oracle, US govt	Multi-year	~\$500 billion	Build ~10 GW of AI data centers in U.S. (later upsized) (www.reuters.com) (www.reuters.com)
OpenAI–Nvidia GPU Supply (mid-2025)	Oracle & Nvidia (for OpenAI)	N/A	~\$40 billion	Purchase ~400,000 Nvidia GB200 GPUs for new Abilene TX DC (www.reuters.com)
Crusoe Texas Data Center Funding (May 2025)	Crusoe (for OpenAI)	N/A	\$15 billion	Expand OpenAI’s largest US data center (Abilene, TX) from 2 to 8 buildings (www.reuters.com)
Broadcom–OpenAI Chip Deal (Oct 2025)	OpenAI & Broadcom	~4 years	Undisclosed	Co-develop 10 GW of custom AI accelerators (deployment by 2029) (www.reuters.com)
Oracle–Meta Cloud (rumor, Sept 2025)	Oracle & Meta (report)	N/A	~\$20 billion	Provide cloud AI computing power to Meta (if finalized) (www.reuters.com)
Amazon–OpenAI Cloud (Nov 2025)	Amazon & OpenAI	N/A	~\$38 billion	AWS supply computing for OpenAI (public reports) (www.reuters.com)
Meta–Google Cloud (Aug 2025)	Meta & Google	6 years	~\$10+ billion	Google Cloud services (servers, storage, networking) (www.reuters.com)

Notes: The Oracle–OpenAI deal dwarfs these other contracts (“one of the largest cloud computing deals in history” (www.tomshardware.com)). Terms, exact start dates, and deliverables vary by agreement.

Oracle–OpenAI \$300B Contract Details

On **September 10, 2025**, *The Wall Street Journal* first reported that OpenAI had entered a **\$300 billion, five-year contract** with Oracle to “procure computing power” (www.reuters.com). This was quickly confirmed by multiple outlets (e.g. DataCenterDynamics and Tom’s Hardware (www.datacenterdynamics.com) (www.tomshardware.com)). The agreement is commonly phrased as “[OpenAI] agreed to buy \$300 billion in computing power over five years from Oracle” (www.datacenterdynamics.com). The deal is slated to **begin in 2027** and run through roughly 2031/2032 (www.datacenterdynamics.com) (www.tomshardware.com).

Scope and Deliverables

While the exact details (e.g. how many machines or what services) are proprietary, public reports provide key figures:

- **Power Consumption:** The contract **assumes 4.5 gigawatts of power** usage year-round on Oracle’s infrastructure (www.tomshardware.com) (www.techradar.com). For comparison, 4.5 GW is roughly enough for “about four million homes” (www.techradar.com). This underscores that OpenAI will be

renting an immense scale of data center uptime and capacity.

- Annual Spend:** The \$300B over five years implies roughly **\$60B per year** in charges to OpenAI. (Earlier in 2025, OpenAI and Oracle had arranged for ~4.5 GW in 2028 at \$30B/year (www.datacenterdynamics.com); the new larger deal essentially doubles that annual commitment.) At ~\$60B/year, Oracle would see tens of billions in annual revenue from this single client (on the order of \$30B–\$60B/year, depending on accounting) (www.techradar.com) (www.pcgamer.com).
- Infrastructure Build-Out:** To meet this demand, Oracle plans to construct entirely new data center campuses. As part of Oracle's preparations, **Crusoe**, a specialized data-center builder, has been contracted to develop large facilities. For example, Oracle and Crusoe are expanding an Abilene, Texas campus (initially 2 buildings) to eventually 8 buildings (www.reuters.com). Oracle has also begun building a data center in Shackelford County, Texas, and acquired an Ohio site for hardware manufacturing (www.reuters.com) (apnews.com). This suggests Oracle will deliver the contracted compute via co-owned cloud data centers (leasing capacity to OpenAI) rather than purely virtual renting of existing infrastructure.
- Technology Stack:** The deal covers “cloud computing power,” which likely means a mix of GPUs/AI accelerators, networking, storage, and custom support services. Oracle plans to populate these centers with cutting-edge GPUs: a May 2025 report indicates Oracle is buying roughly **400,000 Nvidia GB200 GPUs** (the newest “Blackwell” series) for ~\$40B to equip its new Texas data center (www.reuters.com). In addition to Nvidia hardware, OpenAI is diversifying with in-house chip designs (via Broadcom) and deals with AMD (www.reuters.com) (www.reuters.com). Oracle's infrastructure will thus include Nvidia GPUs plus possibly other accelerators as developed by OpenAI.

Table 2 summarizes these contract parameters:

Parameter	Figure	Source/Notes
Contract value	~\$300 billion (5 years)	Per reported Wall St. Journal & Reuters reports (www.reuters.com) (www.datacenterdynamics.com)
Contract term	2027–2032 (approx.)	Reported as a five-year agreement starting in 2027 (www.datacenterdynamics.com)
Annual compute cost	~\$60 billion/year	Implied by \$300B over 5 years (www.datacenterdynamics.com) (www.techradar.com)
Annual power demand	~4.5 GW	Oracle cited 4.5 GW figure (www.tomshardware.com) (www.techradar.com)
Total chips supported	~2–10 million NVIDIA GPUs	2M chips for initial phases (www.reuters.com); Oracle to buy 400k Nvidia GB200 (www.reuters.com)
Oracle revenue impact	+\$30–60B/year (projected)	Oracle disclosed \$317B in new contract rev, much from this deal (www.techradar.com)
Oracle stock reaction	+40% (intraday jump)	Price jumped ~43% on news (www.tomshardware.com) (www.pcgamer.com)

Parameter	Figure	Source/Notes
OpenAI spend per year	~\$60 billion	As above; far above current revenues (www.datacenterdynamics.com) (www.techradar.com)

This deal is **unprecedented in scale**. By comparison, in October 2025 OpenAI contracted to purchase an incremental \$250 billion of Azure services from Microsoft (blogs.microsoft.com) (www.reuters.com), and Amazon closed a smaller ~\$38B OpenAI cloud deal (www.reuters.com). Oracle's contract is larger than either. It is "one of the largest cloud contracts in history" (www.tomshardware.com) (www.techradar.com) and would alone quadruple Oracle's cloud revenue if fully realized.

Strategic Rationale

The deal addresses **both companies' strategic needs**:

- OpenAI Needs:** OpenAI's current infrastructure (primarily Microsoft Azure) was under immense strain from user demand. CEO Sam Altman has said scaling AI inevitably requires multi-hundred-billion-dollar investments (www.datacenterdynamics.com) (www.axios.com). By committing to Oracle, OpenAI secures vast additional capacity beyond what Azure could easily provide. This diversification also buys leverage and flexibility: OpenAI will no longer be wholly dependent on a single cloud provider. In short, OpenAI needed a "landslide" of resources to train ever-larger models and serve millions of users; the Oracle deal is a direct response to that need.
- Oracle Goals:** For Oracle, the deal is a major coup in its bid to become a major AI cloud player. Traditionally trailing the larger hyperscalers, Oracle sees supplying OpenAI as a way to instantly become one of the top cloud vendors (by percent revenue growth at least). Larry Ellison publicly pushed OCI's performance, and this deal dramatically raises Oracle's profile and (potentially) its future revenues. It also justifies Oracle's own massive infrastructure buildout (the "Stargate" data centers) and paves the way for selling capacity to other AI firms. As one observer noted, Oracle's then-CEO Safra Catz highlighted a surge in \$455B of remaining performance obligations, largely AI deals (www.datacenterdynamics.com); this Oracle–OpenAI contract is a big chunk of that backlog.
- State and Competitiveness:** The U.S. government has strongly backed domestic AI infrastructure. The Stargate initiative, announced by President Trump in January 2025, explicitly gathered OpenAI, Oracle, and SoftBank to invest in American data centers (www.reuters.com) (www.reuters.com). This was partly motivated by strategic competition with China. The Oracle–OpenAI pact fits this national strategy: it "aims to solidify the United States' leadership in artificial intelligence" by massively boosting onshore compute capacity (www.reuters.com). In return, the government has signaled support (e.g. regulatory permissions, potential infrastructure aid) to help push these projects forward.

Table 3: Major Compute Capacity and Investment Commitments (as of late 2025)

Project / Initiative	Compute Target (GW)	Investment	Partners/Notes
OpenAI Infrastructure (Altman vision)	~30 GW (total eventual)	~\$1.4 trillion (total)	Sam Altman: aim for 1 GW/week, scaling AI capability (www.reuters.com) (www.axios.com)
“Stargate” AI Data Center Program	~10–11 GW (initial goal)	~\$500 billion (announced)	OpenAI, Oracle, SoftBank (US govt-supported) (www.reuters.com) (www.reuters.com)
Oracle–OpenAI Cloud Deal	4.5 GW (per year)	\$300 billion (5-year contract)	Oracle to supply cloud compute (begin 2027) (www.datacenterdynamics.com) (www.tomshardware.com)
OpenAI–Nvidia GPU Supply (Abilene DC)	~1 GW (projected)	~\$40 billion (chips)	400k Nvidia GB200 GPUs to support Texas data center (www.reuters.com)
Broadcom–OpenAI Chip Project	10 GW (by 2029)	Undisclosed (\$50–\$60B/GW est.)	Co-design of custom AI processors (10 GW; Broadcom) (www.reuters.com)
Crusoe–OpenAI Texas DC (Abilene)	1.2 GW (phase 1)	~\$15.0 billion (funding)	Expand initial build from 2 to 8 buildings (www.reuters.com)
SoftBank–Stargate Ohio Factory	N/A (infrastructure %)	\$3 billion (factory invest)	Lordstown EV plant → data center modules (www.reuters.com)

Each of these commitments interlocks. The Oracle–OpenAI deal is essentially the “payback” arrangement for the compute provided by the Stargate centers: Oracle builds and equips the centers (with Crusoe, Nvidia, etc.), then leases that capacity to OpenAI under the \$300B contract. Meanwhile, OpenAI is also pursuing chip-level solutions (Broadcom, AMD) to control costs.

Financial Analysis and Implications

OpenAI’s Projected Costs vs. Revenues

The \$300B commitment by OpenAI represents a **dramatic escalation in its spending**. Even after its rapid revenue growth, OpenAI’s financial outlook was stretched. As of mid-2025, OpenAI’s annualized revenue was roughly \$10–12 B (www.reuters.com) (www.reuters.com). By year-end, company targets projected around \$20 B in 2025 (www.reuters.com), thanks to growing subscription and API sales. But none of this is profit; OpenAI lost ~\$5 B in 2024 (www.reuters.com) as it scaled up servers. Assuming similar growth rates, OpenAI’s expenses (capex + opex) on compute could easily match or exceed its gross revenues in the short term.

The core issue is that **paying \$60 B per year for compute would dwarf OpenAI’s current income**. Table 4 illustrates the mismatch: even if OpenAI hit \$20 B revenue by 2025 (on track from \$12 B mid-year), that leaves \$40 B/year to cover via other means just to pay Oracle. Analysts have labeled this situation “counterparty risk,” since a few customers (OpenAI, VMware, etc.) now make up a huge portion of Oracle’s RPO (www.reuters.com). OpenAI must therefore find new revenue channels to cover these costs.

Proposed strategies include:

- **Higher Pricing & Ads:** OpenAI is exploring ways to monetize more of its user base. It plans to increase ChatGPT price tiers (e.g. more Pro subscribers), add ads, and open payments features (www.itpro.com). These could push revenue upward, but likely not enough in the near term.
- **Enterprise Deals:** OpenAI seeks large licensing deals (enterprise customers, government) to boost income. Recent reports suggest it is using debt to lease GPUs, raising external capital (SoftBank, etc.), and entering B2B partnerships.
- **Cutting Costs:** By designing its own chips, OpenAI hopes to reduce reliance on expensive third-party GPUs (www.reuters.com). Altman cited potential drops in per-chip costs due to competition (www.itpro.com).
- **Investor Funding:** Continuing to raise capital is part of the plan. SoftBank and others have injected tens of billions (e.g. Vision Fund financing, SoftBank's later efforts) to cover near-term spending (www.reuters.com) (www.reuters.com).
- **Delayed ROI:** OpenAI's leadership, including Altman, has stated that they expect to run at a loss for the foreseeable future (no profit until ~2029) as long as growth continues (www.tomshardware.com) (www.itpro.com). Microsoft's new deal (October 2025) also suggests confidence in a long runway (www.reuters.com).

Table 4: OpenAI Income vs. Infrastructure Spend (2025–2026)

Metric	Value	Source / Note
OpenAI 2024 Revenue	~\$6 billion	Confirmed by CFO Sarah Friar (openai.com)
OpenAI 2025 Revenue	~\$20 billion	Confirmed by CFO (openai.com)
OpenAI ARR (Feb 2026)	~\$25 billion (~\$2B/month)	(openai.com)
OpenAI annualized revenue (Sep 2026)	Almost \$50 billion	Told to investors (www.reuters.com)
OpenAI Projected 2026 Loss	~\$14 billion	Projected (finance.yahoo.com)
Oracle Deal Annual Cost to OAI	~\$60 billion/year	\$300B over 5 years (www.datacenterdynamics.com)
Shortfall (w/o new funds)	~\$35 billion (gap)	Cost minus Feb 2026 ARR (approx.); about \$10 billion against the Sep 2026 run rate
Capital Raised (Mar 2026 round)	\$122 billion	At \$852B valuation (www.cnbc.com)
Long-term Spending (through 2029)	~\$115 billion	Projected (finance.yahoo.com)

Sources: OpenAI revenue confirmed by CFO Sarah Friar (openai.com); Oracle contract terms (www.datacenterdynamics.com) (www.tomshardware.com). Figures are approximate. The \$122B funding round substantially narrows the near-term gap, though OpenAI's long-term compute obligations still require continued revenue growth and future capital raises.

Thus, barring other cash inflows, OpenAI would need **massive external funding** each year to sustain operations. In practice, SoftBank (Vision Fund), Microsoft, and others have filled the gap. SoftBank's investment rounds (mid-2025) raised up to \$40B (www.reuters.com). In Q3 2025, SoftBank in fact realized part of its AI strategy: it boosted OpenAI's valuation from \$300B to \$500B and took profits from earlier chip investments (www.reuters.com) (www.reuters.com). Microsoft's recent recapitalization deal (late Oct 2025) also relaxed financial constraints on OpenAI (www.reuters.com). Cumulatively, these moves aim to align OpenAI's war chest with its trillion-dollar ambitions.

Oracle's Financial Impact

For Oracle, the \$300B contract is a windfall, if delivered, but also demands onerous upfront investment. Oracle's reported **remaining performance obligations (RPO)** have grown rapidly, from \$455B in Q1 FY2026 to **\$553 billion** by Q3 FY2026 (up 325% year-over-year) (investor.oracle.com), reflecting new long-term contracts (including OpenAI's). Oracle Q3 FY2026 was the company's best organic quarter in 15+ years, with total revenue up 22% to \$17.2 billion and cloud revenue soaring 44% to \$8.9 billion (www.cnbc.com). Cloud Infrastructure (IaaS) was the standout, surging 84% year-over-year to \$4.9 billion. Oracle guided for Q4 cloud revenue growth of 44–50% in USD. By Q1 FY2027 (the quarter ended August 31, 2026), RPO had reached \$664 billion, cloud revenue grew 62% to \$11.6 billion and IaaS revenue grew 121% to \$7.4 billion (www.prnewswire.com).

However, analysts have raised concerns. Moody's has warned that Oracle's **debt would grow faster than earnings**, pushing leverage toward 4× EBITDA due to capital spending required for Stargate data centers (www.reuters.com). Indeed, Oracle disclosed that fulfilling these contracts might require raising capital via bond issuances or loans (some already reported by banks (www.reuters.com)). Investors have noted that Oracle's stock, after the initial 40% spike (www.techradar.com) (www.pcgamer.com), pulled back (partly due to profit-taking and partly due to realization of the costs ahead).

In absolute terms, Oracle expects to earn something like **\$30–\$60 B per year** in incremental cloud revenue from OpenAI alone (www.tomshardware.com) (www.pcgamer.com). This could catapult its cloud segment to rival AWS, at least in revenue scale. Oracle's then-CEO Safra Catz hinted that AI-related spending by customers could make Oracle Cloud revenue exceed half a trillion dollars eventually (www.reuters.com). But in the meantime, Oracle must **outlay tens of billions** on facilities and equipment. Reports show Oracle financing its Texas data center via \$9.6 B bank debt and \$5 B of equity contributions (www.reuters.com). Additional bank loans (e.g. \$18 B in November 2025 (www.reuters.com), \$38 B under discussion (www.reuters.com)) keep augmenting the debt pile.

Ultimately, Oracle's return on this investment depends on OpenAI and others using that capacity (OpenAI is reportedly a major tenant in these sites (www.reuters.com)). Moody's vs. Oracle's leadership differ on risk / reward: Oracle bets on outsized growth in AI demand (and has said it will soak up excess chip supply (www.reuters.com)), while raters caution about the leverage and execution risk. Oracle is still executing on earlier plans (such as integrating Oracle Cloud with AWS, Google and other clouds to attract more workloads (www.reuters.com)), but its fortunes are now tightly coupled to OpenAI's success.

Progress and Implementation (to Dec 2025)

By the end of 2025, **key elements of the Oracle–OpenAI “300B” project have moved from plan toward realization**, though full deployment is years away. We examine the state of play on infrastructure buildout, chip supply, financing, and any emerging bottlenecks.

Data Center Construction

OpenAI’s 2025 announcements centered on building physical data centers under the “Stargate” umbrella (OpenAI/Oracle/SoftBank alliance). Originally unveiled in January 2025 as a plan for \$500B of infrastructure, **the focus has sharpened to meet immediate needs**. Notably:

- Abilene, Texas – Primary U.S. Site (Now Live):** The first large Stargate facility is running on Oracle Cloud Infrastructure in Abilene, TX. Crusoe announced the first phase of the flagship campus as live in September 2025 (www.crusoe.ai), and the final building was topped out by Q1 2026 (www.datacenterdynamics.com). The campus spans approximately 4 million square feet with a total power capacity of **1.2 GW**, with full campus completion expected by mid-2026. In March 2026, Oracle and OpenAI dropped a planned 600 MW expansion next to the campus, and a source told Reuters that two of the eight buildings were running at that point (www.reuters.com). It houses “hundreds of thousands” of Nvidia GB200 GPUs (apnews.com) using innovative “H-shaped” hall designs engineered for liquid-cooled GB200 racks. The site employed ~6,000 workers during build-out and will have ~1,700 permanent jobs (apnews.com). In a notable development, **Microsoft** leased 900 MW of capacity from Crusoe’s adjacent Abilene site, an expansion that OpenAI had declined to pursue, in a sign of the two companies’ increasingly separate paths. Crusoe is building two new “AI factory” buildings and a 900 MW on-site power plant for Microsoft, right next to the OpenAI/Oracle campus (fortune.com).
- New Mexico – Planned Campus:** In November 2025, Bloomberg/Reuters reported that **20+ banks arranged an \$18 B project loan** to build a new Oracle–OpenAI data center campus in New Mexico (www.reuters.com). This campus is part of Stargate and is expected to become one of the consortium’s largest sites. It will operate under typical project-finance terms (2.5%+ SOFR, 4-year principal with extensions (www.reuters.com)) and is intended to be a major OpenAI hub (Oracle is expected to be a tenant (www.reuters.com)). The aggregate Stargate plan envisaged 5 new centers by Sept 2025 (www.reuters.com), and New Mexico is one of them (Dona Ana County, see below). In September 2026, Oracle issued a force majeure notice to the developer, a Blue Owl unit, citing potential delays in securing power for the campus, known as Project Jupiter and planned to come online in 2028; Oracle said Project Jupiter remained on its planned schedule (www.reuters.com).
- Lordstown, Ohio – Conversion of a Factory:** SoftBank announced a \$3 B investment to convert a shuttered EV factory in Lordstown into a **modular data center fabrication plant** (www.reuters.com). This facility will produce prefabricated data center units (containers) for OpenAI’s deployments (targeting sites in Texas, Ohio, etc.). The Ohio plant will also include a demonstration data center, with production starting in 2026 (www.reuters.com). This leverages a SoftBank purchase of the site (for \$375 M) announced in Aug 2025.

- **Additional Sites – Expanding Rapidly:** By early 2026, the Stargate project has expanded well beyond its original five-site plan. Oracle is developing additional sites in **Wisconsin**, with Vantage (openai.com), and **Michigan**, where construction of the Saline Township campus was underway by June 2026 (www.oracle.com). The combined capacity from the five sites announced in September 2025, along with the flagship Abilene campus and ongoing projects with CoreWeave, brings Stargate to nearly **7 gigawatts** of planned capacity and over **\$400 billion** in investment over the next three years (openai.com). The three Oracle-developed sites (Shackelford County, Texas; Doña Ana County, New Mexico; and the Midwest site) plus a 600 MW Abilene expansion, since dropped, were expected to deliver over 5.5 GW of capacity. The Milam County, Texas campus is being developed by SB Energy (a SoftBank Group company), providing a 1.2 GW powered infrastructure site with initial facilities expected to enter service in 2026. Crusoe raised \$1.38 B in Oct 2025 (Series E) at ~\$10 B valuation (www.reuters.com) explicitly to fund the Abilene expansion, bringing \$15 B total to date for that Texas project. The next wave of Stargate compute will feature **Nvidia Vera Rubin** next-generation chips, with shovels in the ground now laying foundations for compute expected to come online in 2026–2027.
- **“Scaled Back” Ohio Center:** In mid-2025, the Stargate Ohio plan was reduced to a smaller data center by end-2025 (www.reuters.com). However, SoftBank’s \$3 B Ohio factory conversion (Lordstown) continues as a modular data center fabrication plant, with production starting in 2026.

The takeaway is that **construction is now well underway across multiple states**. The Abilene flagship is live and operational, additional sites are in active construction, and the Stargate program has expanded from an initial \$500 billion commitment to over \$400 billion in near-term investment alone, with OpenAI stating these are just the first set of site selections with additional locations to come (openai.com).

Chip Supply and Technology Partnerships

In parallel with data center construction, OpenAI and Oracle have lined up chip suppliers and innovators:

- **Nvidia GPUs:** Oracle is investing ~\$40 B to buy ~400,000 of Nvidia’s latest AI accelerators (GB200 series) (www.reuters.com). These chips will power the new U.S. Stargate centers (e.g. Abilene DC) and reduce OpenAI’s dependency on Microsoft’s Azure cluster resources. Oracle will lease this capacity to OpenAI, sidestepping Microsoft’s current supply constraints (www.reuters.com). In late 2025 reports noted Nvidia committed up to \$100B in chip supplies to the initiative (www.reuters.com). Such massive GPU pools are critical, since each ~50k-GPU building (estimated cost ~\$3B–\$4B per building (www.reuters.com)) can train and run advanced LLMs.
- **Broadcom Custom Chips (Design Complete, Deploying H2 2026):** In October 2025, OpenAI announced a **partnership with Broadcom** to co-develop its first in-house AI processors (www.reuters.com). The deal aims to deploy over **10 GW of Broadcom-based AI hardware** by 2029. As of early 2026, the two companies have **completed the design phase** of a custom AI inference engine, specifically optimized for OpenAI’s “o1” series inference models and future GPT versions. The chip uses a “system array” design optimized for dense matrix multiplication in the Transformer architecture. First deployment in data centers is expected in **H2 2026** (www.tomshardware.com). The racks will be scaled entirely with Ethernet and other connectivity solutions from Broadcom. Building 1

GW of data center capacity costs approx \$50–60 B, implying the Broadcom project alone corresponds to up to ~\$600 B of capex (www.reuters.com). Broadcom has since expanded its custom chip business, agreeing to deals with Google and Anthropic as well (www.cnbc.com), confirming that custom AI silicon is now a major industry trend.

- **AMD Collaboration:** On October 6, 2025, a week before the Broadcom announcement, OpenAI and AMD announced a 6 GW agreement starting with AMD Instinct MI450 GPUs, with the first 1 GW deployment set to begin in the second half of 2026 and a warrant for OpenAI to acquire up to 160 million AMD shares (ir.amd.com). These complement Nvidia/Oracle supplies.
- **Networking and Systems:** The Broadcom partnership also encompasses high-speed networking (Broadcom's InfiniBand alternatives), which are crucial for scaling GPUs into 1M+ chip clusters. Other vendors (Marvell, etc.) are reportedly involved in networking gear. Oracle is likely procuring turnkey solutions (servers, racks, cooling systems) alongside chips; Crusoe and operators like Cold Fusion provide integrated design/engineering.

In sum, **the hardware pipeline is confirmed** by credible reporting: Oracle is securing hundreds of thousands of GPUs and partnering in advanced chip design (www.reuters.com) (www.reuters.com). This reduces execution risk associated with supply shortages (one reason Microsoft had capped OpenAI GPU capacity in mid-2025). It also diversifies sources and may lower long-term costs. One metric cited is that each gigawatt of capacity costs an estimated \$50–60 billion, indicating the sheer magnitude of capital required to field each gigawatt of AI power (www.reuters.com).

Financing and Investment Flows

Building and equipping these data centers has demanded extensive financing. Multiple funding vehicles have been reported:

- **Debt Syndications:** In November 2025, a consortium of ~20 banks agreed to finance an \$18 B loan for Oracle's New Mexico DC campus (www.reuters.com). Interest pricing was quoted as SOFR+2.5%, 4-year maturity [with extension options] (www.reuters.com). Similarly, banks are negotiating a **\$38 B loan** to Oracle/Vantage (DC developer) for further U.S. site development (www.reuters.com). These loans show traditional finance is backing the venture at sizable scale. (For comparison, even large corporate loans seldom approach these sizes; the scale underscores how data center projects are being funded almost like infrastructure utilities.)
- **Equity and VC Funding:** Data center operator Crusoe secured \$1.38 B in Oct 2025 (Series E) at a ~\$10 B valuation (www.reuters.com). Its investors include major VCs (Valor, Mubadala, Nvidia, Fidelity, Founders Fund). This round specifically cites Crusoe's involvement in building OpenAI's Abilene data center (launched 1.2 GW of capacity) (www.reuters.com). Earlier in 2025, Crusoe had raised ~\$3.9 B total, including earlier funding that financed most of Abilene's build (www.reuters.com). Blue Owl Capital and Valor reportedly provided equity as part of the Oracle-Abilene project (www.reuters.com).

- **OpenAI Funding Rounds:** OpenAI itself is raising more capital. In Fall 2025, OpenAI raised \$15–20 B in new capital (e.g. from SoftBank, Tiger Global, Fidelity) at a ~\$500 B valuation (www.reuters.com). These funds are explicitly to help finance OpenAI's cloud purchases and ongoing ops. Notably, SoftBank has been willing to re-allocate funds (selling its Nvidia stake (www.reuters.com)) and issue bonds to prop up OpenAI investments (www.reuters.com).
- **Government and Policy Support:** U.S. policy has indirectly supported financing. The White House (under President Trump) announced the Stargate initiative as a national priority (Jan 2025) (www.reuters.com). New reports (Nov 2025) suggest OpenAI considered (and Altman later confirmed) seeking federal loan guarantees for domestic chip plants (www.axios.com), though not for data centers. Some states (like Texas and Ohio) have offered tax incentives and grants for data center construction. This environment lowers the effective cost of capital for all players.

Overall, by late 2025 the funding infrastructure is in place. Oracle and partners have drawn on equity, debt, and government support to finance at least **\$40–\$60 billion** in projects (Abilene, New Mexico, the Ohio plant, etc.). However, the total needs (approaching \$1T) far exceed these, so further financing (or build-out delays) can be expected.

Preliminary Outcomes & Current Status

What was the situation as of April 2026? (See the October 2026 update at the top for later developments.)

- **Construction Milestones:** The Abilene, TX flagship campus is **running on Oracle Cloud Infrastructure** and being built out to 1.2 GW (www.crusoe.ai). Crusoe topped out the final building by Q1 2026. Oracle is developing additional sites in Wisconsin and Michigan (openai.com). SB Energy (SoftBank) is developing the 1.2 GW Milam County, Texas campus with initial facilities entering service in 2026. SoftBank's Ohio factory conversion continues for modular data center unit production. The total Stargate program now encompasses nearly **7 GW of planned capacity** with over **\$400 billion** in committed investment.
- **Technology Deployment:** Oracle has delivered Nvidia GB200 racks to Abilene, with hundreds of thousands of GPUs now operational. OpenAI's Broadcom custom chips have completed the design phase and are on track for first deployment in H2 2026 (www.tomshardware.com). Next-generation Nvidia Vera Rubin chips are planned for new Stargate sites coming online in 2026–2027. AMD GPU commitments (6 GW) continue to supplement the hardware pipeline.
- **Financial Results:** Oracle's stock initially jumped 43% on Sep 2025 news but has since pulled back: shares were down ~24% year-to-date in early April 2026 amid broader AI bubble concerns (www.fool.com). However, Oracle's Q3 FY2026 results (Mar 2026) beat expectations: revenue up 22% to \$17.2B, cloud revenue up 44% to \$8.9B, sending shares up 10% after-hours (www.cnbc.com). Oracle's RPO reached \$553B. OpenAI's valuation rose from \$500B to **\$852 billion** following a \$122B funding round in March 2026 (www.cnbc.com), with monthly revenues then at ~\$2 billion.

- **Partnership Dynamics:** The competitive landscape has intensified. Amazon invested \$50B in OpenAI's March 2026 funding round (the largest single investor) (www.cnn.com), alongside its earlier \$38B cloud deal. Microsoft, despite retaining ~27% ownership of OpenAI, is increasingly pursuing AI independently, as shown by its leasing 900 MW from Crusoe in Abilene at a site OpenAI declined to expand (fortune.com). The two companies are literally neighbors on the same tract of land yet pursuing separate compute strategies. Meta's Oracle cloud deal (~\$20B) continues in discussions. Oracle now serves **all of the top five AI models** on its cloud platform.
- **Corporate Restructuring:** OpenAI completed its transition to a **Public Benefit Corporation** in late 2025, with a nonprofit Foundation holding a \$130B stake. The company was then preparing for a potential **Q4 2026 IPO**, though CEO Sam Altman and CFO Sarah Friar reportedly disagreed on timing (www.theinformation.com). OpenAI targets a \$1 trillion valuation at IPO and has communicated \$280 billion annual revenue projections by 2030 to prospective investors (www.indexbox.io).

Summary: By April 2026, Oracle and OpenAI have moved well beyond announcements into operational reality. The Abilene flagship is live, multiple new sites are under construction across the U.S., and the infrastructure for the \$300B contract is being built. OpenAI has secured its financial position with a record \$122B funding round and is generating \$2B/month in revenue. Oracle's cloud business was growing at 44%+ with \$553B in contracted backlog. The full 4.5 GW/year contract delivery is still ramping toward 2027, but the infrastructure foundation is now firmly in place.

Multi-Stakeholder Perspectives and Analysis

OpenAI's Perspective

From OpenAI's viewpoint, the Oracle deal is **necessary insurance** against ceilings on its growth. Sam Altman famously said he expects to spend "trillions" on infrastructure (www.datacenterdynamics.com) and that "30 GW is the target, roughly \$1.4 T in build-out" (www.axios.com). Legitimizing this scale has been a priority: the Oracle contract ensures compute for at least a fraction of that. OpenAI needed to signal to customers, investors, and employees that it won't stall due to lack of hardware.

Concerns for OpenAI include:

- **Affordability:** The money obligations are vast, so OpenAI must satisfy creditors and partners that it can pay Oracle without derailing R&D budgets. The plan appears to be to **grow into the deal**: rely on rapid revenue growth, drop cost curves, and continuous funding rounds (www.itpro.com). Internally, OpenAI likely runs models on "if we hit \$X revenue per user, then lighten cloud usage." Altman and team have said they will use commercial mechanisms (ads, upsells) to underwrite growth (www.itpro.com).
- **Technical Predictability:** Having a dedicated partner like Oracle (with co-located data centers) might simplify IT management versus multiple providers. But OpenAI also must manage the risk of hardware lock-in or underperformance. As a countermeasure, OpenAI is simultaneously investing in its own hardware (Broadcom chips) to hedge against vendor issues (www.reuters.com).

- **Regulatory and Public Pressure:** OpenAI is under public microscope for issues like AI safety. Committing to such extravagant spending draws scrutiny. Altman has defended it by pointing out that strong government and industry backing justifies scale (www.reuters.com). However, critics (including some technology journalists) have argued that the deal may be “smoke and mirrors” if OpenAI cannot justify the costs (theglobalvillager.com). The coming years will test those criticisms as progress (or lack thereof) becomes visible.

Oracle's Perspective

Oracle sees itself confirmed as a **central AI infrastructure provider**. Safra Catz, CEO when the deal was signed and executive vice chair since September 2025 (www.sec.gov), touted hundreds of billions in future cloud deals and sees OpenAI as a catalyst. On the positive side, Oracle's board and investors anticipate massive new revenue streams (www.techradar.com) (www.pcgamer.com). This is also a vindication of Oracle's cloud build-out strategy: their planned 4.5 GW expansion (added in July 2025 (www.reuters.com) (www.reuters.com)) was partly aimed at accommodating OpenAI workloads.

However, Oracle management also manages huge risk:

- **Engineering Execution:** Oracle has typically been slower-moving than Amazon/Azure. Building and filling multiple gigawatt-scale data centers is a steep challenge. Oracle has brought in outside “neocloud” partners (Crusoe, CoreWeave) to do the heavy lifting (www.reuters.com). The pace as of late 2025 appears on track but remains ambitious. If delays or cost overruns occur, Oracle might be on the hook to build facilities it cannot fill quickly.
- **Counterparty Risk:** Moody's flagged that Oracle is pegging its success on a few big clients (www.reuters.com). If OpenAI falters or renegotiates (unlikely publicly but possible), Oracle's gamble could backfire. On the other hand, if the AI boom continues, Oracle's cloud business could skyrocket. Oracle's going from a laggard to a lead player, which is a high-risk, high-reward scenario.
- **Competition and Deal Flow:** The news that Oracle is discussing large deals with Meta (www.reuters.com) and integrating with AWS/Google clouds demonstrates it is aggressively pursuing more AI customers. Each new contract (e.g. a rumored \$20B Meta deal) reinforces Oracle's momentum and dilutes risk concentration. In theory, serving multiple hyperscalers would keep multiple “eggs in the basket”.

Microsoft and Other Cloud Competitors

Microsoft, the longtime OpenAI partner, finds itself in an increasingly complex position by early 2026. While it retains ~27% ownership of OpenAI (worth ~\$135 billion post-restructuring) and continues to capture a share of OpenAI's business (Azure usage, GitHub Copilot, etc.), the relationship has measurably cooled. The most telling sign came in March 2026, when **Microsoft leased 900 MW of data center capacity from Crusoe in Abilene, Texas**, at a site that OpenAI had declined to expand, making the two companies literal neighbors on the same tract of land yet pursuing separate compute strategies (fortune.com). Crusoe is building two new “AI

factory" buildings and a 900 MW on-site power plant for Microsoft, right next to where it built OpenAI and Oracle's campus. Microsoft relaxed its exclusivity clauses in early 2025 to allow OpenAI to pursue new cloud compute sources (www.reuters.com), but the extent of the drift has surprised observers.

Amazon has deepened its OpenAI ties considerably. Beyond the \$38B cloud computing deal announced in November 2025 (www.reuters.com), Amazon became the **largest single investor** in OpenAI's March 2026 funding round at **\$50 billion** (though \$35B is contingent on IPO or AGI milestones) (www.cnbc.com). This positions Amazon as both a cloud provider and major equity stakeholder in OpenAI, a dual relationship that mirrors Microsoft's earlier arrangement.

Google Cloud's immediate role in OpenAI's plans remains limited. While Google's own LLM efforts (Gemini) position it as a competitor, Oracle has integrated OCI with Google Cloud services (www.reuters.com), hinting at cooperative channels. Meanwhile, competitors in the AI space are also making major moves: Anthropic reportedly passed OpenAI in run-rate revenue at \$30B ARR by April 2026, with projected training spend through 2030 about a quarter of OpenAI's (www.the-ai-corner.com), and its quarterly revenue exceeded OpenAI's for the first time in Q2 2026, at \$11.5 billion against \$6.7 billion (www.reuters.com), underscoring that OpenAI's massive infrastructure bet faces competitive pressure from more capital-efficient rivals.

National and Economic Perspectives

The Oracle–OpenAI project is closely entwined with U.S. industrial strategy. It enjoyed political backing (including direct encouragement by President Trump (www.reuters.com)) as a way to ensure American leadership in AI over China. The commitment of hundreds of billions is unprecedented in tech and has raised eyebrows among policymakers. For example, some Democrats and climate/environment advocates question the sustainability of such energy-hungry projects (though Oracle has highlighted green power and closed-loop cooling).

Economically, the project could drive regional development: Abilene's data center promises thousands of jobs (apnews.com), and other sites (NM, OH) are chosen partly for economic transitions (Lordstown's EV plant being repurposed). However, critics warning of "jobless" AI expansion note that most of the work is high-tech and not as labor-intensive as old industries, and much of the funding flows to hardware makers.

Overall, the deal is seen by allies as a bold move to "pull ahead" in AI infrastructure, but by skeptics as a risky mega-bet that assumes perpetual growth. The truth will depend on execution in 2026–2028.

Case Studies and Comparisons

To put the Oracle–OpenAI project in context, we consider other major AI/cloud infrastructure case studies:

- **Meta's Google Cloud Deal (Aug 2025):** Meta signed a \$10+ billion, six-year cloud contract with Google to support its services (www.reuters.com). Like Oracle's deal, it indicates that even the largest tech companies are outsourcing a portion of their compute to specialized cloud providers. The Meta–Google deal involved running core Meta applications in Google's data centers (a partial move). In comparison,

Oracle–OpenAI is purely for AI training/deployment. The existence of these deals shows a trend: big tech stacks need to diversify their infrastructure and are willing to enter multi-year contracts with hyperscalers.

- **Amazon–OpenAI & AWS Pivot (Nov 2025):** In November 2025, AWS won a reported **\$38 billion** contract to supply compute for OpenAI (www.reuters.com). This is dwarfed by Oracle’s \$300B, but noteworthy because Amazon had lagged in AI until then. The seven-year AWS deal covers hundreds of thousands of Nvidia GPUs (www.aboutamazon.com); it came as AWS was scaling Project Rainier, its Trainium2 cluster built with Anthropic (www.aboutamazon.com), and the rollout of its own AI chip Trainium (www.reuters.com). The timing suggests Oracle’s deal spurred Amazon to land a slice of the OpenAI pie and reassert itself in the AI cloud race. The AWS arrangement is not detailed publicly (likely it includes storage, network, and some chips) but it highlights that the Oracle deal shined a light on unmet demand that competitors raced to fill.
- **IBM/Google HPC Example:** Although not AI-specific, large-scale government HPC deals provide a contrast. For example, the U.S. Department of Energy’s Exascale computer contracts (Frontier and its peers) were on the order of ~\$600 million each (www.reuters.com), dwarfed by what Oracle–OpenAI spends in a day. Metaphorically, IBM and Google have undertaken large HPC procurement for air force or science (e.g. IBM’s 2018 Summit supercomputer at Oak Ridge (www.ornl.gov)), but those budgets (\$1–2B range) are tiny compared to modern AI finance. The Oracle deal essentially defines a new “order of magnitude” for computing outlays.
- **Chipmaker Partnerships:** OpenAI’s Broadcom partnership is akin to Google’s TPU development or Apple designing its own chips. Many tech companies (Google with TPUs, Amazon with Trainium/GPU deals, Meta with MTIA (ai.meta.com)) have attempted custom accelerators for performance/cost gains. OpenAI’s 10 GW/term joint project with Broadcom is the largest such example by far (www.reuters.com). It confirms a trend where top AI firms hedge against reliance on external hardware by creating their own silicon.

Table 5 compares major AI infrastructure commitments (beyond Oracle–OpenAI) to show the scale:

Project	Participants	Commitment	Value / Cost	Purpose
AWS/OpenAI cloud deal	Amazon, OpenAI	Hundreds of thousands of Nvidia GPUs, 7 years	~\$38 billion (specifically for OpenAI) (www.reuters.com)	Cloud compute for OpenAI services
Meta/Google Cloud	Meta (Facebook), Google	6-year cloud contract	~\$10+ billion	Host “non-AI” compute and some AI
Microsoft/OpenAI	Microsoft, OpenAI	Multi-year strategic alliance	~\$50–250+ billion (estimated)*	Azure compute + co-development
OpenAI/Broadcom	OpenAI, Broadcom	10 GW custom chips	Perhaps ~\$500–600B (est.) (www.reuters.com)	Co-develop on-prem AI accelerators
SoftBank/OpenAI (Stargate EQ)	SoftBank, OpenAI (and Oracle)	~\$40B funding commitment	\$40 billion via equity/loans (www.reuters.com)	Finance for data centers and chips

Project	Participants	Commitment	Value / Cost	Purpose
Oracle/OpenAI	Oracle, OpenAI	4.5 GW/year for 5 years	\$300 billion	Cloud compute lease
Meta (AI Compute)	Meta AI (LLaMA, etc) + web	Not disclosed (internal)	Opaque (internal budgets)	Build Meta's own AI hardware

* *Microsoft's investments in OpenAI span cash infusions and computing credits. Exact figure is not public, but media estimates years of 10s of billions.* This table underscores that Oracle–OpenAI is the largest **single** financial commitment focused on AI compute. OpenAI's incremental \$250 billion Azure commitment to Microsoft (October 2025) comes closest (blogs.microsoft.com).

Data Analysis and Evidence

Key quantitative highlights gathered from reporting and filings:

- **Revenue vs. Compute Costs:** OpenAI's revenue trajectory has improved substantially since the deal was announced. From ~\$1 B/month in mid-2025, revenues have grown to approximately **\$2 B/month** by early 2026 (openai.com), reaching \$25B ARR. A \$60 B/year compute bill equates to ~\$5 B/month. Thus, paying Oracle would still consume ~2.5× current revenues, a significant improvement from the 5× gap in mid-2025, but still a deficit that requires external funding. The \$122B raised in March 2026 provides a substantial runway, and OpenAI's new advertising business (projected \$2.5B in 2026, targeting \$100B by 2030) could narrow the gap further. Nonetheless, OpenAI is projected to lose ~\$14B in 2026 alone, emphasizing that profitability remains years away. By September 2026, OpenAI told investors its annualized revenue was almost \$50 billion, or roughly \$4 billion a month against the ~\$5 billion a month implied by the Oracle contract (www.reuters.com).
- **Stock Market Impact:** Oracle's stock (NYSE: ORCL) rose ~43% intraday when the deal was announced (www.tomshardware.com) (www.pcgamer.com). Analyst consensus after the dust settled was mixed; some upgraded on anticipated growth, others argued the price already accounted for future earnings. Oracle's market cap (roughly \$680B pre-news) gained \$244B to \$922B that day (www.cnbc.com), briefly making Ellison the richest person (www.tomshardware.com) (www.pcgamer.com).
- **Job Creation:** The five new Stargate centers were projected to generate ~25,000 construction jobs (www.reuters.com). For Abilene specifically, permanent positions are in the low thousands (apnews.com). While substantial, these numbers are tiny relative to America's total workforce (0.01%). However, the economic multipliers (electricity, local services) may be significant locally.
- **Hardware Scale:** Abilene's initial build used ~1.2 GW of capacity (www.reuters.com). Four more such sites at 1–2 GW each would bring the system to 7–10 GW across all phases (matching the Stargate goal (www.reuters.com)). Thus far, ~2 GW is confirmed built/upcoming (Abilene + initial Texas site) and 4.5 GW more announced (www.reuters.com) (www.reuters.com). The plan to add 4.5 more GW

(bringing total to 5+ GW) by end of 2025 was publicized (www.reuters.com), and Reuters reported commitment of 4.5 GW additional in mid-2025 (www.reuters.com). These hard numbers provide evidence of the actual build-up trajectory.

- **Investment vs. Revenue Forecasts:** OpenAI achieved its 2025 revenue target of ~\$20B, then grew to \$25B ARR by February 2026. The company has communicated to prospective IPO investors a target of **\$280 billion in annual revenue by 2030** (www.indexbox.io), implying over 10× growth from current levels. The capital commitments (tens of billions per year) still require revenues many times the current level, and projected long-term spending totals ~\$115 billion through 2029. The March 2026 funding round (\$122B) dramatically bolstered OpenAI's war chest, with investors including Amazon (\$50B), Nvidia (\$30B), SoftBank (\$30B), and Microsoft. Total equity raised to date now exceeds **\$170 billion**, substantially narrowing the gap versus five-year compute costs, though the gap remains significant if all commitments are fulfilled out-of-pocket without further revenue growth.
- **Global Infrastructure Spending:** Industry reports estimate that global data center spending will reach ~\$2.9 trillion by 2028 (www.techradar.com). The Oracle–OpenAI infrastructure (plus related stargate projects) is on that order. It alone may represent >10% of total global data center capex in this period. This underscores how the AI boom is reshaping entire industries (chipmakers, utilities, real estate).

Discussion: Implications and Future Directions

National Competitiveness and Security

The national significance is high. U.S. officials see these investments as critical to staying ahead in AI vis-à-vis China. By bringing such projects onshore, the initiative averts a scenario where AI training runs on foreign soil (where data sovereignty and geopolitical risk exist). The Stargate sites are located in “friendly” U.S. territories, often co-locating with renewable power sources. For example, Abilene's site leveraged local wind and solar power and even a natural-gas plant for grid stability (apnews.com). This model may be extended: loan covenants or permits could require efficiency standards.

On the other hand, the extreme concentration of compute in a few U.S. sites could raise national security concerns (e.g. a natural disaster or cyberattack on these clusters could be catastrophic). The government's role (president's emergency orders) hints they view this infrastructure as strategic as a power grid.

Cloud Industry Impact

For the cloud market, Oracle's ascendance (if realized) would alter dynamics. Before this, AWS, Azure, and Google dominated. Oracle used to be a distant fourth place in market share. A \$300B deal vaults OCI into instant relevance for generative AI workloads. It may induce second-order effects: specialized services (AI ops software, custom AI chips) will be offered by Oracle to capture more of that spending.

Competing clouds cannot rely solely on traditional enterprise workloads; this competition shows they must pivot to support AI. Microsoft announced, as a result of these developments, that it is now allowing expanded use of its compute by OpenAI (and possibly other AI firms) (www.reuters.com) (www.reuters.com). AWS responded by investing \$11B in new DCs and cutting prices on its Trainium chips (www.reuters.com). Google is building GPT-like models to serve its cloud customers. In short, Oracle's deal has catalyzed an industry-wide acceleration in AI infrastructure investment.

Economic and Environmental Considerations

Economically, such multibillion-dollar deals can stimulate jobs and local economies in chosen sites, but raise questions about sustainability. Data centers are electricity-hungry: Oracle itself will need “4.5 GW of power” year-round (www.tomshardware.com). That is a substantial fraction of Texas' entire generation capacity. Environmental assessments (as reported for Abilene) are active concerns. Analysts point out as AI data center demand grows, utility companies may need to upgrade grids and build new generation (likely gas or nuclear), a public cost often overlooked in the initial investments.

There is also a hope that AI demand could help renewables, by providing flexible (shiftable) loads and additional revenue streams for wind/solar projects. The reported use of gas turbine backup at Abilene suggests a hybrid approach: primary renewables + diesel/gas peak coverage. Over time, as more states vie to host these facilities, incentives (e.g. low-cost wind in Texas, geothermal in NM, hydro in Washington) will shape the geography of AI hubs.

Future Outlook

Given current trends as of April 2026, we expect:

- **Continued Mega-Spending with IPO Catalyst:** OpenAI's leadership remains publicly committed to this path. An IPO, once targeted for **Q4 2026** at a \$1 trillion valuation and now pushed beyond 2026, would provide a new source of capital and public market accountability. OpenAI has communicated \$280B annual revenue projections by 2030 to prospective investors, implying enormous growth expectations. New Stargate sites in Michigan and Wisconsin are under development, with next-generation Nvidia Vera Rubin chips planned for 2026–2027 deployment.
- **Custom Hardware Arriving Sooner Than Expected:** OpenAI's Broadcom-designed AI inference chips have already **completed the design phase** and are on track for first data center deployment in **H2 2026**, ahead of earlier projections. The custom chips are optimized for the Transformer architecture and o1-series inference workloads. Broadcom has also expanded its custom chip business to Google and Anthropic (www.cnbc.com), confirming that custom AI silicon is now an industry-wide movement that may reduce Nvidia's dominance over time.
- **Oracle's Financial Tightrope:** Oracle raised \$43 billion of debt and \$5 billion of equity in fiscal 2026 (www.sec.gov), then sold \$20 billion of common stock in Q1 FY2027 (www.prnewswire.com). Fiscal 2026 capex was \$55.7 billion, and fiscal 2027 capex could reach \$95 billion. S&P's July 2026 downgrade left Oracle on the cusp of a junk rating, and Moody's has said leverage could approach 5x

temporarily (www.reuters.com). Barclays had earlier warned Oracle's debt could approach junk status. If Oracle's cloud revenue keeps growing at the 62% pace of Q1 FY2027, the gamble may pay off. But if AI demand slows, the leverage becomes dangerous. Oracle's success now depends not just on OpenAI but on its ability to serve all major AI model providers.

- **Competitive Landscape Shifting:** Anthropic's reported surge past OpenAI, first in run-rate revenue in April 2026 and then in quarterly revenue in Q2 2026, with projected training spend far below OpenAI's, poses a fundamental question about the value of massive infrastructure bets. If more efficient model training and inference approaches gain traction, the economics of multi-gigawatt data centers could shift. However, as AI applications expand to agents, multimodal systems, and real-time inference at scale, compute demand may well outpace efficiency gains.
- **Regulatory Scrutiny:** Antitrust regulators globally are monitoring these alliances. The concentration of AI infrastructure in a few U.S. locations and companies raises competition and national security questions. The EU's AI Act implementation and potential data localization requirements could complicate globally scaled AI infrastructure plans.
- **Innovation and Efficiency:** The massive capital investment creates strong incentives to innovate on efficiency. Crusoe's pioneering use of on-site gas power for AI compute, combined with liquid cooling advances and custom silicon, points to a future where cost-per-FLOP continues declining rapidly. OpenAI's advertising pilot (\$100M ARR in under 2 months) also suggests new monetization paths that could help justify infrastructure costs.

Conclusion

The **Oracle–OpenAI \$300B deal** marks a watershed in the technology landscape. It epitomizes the **AI arms race**: unprecedented amounts of capital being deployed to secure computing power. By October 2026, the project has moved well beyond announcements into operational reality: the flagship Abilene campus is live on Oracle Cloud Infrastructure and being built out to 1.2 GW, additional sites are under construction across Michigan, Wisconsin, New Mexico, and multiple Texas locations, and the Stargate program has expanded to nearly 7 GW of planned capacity with over \$400 billion in investment.

OpenAI's financial position has strengthened considerably since the deal was announced, with revenue growing from ~\$12B ARR in mid-2025 to ~\$25B ARR in early 2026 and a record \$122 billion funding round at an \$852 billion valuation. The company has completed its restructuring into a Public Benefit Corporation and has said it will not go public in 2026. Oracle's cloud business is surging (62% cloud revenue growth and \$664B in contracted backlog as of August 2026), though financial risks remain with \$55.7B of fiscal 2026 capex, up to \$95B planned for fiscal 2027, and mounting debt pressures that have drawn Moody's warnings and an S&P downgrade.

Our research indicates both the *immense potential* and the *serious risks* of this venture. Key developments since the original analysis include: Microsoft's increasing distance from OpenAI (now leasing its own 900 MW site next to the Stargate campus rather than serving as OpenAI's exclusive cloud provider), Amazon's emergence as a major OpenAI investor and cloud supplier, Anthropic's revenue surge past OpenAI (a run rate above \$30B by April 2026 and higher quarterly revenue than OpenAI by Q2 2026), and the design completion

of OpenAI's custom Broadcom chips ahead of H2 2026 deployment. Competition from more capital-efficient rivals like Anthropic raises questions about whether massive infrastructure spending alone guarantees market leadership.

Nonetheless, one thing is clear: the Oracle–OpenAI partnership has already **reshaped expectations** of what infrastructure support for AI can look like. It has set a new benchmark for ambition in the field, against which all future deals and strategic plans will be measured. As the \$300B contract begins to deliver compute in 2027, the next 12–18 months will prove whether this grand experiment pays off in transformative AI capabilities or whether the AI industry's capital efficiency curve renders such massive bets unnecessary.

Key Citations: Industry reporting (Reuters, AP, Tech publications) has been used throughout to substantiate facts about deal terms, power requirements, finance figures, and project status (www.reuters.com) (www.tomshardware.com) (www.techradar.com) (www.datacenterdynamics.com) (www.reuters.com) (www.reuters.com) (apnews.com) (www.reuters.com). Technical publications highlight the scale of AI's energy demands (arxiv.org) (arxiv.org). All views are grounded in these sources and data.

Source links

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

1. <https://www.prnewswire.com/news-releases/oracle-announces-q1-results-driven-by-triple-digit-growth-in-cloud-infrastructure-revenues-302875728.html>
2. <https://www.reuters.com/business/oracle-corp-goes-high-stakes-ratings-gamble-ai-strategy-2026-08-04/>
3. <https://www.reuters.com/business/oracle-openai-end-plans-expand-texas-data-center-site-bloomberg-news-reports-2026-03-06/>
4. <https://www.reuters.com/business/oracle-cites-force-majeure-shield-itself-controversial-data-center-bloomberg-2026-09-24/>
5. <https://fortune.com/2026/09/12/sam-altman-openai-ipo-delay-ill-advised-moment-safety-concerns/>
6. <https://www.reuters.com/business/openais-annualized-revenue-20-billion-less-than-previously-signaled-ft-reports-2026-10-08/>
7. <https://www.datacenterdynamics.com/en/news/openai-signs-300bn-cloud-deal-with-oracle-report/>
8. <https://www.techradar.com/pro/openai-bets-usd300-billion-on-oracle-contract-to-power-artificial-intelligence-expansion-despite-ongoing-losses>
9. <https://www.reuters.com/technology/artificial-intelligence/openai-softbank-oracle-invest-500-bln-ai-trump-says-2025-01-21/>
10. <https://www.reuters.com/sustainability/land-use-biodiversity/altman-touts-trillion-dollar-ai-vision-openai-restructures-chase-scale-2025-10-29/>
11. <https://www.tomshardware.com/tech-industry/openai-signs-contract-to-buy-usd300-billion-worth-of-oracle-computing-power-over-the-next-five-years-company-needs-4-5-gigawatts-of-power-enough-to-power-four-million-homes>
12. <https://www.cnn.com/2026/03/26/openai-ads-pilot-tops-100-million-in-arr-in-under-2-months.html>
13. <https://openai.com/index/a-business-that-scales-with-the-value-of-intelligence/>
14. <https://www.cnn.com/2026/03/31/openai-funding-round-ipo.html>
15. <https://www.reuters.com/business/moodys-flags-risk-oracles-300-billion-recently-signed-ai-contracts-2025-09-17/>
16. <https://investor.oracle.com/investor-news/news-details/2026/Oracle-Announces-Fiscal-Year-2026-Third-Quarter-Financial-Results/default.aspx>
17. <https://www.crusoe.ai/resources/newsroom/crusoe-announces-flagship-abilene-data-center-is-live>

18. <https://openai.com/index/five-new-stargate-sites/>
19. <https://www.reuters.com/business/oracle-buy-40-billion-nvidia-chips-openais-us-data-center-ft-reports-2025-05-23/>
20. <https://www.tomshardware.com/openai-broadcom-to-co-develop-10gw-of-custom-ai-chips>
21. <https://intuitionlabs.ai/articles/gemini-3-healthcare-applications>
22. <https://intuitionlabs.ai/articles/kimi-k2-technical-deep-dive>
23. <https://arxiv.org/abs/2505.05794>
24. <https://intuitionlabs.ai/articles/claude-opus-4-5-healthcare-pharma-ai>
25. <https://arxiv.org/abs/2509.07218>
26. <https://arxiv.org/abs/2506.14551>
27. <https://intuitionlabs.ai/articles/chatgpt-plans-comparison>
28. <https://www.reuters.com/business/media-telecom/openais-annualized-revenue-hits-10-billion-up-55-billion-december-2024-2025-06-09/>
29. <https://www.reuters.com/business/openai-hits-12-billion-annualized-revenue-information-reports-2025-07-31/>
30. <https://www.reuters.com/commentary/breakingviews/openais-valuation-is-tech-booms-key-man-risk-2025-08-07/>
31. <https://www.investing.com/news/company-news/openai-projects-25-billion-in-ad-revenue-for-2026-93CH-4605384>
32. <https://www.reuters.com/world/asia-pacific/softbank-group-posts-166-billion-net-profit-q2-2025-11-11/>
33. <https://techcrunch.com/2026/03/31/openai-not-yet-public-raises-3b-from-retail-investors-in-monster-122b-fund-raise/>
34. <https://openai.com/index/evolving-our-structure/>
35. <https://www.indexbox.io/blog/openai-targets-q4-2026-ipo-with-1-trillion-valuation-goal/>
36. <https://fortune.com/2026/09/16/openai-ipo-sam-altman-vc-funding-valuation-1-2-trillion/>
37. <https://finance.yahoo.com/news/reported-openai-ipo-later-test-165344346.html>
38. <https://www.pcgamer.com/hardware/oracles-usd300-billion-deal-with-openai-briefly-knocked-elon-musk-off-the-worlds-richest-person-spot-replaced-instead-by-the-companys-head-honcho-larry-ellison/>
39. <https://www.cnbc.com/2026/03/10/oracle-orcl-q3-earnings-report-2026.html>
40. <https://www.fool.com/investing/2026/04/06/oracle-shares-are-down-24-so-far-in-2026-amid-ai/>
41. <https://www.reuters.com/technology/microsoft-relaxes-data-center-grip-openai-amid-500-bln-joint-venture-2025-01-22/>
42. <https://www.reuters.com/business/media-telecom/openai-oracle-softbank-plan-five-new-ai-data-centers-500-billion-stargate-2025-09-23/>
43. <https://www.reuters.com/technology/openais-biggest-data-center-secures-116-billion-funding-wsj-reports-2025-05-21/>
44. <https://www.reuters.com/business/openai-taps-broadcom-build-its-first-ai-processor-latest-chip-deal-2025-10-13/>
45. <https://www.reuters.com/business/oracle-talks-with-meta-20-billion-ai-cloud-computing-deal-bloomberg-reports-2025-09-19/>
46. <https://www.reuters.com/business/retail-consumer/amazons-38-bln-openai-deal-shows-it-is-no-longer-an-ai-laggard-2025-11-04/>
47. <https://www.reuters.com/technology/openai-oracle-sign-300-billion-computing-deal-wsj-reports-2025-09-10/>
48. <https://www.reuters.com/business/media-telecom/softbank-invest-3-billion-ohio-factory-openai-data-center-information-reports-2025-11-20/>
49. <https://apnews.com/article/0b3f4fa6e8d8141b4c143e3e7f41aba1>
50. <https://www.reuters.com/business/oracle-openai-add-45-gigawatts-data-center-capacity-stargate-venture-2025-07-22/>

51. <https://blogs.microsoft.com/blog/2025/10/28/the-next-chapter-of-the-microsoft-openai-partnership/>
52. <https://www.axios.com/2025/10/28/openai-1-trillion-altman>
53. <https://www.reuters.com/technology/artificial-intelligence/behind-500-billion-ai-data-center-plan-us-startups-jockey-with-tech-giants-2025-01-23/>
54. <https://www.itpro.com/business/business-strategy/openai-1-trillion-infrastructure-investment-oracle-broadcom-amd>
55. <https://www.reuters.com/business/media-telecom/softbanks-58-billion-nvidia-stake-sale-stirs-fresh-ai-bubble-fears-2025-11-11/>
56. <https://www.reuters.com/business/finance/banks-lend-18-billion-oracle-tied-data-center-project-bloomberg-news-reports-2025-11-07/>
57. <https://www.reuters.com/business/finance/banks-talks-lend-38-billion-fund-sites-openai-ft-reports-2025-11-28/>
58. <https://www.datacenterdynamics.com/en/news/crusoe-tops-out-final-building-at-openai-stargate-data-center-campus-in-abilene-texas/>
59. <https://fortune.com/2026/03/27/microsoft-texas-data-center-open-ai-former-partner-cloud-provider/>
60. <https://www.oracle.com/news/announcement/related-digital-oracle-openai-walbridge-and-governor-whitmer-celebrate-construction-of-stargate-campus-in-saline-township-2026-06-01/>
61. <https://www.reuters.com/technology/ai-data-centre-startup-crusoe-raising-138-billion-latest-funding-round-2025-10-23/>
62. <https://www.reuters.com/business/softbank-openais-stargate-aims-building-small-data-center-by-year-end-wsj-2025-07-22/>
63. <https://www.cnbc.com/2026/04/06/broadcom-agrees-to-expanded-chip-deals-with-google-anthropic.html>
64. <https://ir.amd.com/news-events/press-releases/detail/1260/amd-and-openai-announce-strategic-partnership-to-deploy-6-gigawatts-of-amd-gpus>
65. <https://www.theinformation.com/briefings/exclusive-openai-cfo-questions-2026-ipo-readiness>
66. <https://theglobalvillager.com/2025/09/19/the-openai-oracle-deal/>
67. <https://www.sec.gov/Archives/edgar/data/1341439/000119312525210089/d921500dex991.htm>
68. <https://www.reuters.com/business/openai-oracle-deepen-ai-data-center-push-with-45-gigawatt-stargate-expansion-2025-07-22/>
69. <https://www.the-ai-corner.com/p/anthropic-30b-arr-passed-openai-revenue-2026>
70. <https://www.aboutamazon.com/news/aws/aws-open-ai-workloads-compute-infrastructure>
71. <https://www.aboutamazon.com/news/aws/aws-project-rainier-ai-trainium-chips-compute-cluster>
72. <https://www.ornl.gov/news/ornl-launches-summit-supercomputer>
73. <https://ai.meta.com/blog/next-generation-meta-training-inference-accelerator-AI-MTIA/>
74. <https://www.cnbc.com/2025/09/10/oracle-stock-cloud-backlog-ai.html>
75. https://www.sec.gov/Archives/edgar/data/1341439/000119312526265848/orcl-ex99_1.htm

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.