

Case: Oracle and the Stargate Project 2025

Dr. Lifeng Geng

Associate Professor

Business Program

Lakehead University, Orillia Campus

500 University Avenue, Orillia, ON, L3V 0B9

Canada

Abstract

In 2025, Oracle Corporation embarked on one of the most ambitious investments in modern technology history: the Stargate Project, a multi-year, multi hundred-billion-dollar initiative to build hyperscale AI data centers primarily to support OpenAI's frontier model training and inference needs. The project placed Oracle at the center of the rapidly evolving AI infrastructure ecosystem and represented a defining strategic gamble for the company and its founder, Larry Ellison.

This case examines Oracle's strategic position, motivations, financial risks, and competitive dynamics surrounding the Stargate Project. Students are asked to evaluate whether Stargate represents a rational long-term platform bet or an overextension driven by legacy considerations and AI exuberance.

Keyword: competitive environment, technological paradigm shift, strategic position, business model, financial risk, AI infrastructure, Oracle, Stargate

1. Introduction

In mid-2024, Oracle's Vice Chairman and former CEO, Larry Ellison received a private message on LinkedIn from OpenAI's COO, Brad Lightcap, probing for a potential partnership aimed at building next-generation AI infrastructure for OpenAI's Artificial General Intelligence (AGI) model training. This LinkedIn outreach culminated in a dinner where Altman and Ellison discussed the staggering compute requirements for future AI models, famously described by Ellison as "begging" for more GPU capacity. Larry Ellison faces the most critical strategic decision of his career: whether to commit up to \$300 billion to "Stargate," a massive AI data center initiative designed to support OpenAI's next generation of frontier AI model training.

Oracle was not the first to be contacted by OpenAI. As matter of fact, OpenAI initially reached out to Microsoft about this potential partnership. As OpenAI's longest and largest investor and primary cloud partner (Azure), Microsoft conducted a due diligence analysis and declined to fully finance and own this dedicated infrastructure. The core reason was concentration risk: anchoring such a massive, fixed-asset investment to the technical roadmap and financial health of a single entity, OpenAI, represented an unacceptable asymmetry of risk and reward.

Oracle has dominated the Relational Database Management Systems (RDBMS) segment of enterprise software business for decades. It has been the main service provider of mission-critical transactional workloads (OLTP), systems of record (financials, customer data, inventory, billing), high availability, strong consistency, and regulatory compliance, and large-scale, multi-user enterprise environments for governments, financial institutions, telecoms, airlines, utilities, and other large enterprises.

Historically, Oracle has responded late to critical, disruptive, paradigm shift in IT technologies, including the internet platform era (dominated by Google and Amazon), mobile computing (Apple/Android), and the public cloud (AWS, Azure, Google Cloud). As AI is widely accepted as a transformative, fourth-industrial-revolution technology that may simulate human intelligence using algorithms to automate complex, cognitive tasks across industries, Ellison viewed the Stargate partnership as a potential path to regain centrality in the IT industry after missing previous platform shifts. Meanwhile, the Stargate initiative posed immense capital intensity and significant risk to the company's balance sheet and legacy. Ellison must weigh the risk and reward of the initiative before reaching a final decision.

2. Company Background: Oracle

Before Oracle, Ellison was a brilliant but restless college dropout who worked as a programmer for various companies, including Ampex Corporation. There, he worked on a database project for the CIA codenamed "Oracle." At Ampex, Ellison worked with Bob Miner (a supremely talented senior programmer) and Ed Oates (a visionary systems designer). The trio left Ampex to join a startup, but when that failed, they found themselves as a team looking for their next project.

In 1977, Ellison, Miner, and Oates saw an IBM's paper titled "SEQUEL 2: A Unified Approach to Data Definition, Manipulation, and Control", which described a relational database model based on Edgar F. Codd's (1970) relational model. They realized that that this was a huge opportunity since no company was building a commercial relational database based on this new model. On June 16, 1977, they founded Software Development Laboratories (SDL). Ellison, then aged 33, put up \$1,200 of his own money for a 60% share, and became CEO. Bob Miner (the lead engineer) and Ed Oates each got 20%. Bruce Scott (the first hired employee) would later get a small stake. Their first contract was with the CIA. The agency needed a relational database and had seen Ellison's prior work. The project's code name was, once again, "Oracle." The product they built was named after this contract.

Oracle built the first commercially available SQL-based relational database. Unlike competitors, Oracle focused on commercial customers, not academia, or individual customers, and its product supported multiple operating systems, avoiding hardware lock-in.

Oracle's initial revenue growth was explosive, driven by the rapid adoption of its relational databases by governments, financial institutions, and other large enterprises. In 1982, the company was renamed Oracle Systems Corporation, after its flagship database product. It went public in 1986 and entered the Fortune 500 list in 1991.

From the beginning, Ellison's aggressive sale style was built into Oracle's sales-driven, "sell first, build later" culture (Wilson, 1997). For instance, Oracle deliberately called their first commercial product "Version 2", which was viewed as a masterstroke in marketing psychology by Ellison to imply it was more mature and that a non-existent "Version 1" had already been debugged. Ellison's approach was legendary and ruthless. He sold the vision of a unified database for all of a company's data, often promising features that didn't exist yet. Salespeople sold future capabilities ("vaporware") to win contracts, putting immense pressure on engineers to deliver. This hyper-sales driven culture nearly destroyed the company during the 1990s. Oracle shipped buggy, incomplete software. Customers rebelled, and financial reporting irregularities led to an SEC investigation. The company almost went bankrupt, with its stock plummeting 80%. This forced a painful but necessary corporate restructuring: Ellison brought in professional managers, stabilized finances, and focused on product quality.

In the 2000s, Oracle recognized that databases alone limited their growth and control over business workflows would facilitate their future growth with increased customer lock-ins. It engaged in a series of aggressive acquisitions to transform itself into a full enterprise software vendor in the areas of database, middleware, ERP, CRM, and HCM. In 2005, it acquired PeopleSoft for HR and ERP solutions. In 2006, it purchased Siebel Systems for CRM software. In 2007, it took over Hyperion for Enterprise performance management system and 2008, BEA Systems for middleware systems. With these acquisitions, it provided a complete, bundled enterprise software solution to an expanded addressable market, maintained a locked-in, long-term, customer base, and successfully positioned itself against SAP and IBM.

In 2010, Oracle acquired Sun Microsystems for approximately \$7.4 billion. It gained control over Java, Solaris, MySQL (open-source database), and the SPARC processors. With this acquisition, it established its entire technology stack: Hardware → OS → Database → Applications. Nevertheless, the hardware margins were much lower than software and its SPARC ecosystem never took off. Consequently, enterprise software continued to be its main business and hardware was only a strategic complement, not a core revenue driver.

Oracle's trajectory during the 1990s shaped its conservative financial management with a focus on recurring revenue and stable profitability. This, to some extent, led to its slow response to the disruptive technological change in the IT field. In the Internet platform era, it lost to Google and Amazon. In the mobile computing era, it left ground to Apple's ios and Google's Android ecosystems. This trajectory continued with the major shift from enterprise servers to public cloud. Ellison responded late to cloud computing and was often publicly dismissive of it. This was exacerbated by the leadership change in 2014. In that year, Larry Ellison stepped down as CEO and Safra Catz became Oracle's co-CEO. Safra Catz is one of the most powerful and distinctive executives in the technology industry. Her background is deeply rooted in finance and investment banking, which has profoundly shaped her leadership style at Oracle. As a talented financial expert, her mantra was margins and profitability. She instilled intense financial discipline across Oracle, constantly focusing on cost control, efficiency, and shareholder value. She was known for her mastery of financial details and her relentless focus on the bottom line. She hated projects that stimulate growth with burned expenditures, like cloud computing. Under her leadership and Ellison's influence, the ten critical ten years became Oracle's lost decade. Although Oracle lived nicely with decent profitability, it yielded leadership in cloud computing to AWS, Azure, Google Cloud and witnessed significant decline in its customer base.

Having lost the last three paradigm shifts, the pain of further loss weighs heavily in Ellison's evaluation of the Stargate project. Will Oracle lose again in the fast-progressing AI revolution? Stargate was thus not merely a matter of corporate strategy; it was a high-stakes legacy project for an 81-year-old founder determined to rewrite the final chapter of his career.

3. Oracle's Main Lines of Products/Services and Recent Financial Performance

Oracle's business can be classified into three main categories of product/services:

1. Cloud services, which include both cloud infrastructure (IaaS) and cloud applications (SaaS). Oracle Cloud Infrastructure (OCI) offers over 200 AI and cloud services across on-premises, edge, and public cloud environments, supporting sovereign AI and low-latency requirements. Cloud applications (SaaS) provide a comprehensive suite including ERP (Fusion Cloud ERP, NetSuite Cloud ERP), HCM, CRM, EPM, SCM, CX Commerce, Oracle AI Agent Studio, and Oracle Fusion Data Intelligence.
2. Enterprise Software Solutions, which includes Oracle's foundational Oracle relational database, multi-cloud database services, Oracle Fusion Middleware offerings, and other software in business intelligence, identity management, development software, and operating systems.
3. Enterprise hardware products from Sun Microsystems acquisition.

For the fiscal year of 2025 (ended May 31, 2025), Oracle's cloud services reached \$24.506 billion in revenue, which represented 42.7% of its total revenue, grew 23.93% year-over-year. Revenue from its enterprise software solutions was \$24.724 billion, representing 43.07% of its total revenue, and grew less than one percent over the year of 2024. Its services business generated \$5.23 billion in revenue (9.12% of total revenue) and decreased by 3.65% over the year of 2024. Oracle's hardware business earned \$2.94 billion for the year (5.12% of total revenue) and declined 4.24% over 2024.

This trend continued according to Oracle's Q2 FY2026 financial results released in November 2025. Oracle's quarterly revenue was \$16.1 billion, up 14 percent year-over-year. Its cloud infrastructure stood out with \$4.1 billion in revenue, up 68% over the second quarter of 2025. Its software revenue was down 3% to \$5.9 billion. Services grew 7% to \$2.78 billion, and hardware grew 4.56% to \$1.45 billion (Oracle, 2025).

4. The Enterprise Software Industry

Enterprise software constitutes the digital backbone of modern business—systems that manage core operations, data, and processes for organizations. It comprises software products and services used by businesses and institutions to manage operations, data, customer interactions, processes, and decision-making. It is a vast, dynamic ecosystem that has a market size of approximately \$800 billion in 2025, growing at 10-12% annually. It includes the following segments:

- Database Management Systems (DBMS): This is the historic heart of Oracle's empire. The market is split between:
 - Relational/SQL Databases: Oracle Database, Microsoft SQL Server, IBM Db2, open-source PostgreSQL and MySQL.
 - Cloud-Native & NoSQL Databases: Amazon Aurora/DynamoDB, MongoDB, Google Cloud Spanner, and Snowflake's cloud data warehouse.
 - Enterprise Applications (SaaS): Enterprise Resource Planning (ERP): Financials, HR, supply chain. Key rivals: SAP (the dominant leader), Workday (strong in HCM), Microsoft Dynamics 365.
 - Human Capital Management (HCM): SAP SuccessFactors, Workday.
 - Supply Chain Management (SCM): SAP.
 - Industry-Specific Verticals: All major players offer specialized solutions.
- Cloud Infrastructure & Platform Services (IaaS/PaaS): This is the new battleground, which is dominated by the hyperscalers: Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Oracle Cloud Infrastructure (OCI) is a focused challenger.
- Application Development & Middleware: Platforms for building, integrating, and deploying applications.

Over the past 15 years, the industry has undergone a fundamental architectural and business model transformation, moving from on-premises licenses to cloud based subscriptions model of business. From this perspective, the enterprise software industry could be broken into the on-premises licenses software segment, and the cloud-based software segment.

The on-premises segment was a high-barrier, high-margin oligopoly built on proprietary lock-in, fear-based sales, and complex implementations. Oracle, along with SAP and IBM, defined the landscape of segment through several distinctive features:

(1). Business Model

- Perpetual Licensing: Customers paid large upfront fees (often millions) for the right to use software indefinitely on their own servers.
- Annual Maintenance & Support: Typically, 22-25% of the initial license fee, paid yearly. This was the industry's profit engine, providing:
 - Bug fixes and patches
 - "Tax" on access to new versions
- Margins: License sales had ~90%+ gross margins. Maintenance had ~85-90% margins—creating an incredibly lucrative, recurring revenue stream with minimal cost.

(2). Technology & Architecture: The On-Premise Monolith

- Client-Server Architecture: Software installed on company-owned hardware in their own data centers.

- Vertical Integration: Vendors sold the entire stack—database (Oracle), middleware (Oracle Fusion Middleware), and applications (E-Business Suite).
- Proprietary Standards: Lock-in through proprietary programming languages (PL/SQL), data formats, and protocols.
- Customization-Heavy: Implementations required massive customization, often via professional services, making systems unique to each company.

(3). Sales & Customer Dynamics

- Enterprise Sales Driven: Focused on winning CIO/CTO through: Fear, Uncertainty, Doubt (FUD): "No one gets fired for buying Oracle/SAP/IBM."
- Multi-Year Enterprise Agreements: Front-loaded deals locking enterprises for 3-5 years.
- Top-Down Approach: Selling to executives rather than end-users.
- High Switching Costs: Once installed, switching was almost unthinkable due to:
 - Data migration complexity
 - Retraining thousands of employees
 - Rebuilding custom integrations
 - Vendor audit risks

(4). The Competitive Landscape (The "Big Three" Oligopoly)

- Oracle: Dominated the database layer (the most strategic piece, as all applications run on data). Used database dominance to push applications.
- SAP: Dominated the application layer (ERP), especially in Fortune 500 manufacturing and logistics.
- IBM: Dominated services and middleware, often the integrator between Oracle, DB, and SAP apps.
- Microsoft: Played in mid-market with SQL Server and Dynamics but wasn't an enterprise heavyweight in this era.

(5). Implementation Ecosystem

- System Integrators (SIs): A massive industry of firms like Accenture, Deloitte, IBM Global Services that:
 - Implemented software over 12-24 month cycles
 - Charged millions in consulting fees
 - Created vendor-customer dependency (only they understood the customized system)
- Partner Economics: Vendors maintained arm's-length relationships with SIs—They sold licenses, SIs sold services.

Oracle dominates the on-premises segment with approximately 32% of the market share, followed by SAP (22%) and IBM (15%). This segment represents the old, the traditional, and the past. It grew slowly at the rate of 1%-3% CAGR for the past decade and is projected to decline at a rate of -4% - -6% CAGR from 2025 – 2030. Nonetheless, the on-premises enterprise software segment is one of the most remarkable "cash cow" businesses in technology history - highly profitable, recurring, and resistant to change due to extraordinary switching costs. It still makes up 60-70% of the enterprise software market. Most enterprises, government, and defense workloads remain on-premises, creating a long tail for traditional support revenue. Many enterprises still run customized Oracle 11i/R12 or SAP ECC6—systems, which are too expensive to replace. Millions of Oracle DBAs, SAP consultants, and legacy integrators still maintain these systems. Besides, the audit-and-compliance business remains lucrative as vendors police license compliance in hybrid environments.

The cloud-based enterprise software segment has fundamentally changed the business model of the enterprise software industry. It represented shift from large, upfront license sales to recurring SaaS and consumption-based models. In contrast to the on-premises segment, this segment is characterized by continuous improvement and deployment, self-service adoption, open APIs, and interoperability. This segment generates massive revenue but moderate margins. This changes revenue streams, customer relationships, and requires continuous innovation. It has gone through four historical growth phases:

- (1). Emergence Phase (2006-2010): 40-50% CAGR
During this period, Salesforce establishes SaaS model. AWS launches IaaS (2006). Customers were typically early adopters in the tech/startup sectors
- (2). Adoption Phase (2011-2015): 30-35% CAGR
During this period, enterprises on the Fortune 500 list began cloud migration. Microsoft Azure gains traction and SaaS became mainstream for CRM and HCM
- (3). Dominance Phase (2016-2020): 25-30% CAGR
During this period, core ERP moved to cloud (Oracle Fusion, SAP S/4HANA). The covid-19 pandemic accelerated the adoption.
- (3). AI/Intelligence Phase (2021-2024): 18-22% CAGR
During this period, the market started to mature but remained robust. AI/ML integration became standard and industry-specific clouds started to proliferate.

The cloud-based enterprise software segment can be further classified into sub-segments:

Sub-Segment 1: Platform as a Service (PaaS)

This sub-segment has a market size from \$90-100B (2024) with a growth rate of 25-28%. It provides software services in the areas of Databases, analytics, AI/ML services, development tools.

- Cloud Databases: \$25B (dominated by Snowflake and Databricks)
- AI/ML Platforms: \$20B (fastest growing at 35%+)
- Integration/API Platforms: \$18B
- Development Tools: \$15B

Sub-segment 2: Software as a Service (SaaS)

This sub-segment has a market size from \$280-310B (2024) with a growth rate of 16-18%. It provides software services in the areas of CRM, ERP/Finance, Collaboration, HCM, Vertical SaaS, and other applications

- CRM: 22% (\$62-68B) - Salesforce dominant
- ERP/Finance: 20% (\$56-62B) - SAP, Oracle, Workday
- Collaboration: 18% (\$50-56B) - Microsoft, Google
- HCM: 12% (\$34-37B) - Workday, Oracle, SAP
- Vertical SaaS: 15% (\$42-47B) - Industry-specific
- Other Applications: 13% (\$36-40B)

Currently, the enterprise software industry is undergoing another major architectural shift in 50 years: from cloud-based business models to full-stack AI embedded software services. AI Integration is sweeping through every corner of the enterprise software industry and every application is going through a rebuilding process around the AI-first assumptions. Data + AI services are likely to capture disproportionate value and the lines between infrastructure, platform, and application software are blurring. This industry rewards scale, innovation, and ecosystem power, making the competitive dynamics among Oracle, Microsoft, SAP, and the hyperscalers the defining story in the enterprise software industry.

5. The AI Data Center Industry

AI data center is the cornerstone of the AI infrastructure. AI Data Centers are specialized facilities designed to support the unique, extreme demands of training and deploying large-scale AI models. The AI data center industry is the high-stakes, capital-intensive "picks and shovels" market of the AI gold rush. It's characterized by extreme demand, intensive capital expenditure.

AI data centers are optimized specifically to host:

- Large language models (LLMs) and other deep learning architectures
- High-performance computing (HPC)-like workloads
- Generative AI training and real-time inference
- Massive data throughput and low latency networking

They differ from conventional enterprise data centers in their emphasis on:

- Massive computing density
- Specialized accelerator hardware (GPUs, AI cores, TPUs)
- Enhanced power and cooling infrastructure
- High-bandwidth interconnects (NVLink, InfiniBand, 400 Gbps+ Ethernet)

This industry is characterized by the following features:

- **Extreme Cost & Complexity:** AI data centers demand high upfront investment in land acquisition & preparation; power infrastructure upgrades and substations; building construction tailored for high-density loads; procurement of expensive GPU, TPU chips, networking equipment, and racks. A single, state-of-the-art AI data center can cost \$1-12 billion to build and equip. Building and maintaining AI clusters is capital-intensive and requires rare expertise in distributed systems. This creates a high barrier to entry and favors well-funded players

- **Energy Consumption:** AI data centers have massive power demands. A rack of 8-16 NVIDIA H100 GPUs can draw 70-100kW, compared to ~10kW for a standard server rack. A single AI data center campus can consume 500-1,000+ megawatts, posing significant environment and grid stability challenges. The projected energy consumption for the five data centers Oracle committed to build for the Stargate project will exceed 4.5 GW, more than the energy consumption for the whole Chicago city.

- **Rapid Obsolescence:** Hardware and software stacks evolve at a breakneck pace. The 2-5 year AI chip cycle means facilities risk being outdated before they are fully depreciated (traditionally 10-15 years). Design must be modular and forward-looking.

- **Talent Shortage:** Acute shortage of engineers who understand deep learning, distributed systems, and hardware optimization.

- **Vendor Lock-in:** Fear of being locked into a single cloud provider's or chipmaker's ecosystem (e.g., CUDA) is pushing toward open standards.

Driven by the generative AI boom, the AI data center industry is explosive with a projected 25-35% rate of growth over the period of 2025-2030. This outpaces the general data center market (~10-15% CAGR) by a wide margin. McKinsey (Chui, Et al, 2023) projected that the total addressable market (TAM) for AI data center to reach \$300-500 billion by 2030. NVIDIA's data center revenue alone (selling GPU chips that go inside AI data centers) was over \$115.2 billion for the fiscal year ending January 2025, representing an increase of ~142% year-over-year (Nvidia, 2025). It reflects the massive market size and significant growth for the AI data center industry.

The AI data center industry has an oligopoly structure with a few hyperscalers dominating the industry with a combined market share of 65% - 75% (Tilawat, 2025). Among them, AWS (Amazon) has a market share of roughly 30%; Microsoft's Azure has a market share in the 20%; Google Cloud owns 13%; Meta & others (enterprise/private) has 5-7% (combined).

In the AI data center industry, chip makers (NVidia, AMD, Google TPU division) exert strong influence due to limited alternatives for high-performance accelerators. Power and utility partners influence site selection and operating costs. Government regulations, security and compliance requirements differ by jurisdiction. Besides, stricter data localization laws push providers to build region-specific capacity.

6. The Stargate Project

The Stargate Project is a massive infrastructure initiative. It aims to create a foundational computing network capable of powering advanced AI research and deployment — especially large-scale models that require massive training and inference capacity beyond the limits of conventional cloud setups. It is designed to:

- Provide vast high-performance computing resource specifically for AI development
- Support OpenAI's model training, evaluation, and inference at hyper-scale
- Strengthen U.S. competitiveness and technological leadership in AI
- Create significant economic and employment impacts across multiple U.S. states
- Build next-generation data centers optimized for AI workloads that far exceed typical enterprise requirements

The project plans to invest up to \$500 billion over four to five years. The initiative targets a total capacity of 10 gigawatts (GW) of AI computing power—enough to power approximately 7.5 million homes. It is estimated to create over 100,000 jobs across the U.S. in construction, operations, and technical sectors.

Its main partners include OpenAI, Softbank, Oracle, and Abu Dhabi-based investment firm MGX. OpenAI serves as operational lead of Stargate. SoftBank (led by Chairman Masayoshi Son) oversees the complex financing and debt-raising efforts. Oracle participates as a key member of the management investor group and is primarily responsible for the construction of physical infrastructure and providing Oracle Cloud Infrastructure (OCI) as the computing foundation.

Beyond the core owners, several companies hold critical "technology partner" status:

- **NVIDIA:** Collaborates closely on the "build and operate" phase, supplying GB200 hardware.
- **Microsoft:** Serves as a key technology partner and retains a "right of first refusal" to host OpenAI workloads, despite the project's own on-premises infrastructure.
- **Arm:** Provides the underlying CPU architecture for the specialized AI servers.

The initiative focuses on massive "hyperscale" campuses:

- **Stargate I (Abilene, Texas):** The flagship site, developed by Oracle and Crusoe Energy, is already partially operational and scaling toward 1 GW of capacity.
- **Expansion:** Five additional U.S. sites in Texas (Shackelford and Milam counties), New Mexico, Ohio, and a Midwest location.

Stargate also has an ambitious plan to the UAE (1 GW project), Norway (renewable-powered site in Narvik), the UK, and Argentina.

7. Oracle's Partnership with TikTok

The Oracle, TikTok partnership results from sustained U.S. regulatory pressure over national security concerns tied to TikTok's Chinese parent, ByteDance. U.S. law required major Chinese-owned tech platforms to divest control of their U.S. operations or face a ban. Oracle emerged as a key partner in structuring a compliant solution that keeps TikTok active in the U.S. rather than shut down.

TikTok is a short-form video social media platform centered on algorithmically curated content feeds. Users create, share, and consume videos typically ranging from seconds to a few minutes, featuring music, speech, filters, and AI-driven effects. As of the mid-2020s, TikTok is one of the most influential global social media platforms. It is a dominant force in digital culture, media consumption, and advertising.

TikTok's defining feature is its algorithmic recommendation engine, which prioritizes content relevance over social graph, learns user preferences extremely quickly, and continuously optimizes engagement signals (watch time, replays, interactions). Unlike Facebook or Instagram, TikTok users do not need to follow anyone to get engaging content, and viral reach is possible even for new creators. This makes TikTok fundamentally a content discovery platform, not just a social network. Its global monthly active users are around 1.5 billion (estimates vary by source) and its US users are around 170million. It has an extensive reach among generation Z millennials, and was adopted quickly even by older demographics, and small businesses and creators. TikTok has become a primary source of entertainment and posed as a major competitor to YouTube, Instagram, and even search engines in some contexts.

TikTok was developed by ByteDance, a Beijing-based technology company founded in 2012 by Zhang Yiming. ByteDance initially focused on AI-driven content recommendation, most notably through its Chinese news app Toutiao. In 2017, ByteDance launched TikTok internationally. In 2018, ByteDance acquired Musical.ly (a popular lip-sync app) for \$1 billion and merged it into TikTok. This merger accelerated TikTok's adoption in the U.S. and Europe.

TikTok's ties with ByteDance caused several challenges and controversies around the globe. Governments, especially the U.S., have raised concerns about national security, data access by foreign governments, and Algorithmic influence over public opinion. In 2019, Concerns first gain prominence in the U.S., driven by statements from Senators Marco Rubio and Chuck Schumer, focusing on ByteDance's obligation under China's 2017 National Intelligence Law to cooperate with state intelligence work. In 2020, The Committee on Foreign Investment in the United States (CFIUS) launches a formal review of ByteDance's 2017 acquisition of Musical.ly (which became TikTok), citing risks that user data could be accessed by the Chinese government. In 2020, President Trump issues executive orders attempting to ban TikTok or force its sale to a U.S. company (Oracle/Walmart). These orders are challenged in court and later put on hold by the Biden administration.

In 2021, TikTok publicly announced its partnership with Oracle and the "Project Texas". According to this announcement, Oracle will store and manage TikTok's U.S. user data on its cloud infrastructure (OCI).

This is a mutually beneficial partnership. For TikTok, localized data storage with trusted cloud provider increases transparency, oversight, and, to some extent, alleviates privacy and data residency concerns among regulators. For Oracle, securing workloads from TikTok could meaningfully increase its cloud utilization, reinforce Oracle's profile in secure cloud and AI infrastructure, and help Oracle compete with major hyperscalers by showcasing regulatory and security credentials.

Most importantly, before TikTok, Oracle had never served a client like OpenAI. The partnership with TikTok implies a boot camp of hell training for Oracle. The cloud computing and algorithm requirements from TikTok are completely different from Oracle's customers. TikTok is a company with exceptionally extreme - even insatiable - demands when it comes to algorithms and cloud computing power. In providing cloud services to TikTok, Oracle was forced to learn how to manage how to manage ultra-large-scale GPU clusters, how to handle massive concurrent parallel computing, and how to fulfill the extraordinary requirements from the most demanding clients. These echoes the characteristics of OpenAI's cloud computing demands. To some extent, the partnership with TikTok prepared Oracle for its bold endeavor into AI computing today. As Oracle's current Co-CEO Clayton Magouyrk told one ByteDance senior executive, "without you, we would not have achieved what we have today."

8. Conclusion

The Stargate Project represents a defining moment for Oracle. It is simultaneously a platform bet, a financial gamble, and a personal legacy decision. Whether Stargate becomes Oracle's reinvention or its most expensive miscalculation remains uncertain.

Ellison immersed himself in deep thought in front of all the information he had received. While AI boom advanced rapidly with widespread adoption and massive investment in the AI infrastructure, a fierce debate emerged about the AI bubble, which centered on whether the current surge in AI investment—particularly generative AI, large language models (LLMs), and AI data-center infrastructure—reflects sustainable value creation or a speculative bubble analogous to the dot-com era. This debate is not about whether AI is useful, but whether current investment levels and valuations are justified by near- to medium-term monetization. Wall Street financial analysts pointed out several major problems with the AI boom:

A. Capital Expenditure Far Ahead of Revenue

While AI infrastructure spending is growing at 20–30%+ CAGR, Many AI companies and startups generate limited revenue, operate at significant losses, and they depend on continuous external funding to survive.

B. Weak Unit Economics for Many AI Applications

LLM inference is expensive with high marginal cost per query and high ongoing model retraining costs. Many application cases have improved productivity but failed to generate new revenue streams.

C. Valuations Driven by Narrative, Not Cash Flow

Most AI companies trade at extreme revenue multiples with no clear profitability timeline.

D. Overcapacity Risk in AI Data Centers

While hundreds of billions of dollars, multi-gigawatt AI data center projects announced globally, they pose significant risks for capability underutilization if AI adoption slows, and risks of stranded capital if cloud computing prices fall rapidly.

These risks are particularly severe for Oracle's stakes in the Stargate project. The proposed role for Oracle is a key technology, cloud infrastructure, and building partner for the project. Public filings indicate the scale of Oracle's participation may involve hundreds of billions of dollars in multi-year commitments under long-term supply and capacity. Oracle's 2025 capital commitment alone for AI data center is \$7.0 billion.

Oracle's old business model was characterized by stable streams of revenue, high margin, diversified (yet sticky) clients, and healthy cash flow. Now, to participate in the Stargate project, Oracle must incur massive debt and invest billions upfront to procure GPUs. These GPUs typically have a life span of 2-5 years and will become obsolete or need to be replaced before 2029 when Stargate project is scheduled to be completed. In this project, Oracle's cloud growth would become heavily dependent on a single, albeit massive, customer, OpenAI. Any slowdown or cancellation in later phases of the project would leave Oracle with stranded assets and revenue shortfalls. These factors, combined with the low margin AI data center business may hurt Oracle's balance sheet and drain its cash flow.

Ellison realized that Oracle operates in a fiercely competitive environment for its traditional business, where it must defend its core database and applications business against hyperscalers and best-of-breed attackers, attack the cloud infrastructure market dominated by larger players, and innovate faster than ever by embedding AI across its entire stack. Its success hinges on executing a difficult multi-front war: using its enduring strengths in mission-critical software to fund and fuel its growth in the cloud and AI eras, while constantly evolving its culture and technology to match the pace of AI cloud competitors.

This converges well with its participation in the Stargate project. The AI data center industry is not just a construction boom; it is fundamental re-architecting of the IT infrastructure. It represents a convergence of high-tech semiconductor engineering, heavy industrial power systems, and cutting-edge thermodynamics. Whoever master the AI data center stack will control the physical bottlenecks of the AI era. This industry is a critical battleground where tech giants, energy companies, and nations are competing for supremacy, as it directly underpins economic and strategic advantage in the age of artificial intelligence.

Having lost the last three technological paradigm shifts in IT, the pain of further loss weighs heavily in Ellison's evaluation of the Stargate project. Will Oracle lose again in the fast-progressing AI revolution? Stargate was thus not merely a matter of corporate strategy; it was a high-stakes legacy project for an 81-year-old founder determined to rewrite the final chapter of his career. The Stargate Project represents a defining moment. It is simultaneously a platform bet, a financial gamble, and a personal legacy decision. Whether Stargate becomes Oracle's reinvention or its most expensive miscalculation remains uncertain.

9. Discussion Questions

1. Was Microsoft rational to reject full ownership of OpenAI's infrastructure needs?
2. Does Oracle possess a sustainable competitive advantage in AI infrastructure?
3. How should investors evaluate AI data center risk differently from traditional cloud CapEx?
4. Is Stargate a rational response to missing previous platform shifts — or an emotional overcorrection?
5. What metrics would indicate success or failure within the first 5 years?
6. What has Oracle gained from its partnership with TikTok?

10. Reference List

- Chamberlin, D. D., et al. (1976). "SEQUEL 2: A Unified Approach to Data Definition, Manipulation, and Control." *IBM Journal of Research and Development*, 20(6), 560-575.
- Chui, M., Hazan, E., Roberts, R., Singla, A., Smaje, K., Sukharevsky, A., Yee, L., & Zempel, R. (2023, June 14). The economic potential of generative AI: The next productivity frontier. McKinsey & Company. <https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-economic-potential-of-generative-AI-the-next-productivity-frontier>
- Codd, E. F. (1970). "A Relational Model of Data for Large Shared Data Banks." *Communications of the ACM*, 13(6), 377–387.
- Nvidia (2025), Press release Feb. 26, 2025: NVIDIA Announces Financial Results for Fourth Quarter and Fiscal 2025. <https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2025>
- Oracle, (2025), Press release December 10, 2025, Oracle Announces Fiscal Year 2026 Second Quarter Financial Results, <https://investor.oracle.com/investor-news/news-details/2025/Oracle-Announces-Fiscal-Year-2026-Second-Quarter-Financial-Results/default.aspx>
- Tilawat, M. (2025), AI Data Center Statistics 2026: The \$200 Billion Revolution in Global Infrastructure, <https://www.allaboutai.com/resources/ai-statistics/ai-data-centers>.
- Wilson, M. (1997). The difference between God and Larry Ellison: God doesn't think He's Larry Ellison. William Morrow.