■ Signal From Noise   February 22, 2024

Artificial Intelligence

"AI can benefit humanity in the form of free doctors, free tutors, free education, more media and entertainment, better wealth advisors and lawyers, more robotics, better digital health, better materials design; much better resource discovery: lithium, cobalt, nickel – the hard-atom stuff that tech optimists usually don't talk about." – Vinod Khosla, November 12, 2023

"Accelerated computing and generative AI have hit the tipping point. Demand is surging worldwide across companies, industries, and nations." – Jensen Huang, February 22, 2024

AI-Related Granny Shots vs. the S&P 500

NVIDIA (+2,298%) · TSLA (+645%) · AAPL (+355%) · META (+201%) · MSFT (+172%) · GOOGL (+171%) · ANET (+58%) · AMD (+53%) · CDNS (+66%) · AMZN (+10%)

Source: Bloomberg, Fundstrat. Performance in parentheses is absolute since date added to the Granny Shots portfolio. All charts rebased to the date the stock was added to the Granny Shots portfolio. Arista Networks is a Super Granny, Tesla is a Sleeper Granny.

The landscape

The players in the world of AI are, without exception, characterized by overweening ambition and a hypercharged amalgam of [FOMO](#) and [YOLO](#). AI as a moniker (in the sense of adding "AI" to anything and everything) is, in part, a fad, but AI as machine-learning models powered by GPU training***** at the bleeding edge certainly has every chance of being transformational for many industries.

Among consumer-facing companies there are AI startups that focus on leveraging AI to get content copyright right and startups that offer AI tools to help companies better communicate with their customers. Of course, there is AI to help you convert voice to text (Otter.ai) and vice versa ([IBM's voice chatbots](#)), and AI to convert text to more text (Claude, [Gemini](#) (formerly known as Bard; developed by Google), Poe, [Perplexity](#)), or to images ([Midjourney](#)) or to videos ([Sora](#)). There is even AI that labels images or text in order to teach AI models to recognize patterns (Scale AI).

With a few notable exceptions, many of the startups offering these services are not publicly traded, but enough of the AI-infrastructure companies are and, as we know, with any gold rush, it pays to invest in the suppliers of the mining equipment.

Enterprise AI firms that should be on everyone's radar are: AMD, Arm, TSMC, Cadence Design Systems, ASML, Arista Networks, ServiceNow, Snowflake, Autodesk, and Adobe. There is also the Global X Artificial Intelligence & Technology ETF ([AIQ](#)).

The transformative power of AI

At this point, few people doubt that AI could mean a revolution in the way people communicate and business is done – similar to the one brought about by the arrival of the internet in the mid-nineties. Some say AI will take us even further. The question is what type and scale of transformation AI might bring about. While forecasts are difficult, it is a more realistic and, arguably, profitable proposition to point to those companies that will likely get us there.

Dan Ives of Wedbush Securities speaking to Fundstrat's Tom Lee on the broad significance of AI

As the basic listing of some of the major AI players above intimates, Amazon (AMZN), IBM, and Google should be on the short list of anyone seeking to invest in the groundswell of machine-learning-driven AI tools. Meta is another solid, if *sui generis*, contender. Microsoft (MSFT) is a juggernaut that – largely thanks to its prescient investment in and partnership with OpenAI (whose latest capital raise took place this week, valuing the company at $80 billion and paving the way to an IPO in the near to medium term, at a much higher valuation) – as well as continued curiosity about the newest approaches and models (see its investment in Mistral AI), should be in a category all its own when it comes to AI ambitions.

Research only available to banks & hedge funds, now available to individuals

WEEKDAYS	ON MARKET EVENTS
Tom Lee's First Word	**Tom Lee's Intraday Word**
The best market research out there	These are released periodically

— on a near-daily basis.

after crucial market developments.

DAILY

Daily Technical Strategy

Daily commentary on market trends and short term opportunities.

DAILY

⚡ Flash*Insights*

By-the-minute notifications and updates on breaking market news on an as-needed basis from FS Insight's thought leaders.

WEEKLY

US Policy

Weekly update on developments in Washington DC and how they may impact Wall Street.

WEEKLY

Weekly Roadmap

Weekly summary of market events and updates on our key theses.

EARNINGS SEASON

Earnings Daily

A multi-faceted update of earnings developments.

GET A FREE TRIAL AND VIEW ALL OUR MACRO RESEARCH!

Gray old lady teams up with sexy ingenue

For one, Microsoft is addressing the entire AI stack, starting with the inputs – and we don't mean the training data. Appreciating how massive the electricity needs are – per Constellation Energy's data, AI's new demand for power could be five-six times what's used for charging EVs – Microsoft's medium- to long-term solution is to *go nuclear*. A carbon-free way to output round-the-clock power is a win-win – except, that is,

for the long, complex, and expensive U.S. nuclear regulatory process project developers face. As Microsoft's all-in AI strategy means great demand for electricity, the tech giant is putting a chicken-and-egg-inspired twist on efforts to obtain it.

As you are reading this, Microsoft is *continuously asking generative AI to help it streamline the approval process for nuclear power projects, so it could generate better AI*.** (We've written about this in **one of our First To Market daily updates**).

Research only available to banks & hedge funds, now available to individuals

DAILY
Crypto Market Update
Stay on top of the crypto markets with macro commentary, technical analysis, and the most relevant news.

WEEKLY
Crypto Technical Strategy
From Bitcoin, Ethereum and Solana to other coins, our Head of Technical Strategy highlights important price movement to help give you an edge.

WEEKLY
Crypto Strategy
A must-read summarizing our key thoughts about the crypto market, price action, and developing trends.

WEEKLY
Funding Fridays
Are you looking for investment opportunities in the crypto space? We explore early-stage crypto opportunities in this note about where the smart money is being invested!

WEEKLY
Crypto Events

INTRA-MONTH
Deep Research Reports

Never miss another conference, a high-impact software update, or an event that could be market-moving!

Our consulting business delivers incredible deep research reports to highlight the economics and viability of exciting crypto projects.

GET A FREE TRIAL AND VIEW ALL OUR CRYPTO RESEARCH!

Nvidia (NVDA)

Nvidia is, of course, the dominant technological engine and the sine qua non behind the rise of AI applications. The company's GPU (graphics-processing unit) chips are, for now, the irreplaceable standard for microprocessors powering AI functionality, and when paired with Nvidia's CUDA toolkit, which offers a development environment for creating high-performance GPU-accelerated applications, they have also proven to be irresistible – and a kind of currency*** – powering Nvidia's stock to newer and newer heights. More than $1.4 trillion of its capitalization has been added in the past 12 months, and Nvidia is currently the third most valuable company by market cap, behind Microsoft and Apple (**AAPL**), while Goldman Sachs' trading desk **calls it** "the most important stock on Earth".

Treated as a proxy for AI demand, Nvidia stock is now also the **most traded** by value, dethroning Tesla for the title. As if Nvidia's market share and mindshare were not staggering enough, CNBC's Jim Cramer said yesterday that Nvidia CEO Jensen Huang is a bigger visionary than Elon Musk, who Cramer says is the P.T. Barnum of our time, while seeing
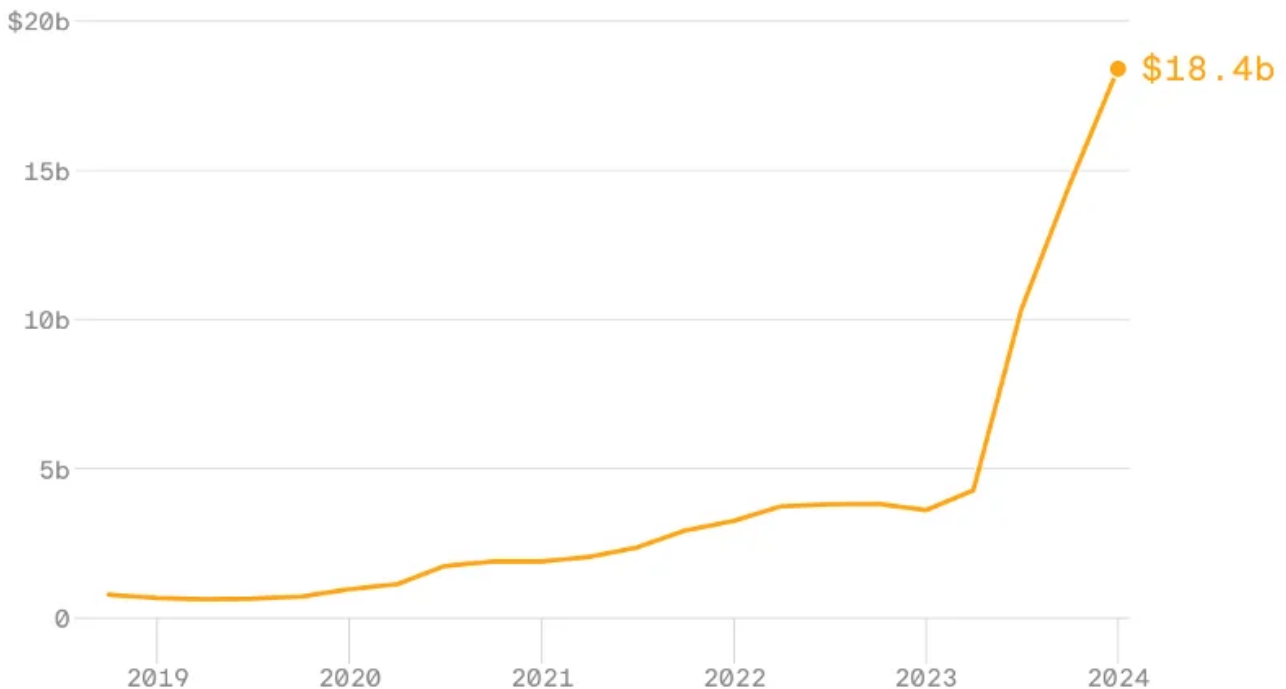
Huang as more like Leonardo da Vinci and Taylor Swift, in the way their success is unparalleled in their respective fields.

It's important to expand one's view of AI beyond just chips and into services. Nvidia made more in Q423 from data-center revenue than it did during the entirety of its previous quarter.

Quarterly Nvidia data center revenue

Q4 2018 to Q4 2023

Data: FactSet; Chart: Axios Visuals

Nvidia seeks to provide a full stack for AI training via the Nvidia AI Enterprise Engine, for which it charges customers per GPU per year (just like for an operating system), with customers then able to run anything Nvidia creates and enables. Nvidia's services include software-based management, optimization, and patching.

Intel (**INTC**). Announcing plans a few days ago to become the Number 2 chip foundry by 2030, Intel has also just announced that Microsoft will use its next-generation 18A technology (14A is the current-generation tech) to make an undisclosed chip, and that it now expects $15 billion worth of foundry orders, up from the $10 billion the company had earlier told investors to expect.

AMD offers processors, graphics cards, accelerators, and adaptive SoCs (systems on chip) for running generative AI LLMs (large language models), video analytics, recommendation engines, and other data-intensive AI workloads in Automotive, Aerospace, Healthcare, Industrial, Telecom, and advanced research.

Arm Holdings (**ARM**) offers microprocessor cores, graphics and camera technology, pre-verified systems and system IP for SoCs, physical IP and processor implementation solutions for Automotive, computing infrastructure, Consumer tech, and IoT (internet of things).

TSMC(Taiwan Semiconductor Manufacturing Co.), a Taiwanese multinational semiconductor contract manufacturing and design company, is the second most valuable semiconductor company in the world; also the world's largest independent and dedicated, i.e. pure-play, semiconductor foundry.

ASML, based in the Netherlands, isthe dominant supplier of lithography equipment used to manufacture advanced logic and memory chips, and enjoys strong demand from leading chipmakers. TSMC uses ASML's EUV (extreme ultraviolet) as well as older deep ultraviolet (DUV) systems to manufacture chips for fabless chipmakers such as Apple, Nvidia, Qualcomm, and AMD. ASML 's EUV lithography machines are required to

produce next-generation AI accelerators and high-bandwidth memory modules.

IBM's AI solutions include IBM watsonx.ai, a data platform with a set of AI assistants, watconx.data, which helps scale responsible AI across the enterprise, and watsonx.governance, which directs, manages, and monitors customers' AI projects by integrating responsible workflows for generative AI.

Industrial robotics, industrial automation, simulating proteins, forecasting weather – moving into more and more of the world's modalities is the next big opportunity for generative AI. Linkages between semiconductors, software, and systems will be crucial going forward. Many components, from data center design to software applications to privacy systems, have to work together to handle increasingly complex AI use cases. The crucial control points will likely be closer to the silicon foundation of the AI infrastructure.
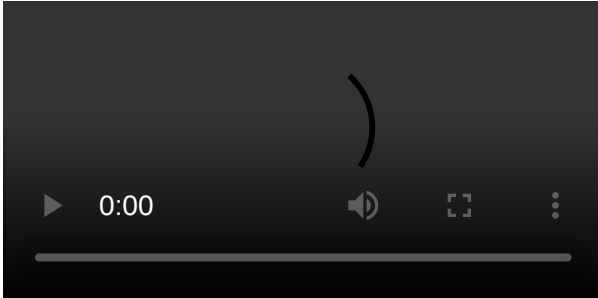
More and more companies are fabless chipmakers, i.e. they design their own chips. Not content with relying on standard chips that are in high demand, and are therefore fiercely competed for, major tech firms have been trying to control their computational fate for years. Perhaps the most notable example came in November 2020 when **Apple announced** it was moving away from Intel's x86 architecture to make its own M1 processor, which now sits in its new iMacs and iPads. Apple, Amazon, Facebook, Tesla, and Baidu (**BIDU**) are all shunning established chip firms and bringing certain aspects of chip development in-house. Both Amazon and Microsoft have developed designs for networking

chips, which improve performance of the cloud servers and AI services they sell to customers. Still, for the moment, none of the tech giants are looking to do all the chip development themselves, as setting up an advanced foundry, like TSMC's in Taiwan, costs around $10 billion and takes several years.
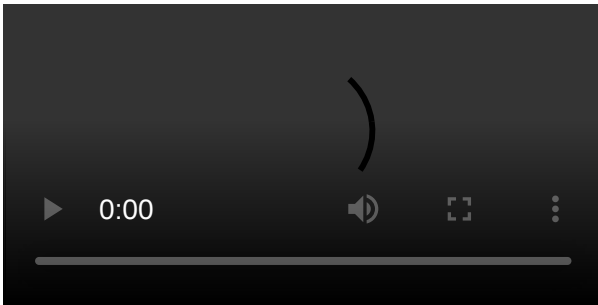
Tesla (**TSLA**) – the automaker designs its own chips and then has major chip makers manufacture them. Previously a client of Samsung, Tesla switched to TSMC for chips governing its Full Self Driving (FSD) computer, which is a classic example of AI/ML (machine learning). Another major TSMC customer is Appleˆ.

Google (**GOOG**) – Google, which has taken a turtle's approach to developing AI tools, still does incorporate machine learning in certain products. For instance, the AI in Google Maps analyzes real-time traffic data to provide up-to-date information about traffic conditions and delays. Its most capable model, and the basis of a growing ecosystem, is Gemini, which powers the Gemini chatbot and includes Firebase extensions that allow users to build features using the Gemini API.

META (**META**) says it seeks to build artificial general intelligence systems that will give its products more human-like capabilities in coming years. It announced that it had begun training Llama 3, the next generation of its primary generative AI model, in mid-January, reiterating a commitment to releasing its AI models via open source — when possible.

Webinars!

We host monthly webinars throughout the year with exclusive insights from our research team & high-profile special guests.

Become a member today and get full access to ALL webinars & future events!

Our Upcoming and Latest Webinars!

Granny Shots & Market Update Webinar – July 2024

Hyperion Decimus: A Diversified Quantitative Digital Asset Manager

Granny Shots & Market Update Webinar – June 2024

Mark Newton's Live Technical Analysis – June 2024

Hyperion Decimus: A Sophisticated and Risk-Managed Way of Investing in Liquid Digital Assets – May 2024

ACCESS OUR WEBINARS WITH A FREE TRIAL!

Government.ai

AI infrastructure is in high demand and short supply, and that's pushing tech giants and chipmakers into a shotgun alliance with governments both in the U.S. and around the world to try to boost output and free up chokepoints. This new AI industrial complexis

rushing to spend fortunes to avoid overdependence on Taiwan's dominant chip factories.

Despite Silicon Valley's libertarian rhetoric and free-market preferences, government and high tech industries have long been intertwined, as illustrated by the internet itself, which started at DARPA, as a government research project.

Following an executive order that mandates that agencies advance their use of AI and build rules for it, AI systems are already in use or will be rolled out in more than 1,200 distinct ways within the U.S. government, from monitoring the border to studying volcanoes. As parts of U.S. law define AI differently, and given the range of tasks where the software can come in handy, getting government offices to agree on a definition of AI will be a challenge.

Still, the U.S. Department of Defense plans to use artificial intelligence to develop more sophisticated autonomous weapons, including drones that need less human direction. Around the world, militaries are increasingly turning to drones for reconnaissance and attack, with Ukraine deploying 200,000+ last year. The Pentagon wants to set international norms for the development of AI-powered military tools. OpenAI announced recently that it is dropping its ban on making military tools in order to work with the Pentagon, but insisted it still prohibits weapons development.
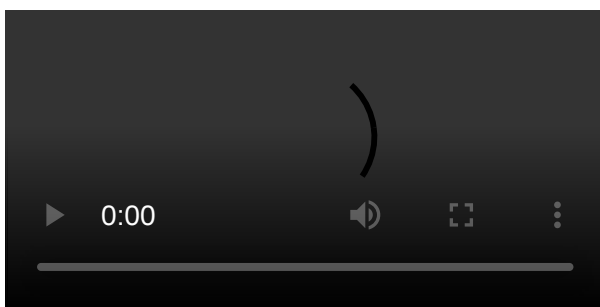
One company in the defense space that should be watched very carefully is Palantir (**PLTR**). **Ranked** a leading AI, data science, and machine learning vendor by market share and revenue, Palantir offers secure LLM access, multiple end-user applications, and an integrated architecture that can potentially bring AI to every decision. Palantir CEO Alex Karp has **spoken** about the two separate AI arms races now happening: the domestic and the peer-to-peer one, explaining why it is imperative that we lose one and win the other.
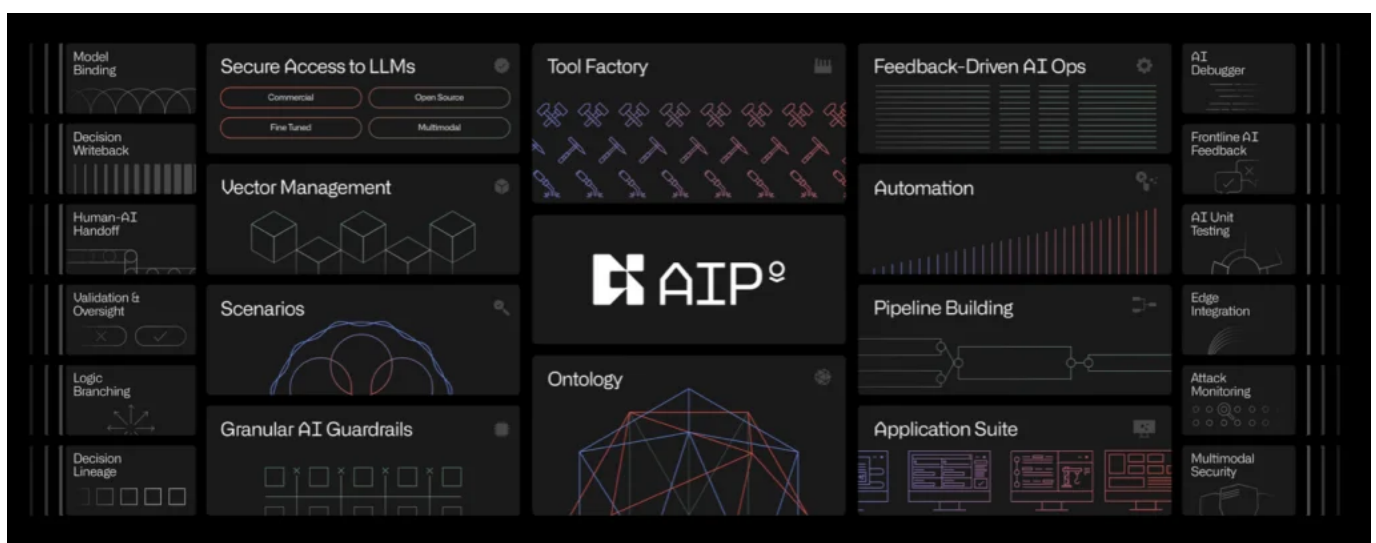
Source: Palantir

The guts of AI

In the meantime, no one has enough AI infrastructure, i.e. the advanced microprocessors that power generative AI systems. CEOs are hoarding Nvidia chips while a number of companies are promising to build chip factories in the U.S., and state-owned or -funded efforts around the world are seeking to build and commercialize AI models in a range of languages.

Although we have all become very familiar with the Artificial Intelligence darlings like Nvidia and Microsoft, there are a handful of companies that are quietly powering this revolution, creating the tools and infrastructure that enable AI's remarkable capabilities. While AI advancements grab headlines, these companies are all established, successful firms, playing a crucial yet often unseen role in its advances. Here are a few of the best names, both of the hardware and software kind.

Cadence Design Systems (**CDNS**) – Cadence provides tools that manufacturers of complex electronics use to improve the design and performance of sophisticated chips. Its clients include Samsung, Palantir, and NASA.

Synopsis (**SNPS**) provides chip-design software. Its list of traditional customers, mostly semiconductor companies, has now been supplemented by superscalers such as Google, Microsoft, et al, as well as chip fabricators such as Intel.

Arista Networks (**ANET**) – The company designs and sells multilayer network switches to deliver software-defined networking for large datacenter, cloud computing, high-performance computing, and high-frequency trading environments. Specific to AI, Arista provides solutions for GPU and storage interconnects driving AI/ML workloads using high-performance IP/Ethernet switches.

ServiceNow (**NOW**) – ServiceNow offers a cloud computing platform to help companies manage digital workflows for enterprise operations. It surfaces context-aware recommendations and delivers improved self-service through conversational AI. It also uses AI to anticipate and

prioritize customer resources and drive improvements through real-time analytics.

Snowflake (**SNOW**) – Snowflake is a data cloud company offering a cloud-based data storage and analytics service, generally termed data as a service. It has successfully used AI tools to provide data storage, processing, and analytic solutions.

Autodesk (**ADSK**) – Autodesk offers software and services for the architecture, engineering, construction, manufacturing, media, education, and entertainment industries. Its AI tools allow clients to perform predictive analysis for wind, noise, and operational energy in real time, predict flood maps, explore additional design options, and automate repetitive tasks.

Adobe (**ADBE**) – Adobe employs machine learning and deep learning to enhance content understanding (including images, videos, etc); recommendations and personalization; search and information retrieval; prediction and journey analysis; content segmentation, organization, editing, and generation. For instance, the AI-powered features in Acrobat and Acrobat Reader products improve the comprehension of lengthy PDF content.

Where AI is going

Futurists predict that enterprise software companies will benefit from millions of AI agents.

OpenAI CEO Sam Altman **wants to raise trillions** (although it may well end up being an order of magnitude less than that) to create a network of chip factories to produce advanced semiconductors, most of which today — including Nvidia's coveted GPUs — are

manufactured in Taiwan. This approach, involving the construction and maintenance of semiconductor fabs, marks a departure from the strategy of OpenAI's industry peers like Amazon, Google and Microsoft, which typically design their silicon and outsource production. Building such facilities requires substantial investment, with a single plant costing tens of billions of dollars.

Early this year, while in Davos, Altman said that "none of the pieces are ready" for delivering AI infrastructure "at the scale that people want it."

OpenAI is in talks to raise new funding at a valuation of $100 billion or more. It has also just raised $80 billion as part of a tender offer.

Microsoft flagged chip shortages as a risk to its bottom line in its **2023 yearend report**.

Taiwan's role as the world's pre-eminent supplier of high-end chips could be disrupted by an invasion, blockade, or other military interference from China, and that's a direct risk for TSMC, and a reason more and more companies will seek to rid themselves of their dependence on this foundry. For example, U.S.-based AMD, Nvidia's closest competitor in the AI chip market, has been **looking to reduce its dependence** on the industry-leading TSMC, which fabricates most AMD-designed products.

TSMC itself is hedging its bets. It's working on a two-plant, $40 billion project in Arizona. GlobalWafers, another Taiwan-based chipmaker, broke ground on a Texas plant in November, after CEO Doris Hsu warned that supply chains would break down within weeks if China invaded Taiwan.

AI investment is likely to grow in the next three years

Private AI investment (dotted lines show GS revenue projections*)
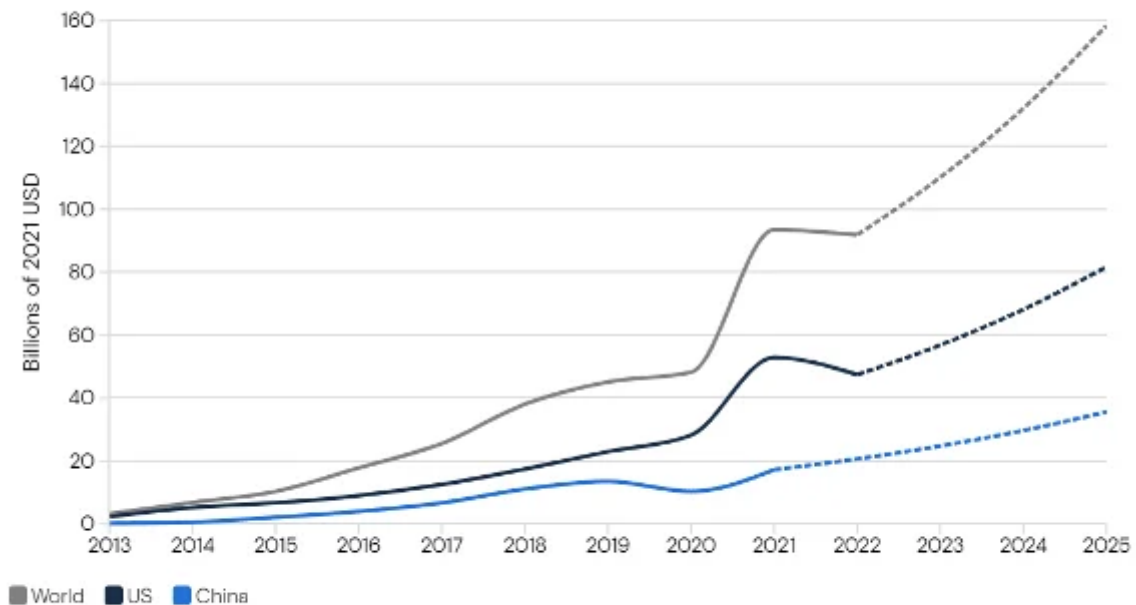
Source: Stanford Institute for Human-Centered Artificial Intelligence, Goldman Sachs Research · *Average of GS Research 2022–2030 revenue growth estimates for Microsoft Azure, NVIDIA, Google Cloud, and Amazon Web Services (when available)

Goldman Sachs

In the meantime, $200 billion has been committed to new U.S. chip manufacturing infrastructure, across several dozen sites, per the Semiconductor Industry Association. The biggest investments come from the likes of Intel, Micron, and TSMC — and the Biden Administration has also announced funding for **25 National AI Research Institutes**. $1.5 billion of CHIPS Act funding was recently awarded to GlobalFoundries (**GFS**), the third-largest semiconductor foundry by revenue.

Global investments are also piling up. Tiny U.A.E., with fewer than 10 million citizens, has launched a new state-backed artificial-intelligence company to commercialize sectoral versions of the country's high-performing Falcon model. Per the country's AI minister, Omar Sultan Al Olama, the goal is to use AI to turn U.A.E. into a global power player.

U.A.E.'s $1.5 trillion sovereign wealth fund controls GlobalFoundries, one of the five biggest semiconductor fabricators, post its 2021 U.S. IPO.

Japan's effort to reinvigorate its flagging chip sector has drawn the backing of American companies, including IBM.

National AI stockpiles of GPUs are being created in India and the U.K. For example, the Indian government has purchased **24,500 GPUs for startups and academics** to use at 17 dedicated centers, while the British government will invest over $600 million to provide **advanced chip access to researchers, non-profits and startups**. The Indian government will also fund 500 AI/deep tech startups at the product development stage, mirroring French President Emmanuel Macron's plan to pump **$600 million into creating local "AI champions."** Ministers in India, France and the Middle East see homegrown AI models as ways to preserve and promote their languages and culture. Krutrim, an Indian start-up, launched a 10-language LLM in December, with CEO Bhavish Aggarwal saying it will help make "India the most productive, efficient and empowered economy in the world."

As usual, Signal From Noise should serve as a starting point for further research before making an investment, rather than as a source of stock recommendations. Although the names mentioned above each have the potential to benefit from government investment in infrastructure, this alone should not be the basis of a decision to invest.

We encourage you to explore our full Signal From Noise library, which includes deep dives on the path to automation and opportunities arising from the ever-increasing global water crisis. You'll also find a recent discussion of the Magnificent Seven and the rise of Generation Z.

* Graphics processing units (GPUs) were originally developed for accelerating graphics processing. They can dramatically speed up computational processes involved in deep learning and are an essential part of a modern **artificial intelligence infrastructure**. New GPUs are being developed and optimized specifically for deep learning.

** First comes the training. Given a six-month timeline, Microsoft engineers have set themselves to training an LLM (large language model) on U.S. nuclear regulatory and licensing documents. The goal is to speed up the process of filling out and filing the paperwork required for such approvals. Normally, the process takes years and costs hundreds of millions of dollars. A partnership with Terra Praxis, a nonprofit that encourages the reuse of former coal plants as sites for smaller-scale reactors, divides the task into two streams: Microsoft does the coding, Terra Praxis supplies the regulatory expertise. Routine work is automated, allowing narrow domain experts, i.e. regulators and developers, to do work that needs the attention of people, not robots.

*** Mark Zuckerberg announced recently that Meta had amassed 340,000 of the Nvidia H100 GPU chips (which powered the development and use of ChatGPT); including other chips, Meta projects it will have the equivalent of nearly 600,000 H100s by the end of the year. Not taking any chances, Zuckerberg says his H100 buying spree "may be larger than any other individual company['s]", and that it will continue, whether others "appreciate that" or not. (Shades of Elon Musk, who said of the then-impending debut of the Cybertruck, "It will be the biggest product launch of anything by far on Earth this year.") This despite the fact that, as far back as 2019, having determined that existing chips were largely inadequate for deep learning – and realizing that future chips used for deep-learning algorithm training, which underpins most of the recent progress in AI would require the ability to manipulate data without having to break it up into multiple batches – Meta was working on its own breed of semiconductor.

Disclosures

Conflicts of Interest

This research contains the views, opinions and recommendations of FS Insight. At the time of publication of this report, FS Insight does not know of, or have reason to know of any material conflicts of interest.

General Disclosures

FS Insight is an independent research company and is not a registered investment advisor and is not acting as a broker dealer under any federal or state securities laws.

FS Insight is a member of IRC Securities' Research Prime Services Platform. IRC Securities is a FINRA registered broker-dealer that is focused on supporting the independent research industry. Certain personnel of FS Insight (i.e. Research Analysts) are registered representatives of IRC Securities, a FINRA member firm registered as a broker-dealer with the Securities and Exchange Commission and certain state securities regulators. As registered representatives and independent contractors of IRC Securities, such personnel may receive commissions paid to or shared with IRC Securities for transactions placed by FS Insight clients directly with IRC Securities or with securities firms that may share commissions with IRC Securities in accordance with applicable SEC and FINRA requirements. IRC Securities does not distribute the research of FS Insight, which is available to select institutional clients that have engaged FS Insight.

As registered representatives of IRC Securities our analysts must follow IRC Securities' Written Supervisory Procedures. Notable compliance policies include (1) prohibition of insider trading or the facilitation thereof, (2) maintaining client confidentiality, (3) archival of electronic communications, and (4) appropriate use of electronic communications, amongst other compliance related policies.

FS Insight does not have the same conflicts that traditional sell-side research organizations have because FS Insight (1) does not conduct any investment banking activities, and (2) does not manage any investment funds.

This communication is issued by FS Insight and/or affiliates of FS Insight. This is not a personal recommendation, nor an offer to buy or sell nor a solicitation to buy or sell any securities, investment products or other financial instruments or services. This material is distributed for general informational and educational purposes only and is not intended to constitute legal, tax, accounting or investment advice. The statements in this document shall not be considered as an objective or independent explanation of the matters. Please note that this document (a) has not been prepared in accordance with legal requirements designed to promote the independence of investment research, and (b) is not subject