

Bittensor: The Internet of AI

April 16th, 2025

Tom Couture^{AC}
Twitter: @Tom_Couture17



Table of Contents

1 Artificial Intelligence

2 What is Bittensor?



3 Network & Performance

4 Subnet Highlights

5 Looking Ahead

Slowly, Then All at Once

- AI technology has rapidly evolved from rule-based systems to powerful machine learning models capable of understanding language, generating content, and making complex decisions, with large-scale models making significant leaps in general-purpose AI capabilities.





Artificial Intelligence is Everywhere


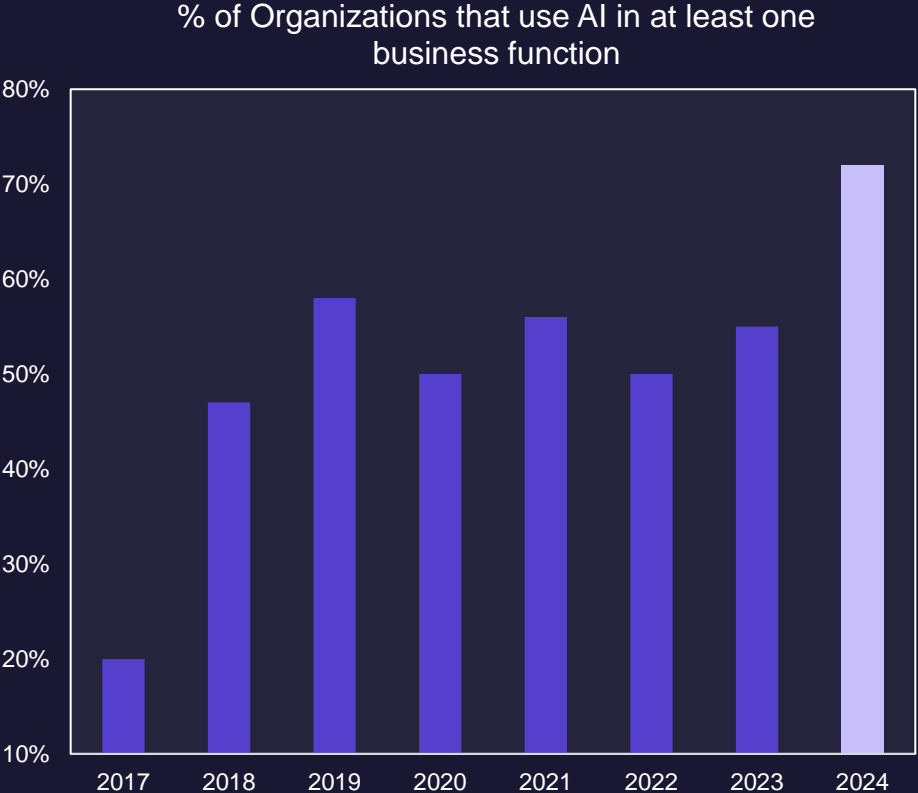
- Advances in data availability, computing power, and neural network design have driven exponential growth in AI capabilities, enabling machines to perform tasks once thought exclusive to humans — including language translation, image recognition, and even creative writing.
- AI is expected to become deeply integrated into nearly every sector, from healthcare to education to national security, raising opportunities for productivity and innovation.





Disrupting The Business World

- After remaining stagnant for the last few years, the number of firms using AI in at least one business function jumped in 2024 from 55% to nearly 75% of organizations.
- In an enterprise survey of over 1,600 C-Suite executives and their employees conducted by the Writer, it was found that AI usage has disrupted work activities for all employee levels, with C-suite executives reporting higher usage than employees in nine out of the ten surveyed categories.







Market Magnitude

- According to Grand View Research, the global AI market was valued at \$279 billion in 2024 and is projected to grow at a CAGR of 35.9% from 2025 to 2030, reaching a total value of \$1.76 trillion by 2030.
- PWC expects that by 2030, AI will have cumulatively added \$15.7 trillion to global GDP due to productivity gains and innovative products and services.
- Companies are rushing to capture market share by heavily investing in AI, with the four largest hyperscalers projected to spend \$325 billion on capital expenditures in 2025.

Global AI Projected Market Size (\$B)



Hyperscaler* Annual Capex (\$B)





The Foundation of it AI

- Foundational models are large-scale AI models trained on vast and diverse datasets (like text, code, or images), enabling them to perform a wide range of tasks with minimal specific tuning. They serve as the "base layer" for many AI applications, allowing developers to build on top of them for specialized use cases.
- Foundational model investment has been prolific with companies like OpenAI and Anthropic raising 10's of billions of dollars at 12-figure valuations.

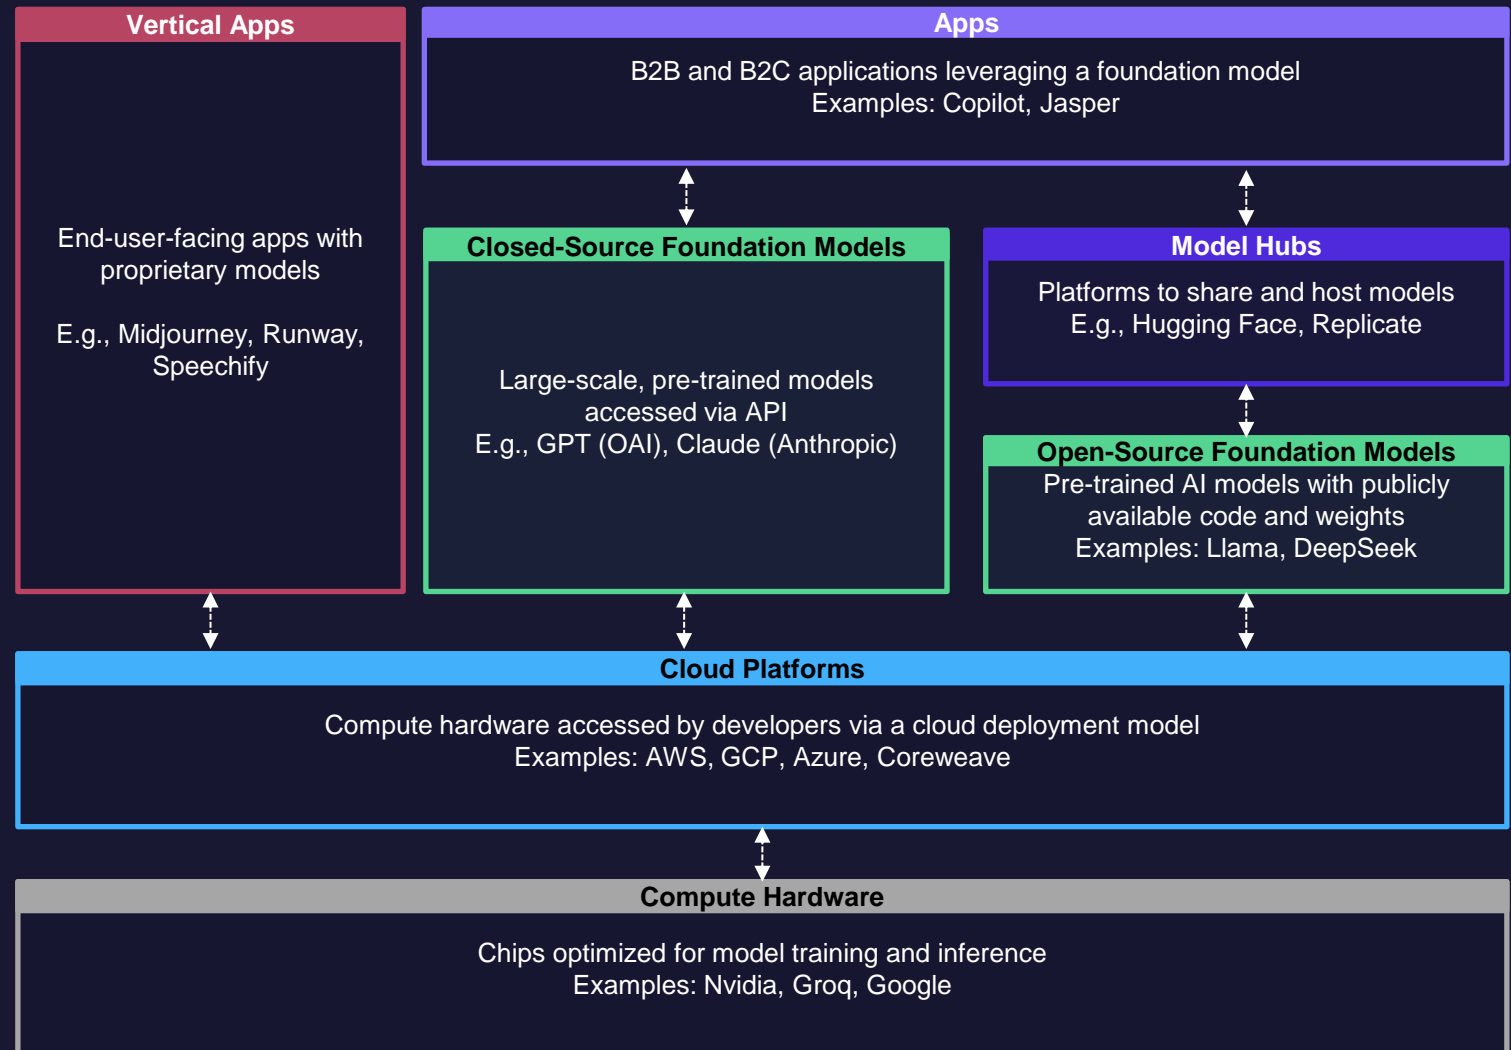
Metric	OpenAI	Anthropic	Google	Meta	Amazon	DeepSeek	Mistral AI	xAI
SOTA Model	OpenAI o3	Claude 3.7 Sonnet	Gemini 2.0 Pro Experimental	Llama 3.3 70B	Nova Pro	DeepSeek R1	Mistral Au Large	Grok 3
AA Intelligence Index	70	48	49	41	37	60	N/A	66
Reasoning Model?	Yes	No	No	No	No	Yes	N/A	N/A
Open-weight?	No	No	No	Yes	No	Yes	Yes	Yes
Go-to-Market	Apps & API	Apps & API	API & Existing Products	Existing Products	API & Existing Products	Apps & API	Apps & API	Apps & API
Funding	\$24B	\$16B	Public	Public	Public	N/A	\$1.2B	\$12B
Valuation	*\$300B	*\$62B	Public	Public	Public	N/A	\$6.5B	\$50B

 = Open Source





Breaking Down the AI Stack

- Much of the value generated from AI has accrued to foundational models and infrastructure providers such as semiconductor companies (Nvidia) or cloud providers (Amazon, Google, Microsoft).
- Compute costs are a significant barrier to entry, with advanced models needing thousands of high-end GPUs.
- Foundational models also require significant data storage and energy costs to run at scale.
- As the cost of inference declines, it's likely that value will shift from infrastructure and foundational models to consumer or vertical apps.



Risks of Centralized Artificial Intelligence

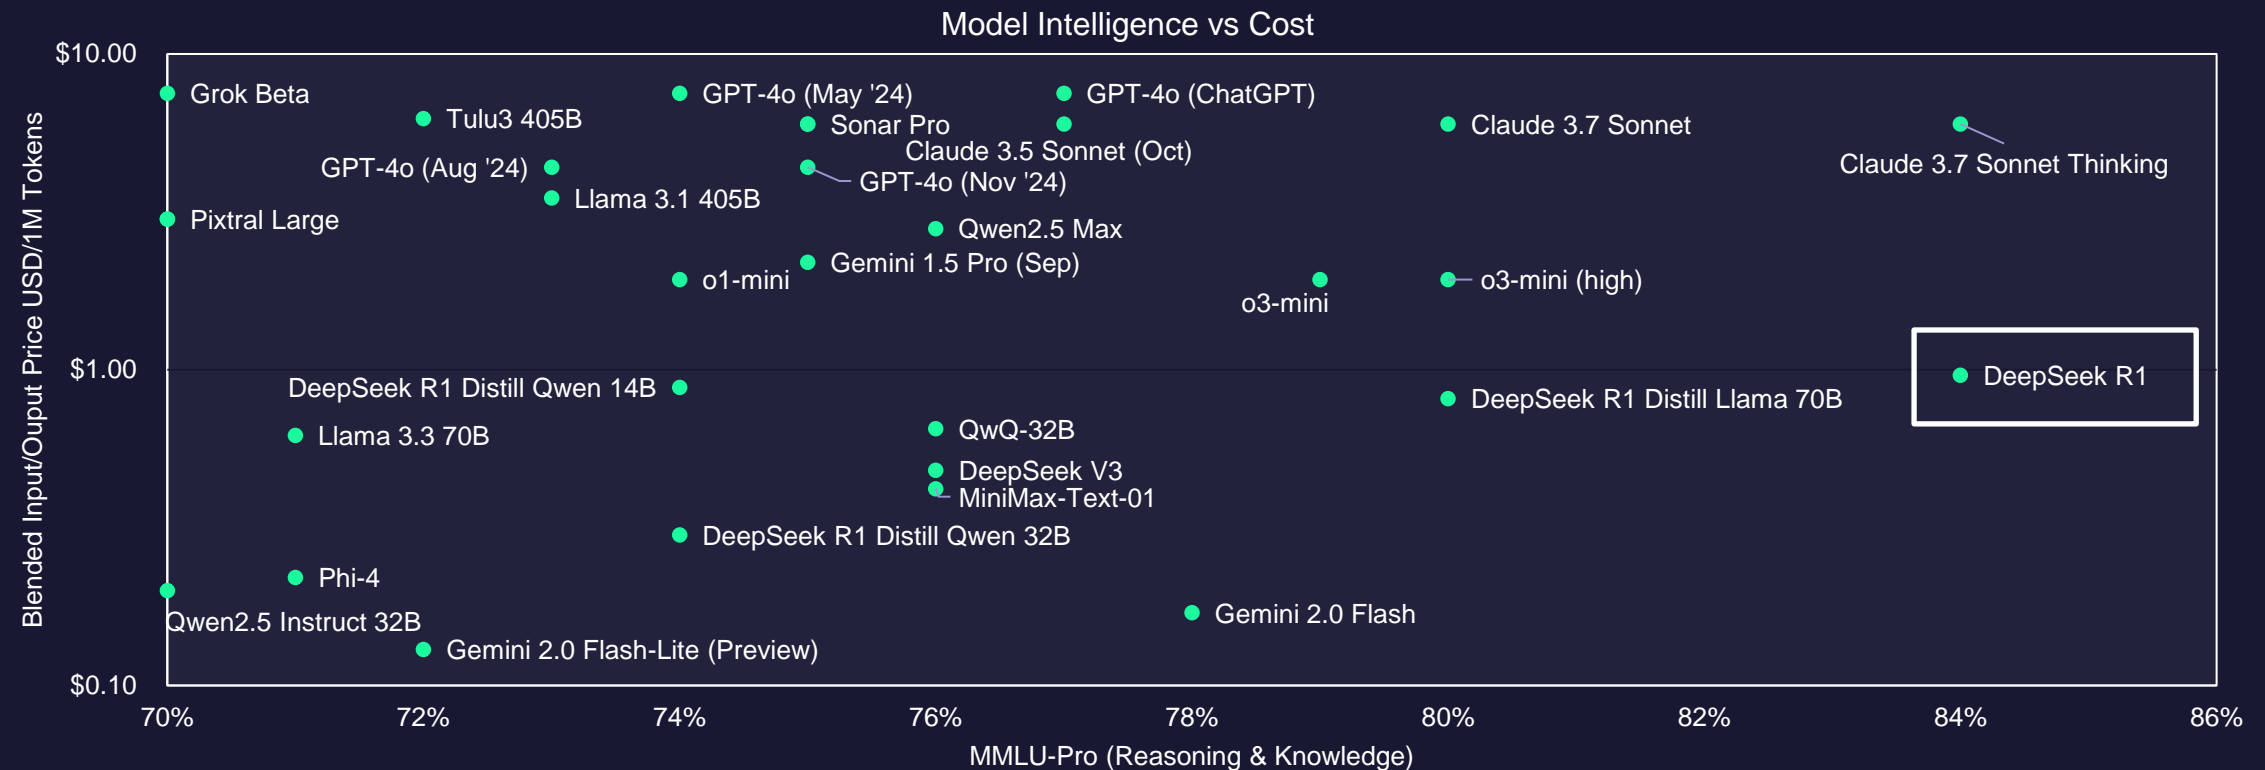
- Intensive capital requirements and competitive advantages have given large tech companies an upper hand in the AI development race. Increasing concentration of user data and model outputs among a few companies has raised concerns about AI governance and transparency.
- Foundational models such as Google's and DeepSeek's have shown biases, historical inaccuracies, and political leanings, displaying the human element that's still evident in proprietary AI development. As AI continues to embed itself in society, the need for transparent objective models is paramount.
- Additionally, developers outside of large-tech companies have limited access to advanced tools and models, stifling competition and potentially leading to monopolization.





The DeepSeek Moment

- Earlier this year, China released an open-source AI model, known as DeepSeek, that exceeded the performance of leading American models such as OpenAI's o1 model. DeepSeek was allegedly created and trained with far fewer resources, with some estimating that only five to six million was spent to train the model, compared to the estimated hundreds of millions for OpenAI's models.
- The DeepSeek release was a pivotable moment, causing a historical sell-off in technology names as investors questioned American technological dominance and igniting a global AI race.

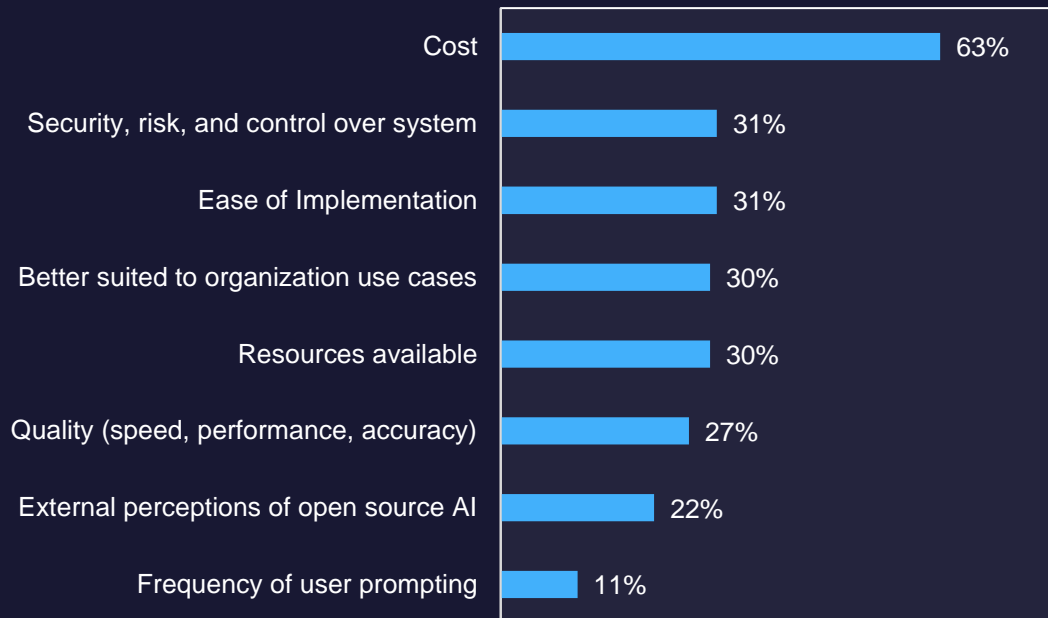





Shift To Open-Source Development

- The DeepSeek situation shed light on the power of open-source development, and the ability of open-source models to compete with centralized technology behemoths. Companies are already beginning to integrate open-source technologies into their businesses.
- In a McKinsey business leader survey of over 700 technology companies, they found that more than 50% of respondents are already leveraging open-source AI, with 76% expecting to increase open-source usage over the next few years.
- Additionally, organizations that view AI as important to their competitive advantage are 40% more likely to use open-source AI models and tools than other organizations.


Reason for Open-Source Preference



>50% *of respondents are already using open-source AI*



76% *of respondents expect to increase open-source AI usage*





Closed vs. Open-Source Models

- Open-source AI models offer several advantages:
 - **Transparency:** Their open nature allows for scrutiny, reducing the risk of unintended behaviors.
 - **Customization:** Developers can tailor models to specific needs without relying on third-party providers.
 - **Censorship Resistance:** Decentralization diminishes the likelihood of content suppression, promoting free access to information.

AI model alignment refers to the congruence between an AI system's objectives and human values or intentions. Misaligned models can produce undesirable or harmful outputs. Additionally, centralized AI models may be subject to censorship, limiting their accessibility and utility in certain contexts.

Category	Open AI	Closed AI
Models	Publicly available	Kept private
Data	Transparent training & data sharing	Training data specifics kept secret
Code	Open-source code	Closed source code
Development Cycles	Slower	Faster
Ease of Use	Less support services and infrastructure	Infrastructure and support services provided
Control	Community	Internal controls
Commercial Benefits	Shared innovation	Maintains competitive edge, commercial benefits
Scrutiny	Broader community identification of problems	Limited to internal resources for identification
Recruitment	Attracts innovative developers	May attract developers seeking proprietary work
Bias Identification	Transparency in training sources	Difficult to identify bias due to lack of transparency
Scalability Improvements	Collective innovation accelerates scalability	Closed models may need to rely on internal efforts



Crypto Can Alleviate Open-Source Pain Points

- Historically, open-source models have faced funding and coordination challenges, with no clear ways to support development costs and incentivize contributions.
- Token incentives provide a potential solution, offering an innovative way to fund development and coordinate innovative contributions, while distributed infrastructure can reduce capital requirements and enhance scalability, possibly facilitating a new era of open-source development.
- Introduction of DAOs can decentralize governance decisions to ensure transparency and stakeholder alignment.

Open-source Challenges	How Crypto Can Help
Incentive Misalignment	Shared token economies can be used to align incentives for developers and end-users. Token rewards can support developer costs and incentivize continual contributions.
Resource Accessibility	Open-source developers often lack access to the proper hardware or advanced development tools needed for advanced AI innovation. A permissionless network of decentralized development tools and distributed infrastructure lowers barriers to entry for developers and is supportive of accelerated AI innovation.
Governance & Coordination	DAOs allow for transparent and decentralized decision-making, helping to coordinate objectives of the entire ecosystem and support cross-project governance.



Table of Contents

1 Artificial Intelligence

2 What is Bittensor?

3 Network & Performance


4 Subnet Highlights

5 Looking Ahead



Bittensor – The Internet of AI

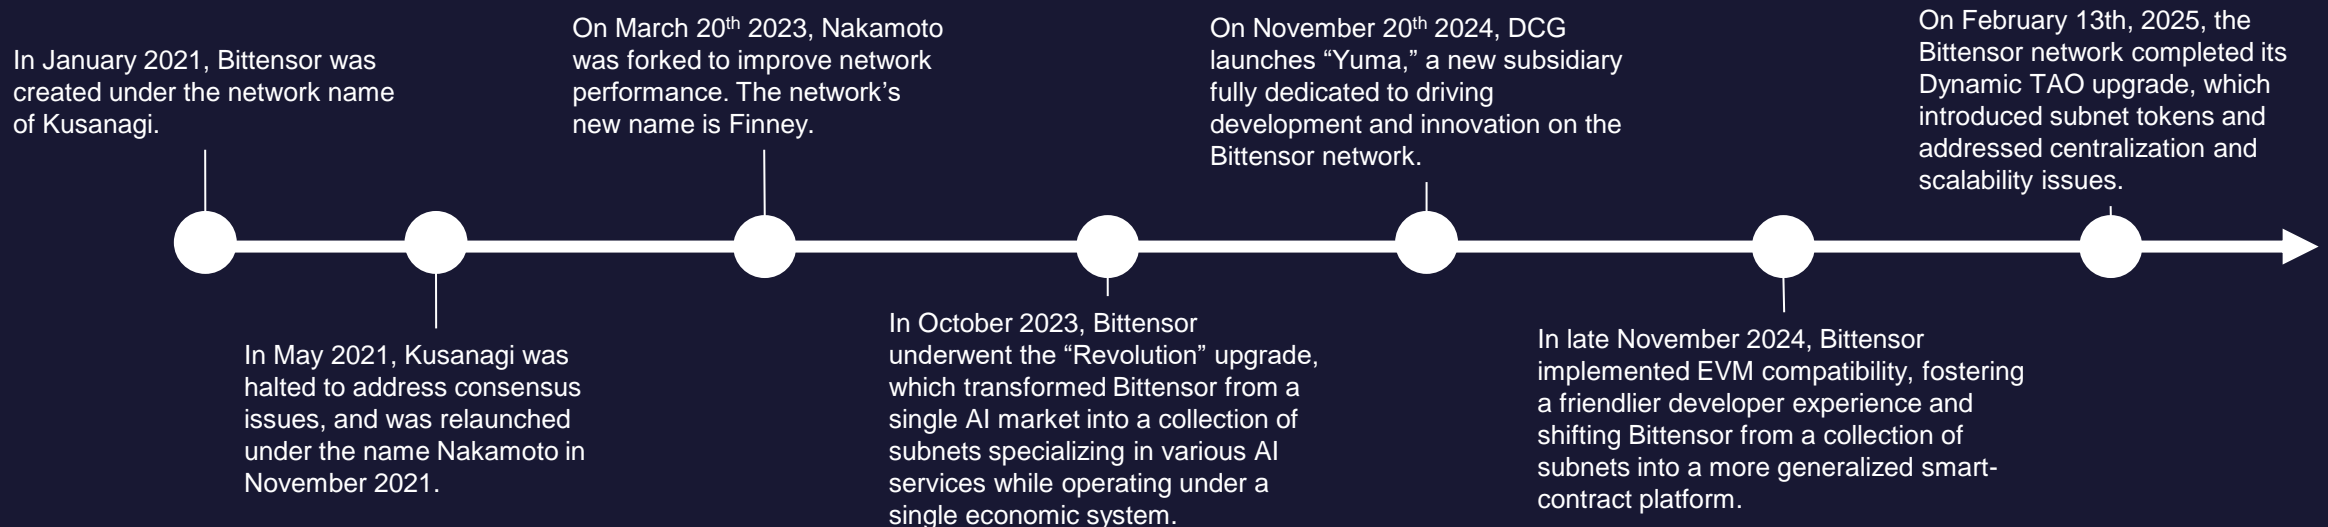
- Bittensor is a decentralized artificial intelligence protocol that utilizes a crypto economic model to incentivize machine learning, AI model creation, and commoditized inference to create the "Internet of AI."
- Bittensor envisions a future of a distributed network of interconnected AI "subnets," forming a credibly neutral global AI platform enabling permissionless access to AI resources and applications.





Evolution of the Bittensor Network

- Inspired by Bitcoin, Bittensor was created in 2021 to create an incentivized network for machine learning and artificial intelligence. Instead of incentivizing SHA256 hashes, Bittensor would use the TAO token to incentivize productive AI services, starting with a single text-prompting network.
- In October 2023, Bittensor underwent the “Revolution” upgrade, which transformed Bittensor from a single AI market into a collection of subnets specializing in various AI services while operating under a single economic system.
- Earlier this year, Bittensor underwent one of its most significant upgrades, with the introduction of Dynamic TAO (dTAO), which is the iteration of its tokenomics and governance model, directly addressing centralization and scalability issues as it pertains to evaluating subnet contributions within the network.





Bittensor Components & Contributors

- There are multiple ways to participate in the Bittensor ecosystem, with the different roles working in conjunction across Bittensor components to create a self-sustaining decentralized AI ecosystem.
- The Bittensor blockchain is responsible for recording balances and transactions for miners, validators and subnet creators, and allows for arbitrary parties to stake currency into subnets to support their work.

Components


1. Bittensor Blockchain (Subtensor) - The Bittensor blockchain serves as a system of record, and its token TAO serves as incentives for participation in subnet activities. Miners and validators harvest emissions based on their performance within subnets.
2. Subnets - Each subnet is an incentive-based market that produces a digital commodity related to artificial intelligence. It consists of a community of miners who produce the commodity, and a community of validators who measure the miners' work to ensure its quality.
3. Yuma Consensus - An algorithmic process used to score the performance of validators and miners, and in turn compute corresponding emissions.
4. Bittensor SDK – A software development kit that supports interactions between miners and validators within subnets and allows all parties to interact with the blockchain as necessary. Bittensor provides all the open-source tools, codebase, and documentation for developers.

Contributors

1. Subnet Creators — Manage the incentive mechanisms that specify the work miners must complete and what validators are evaluating.
2. Miners — Work to produce digital commodities (inference, compute, machine learnings, etc)
3. Validators — Evaluate the quality of miners' work and relay responses to users.
4. Stakers — TAO holders can support specific validators by delegating TAO to them.
5. End-Users – Consumers can use specific products across various subnets for personal use or leverage outputs to build other applications within the Bittensor ecosystem.



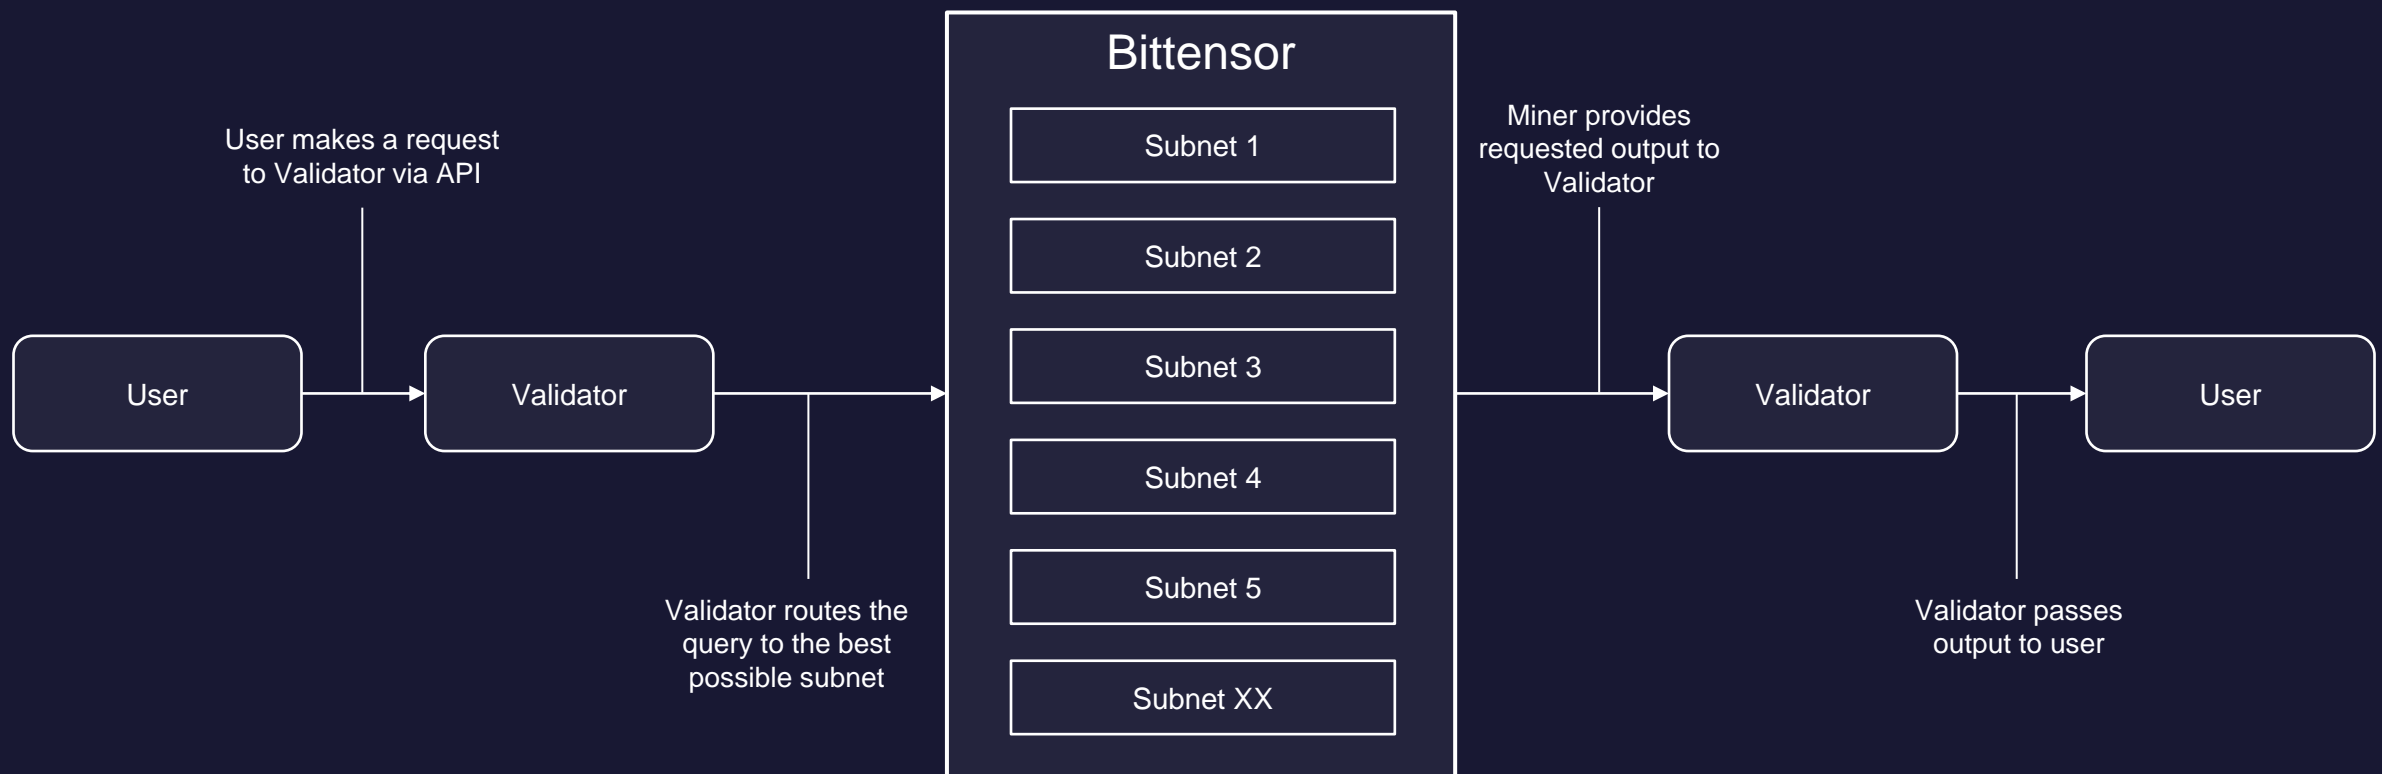
Bittensor Architecture





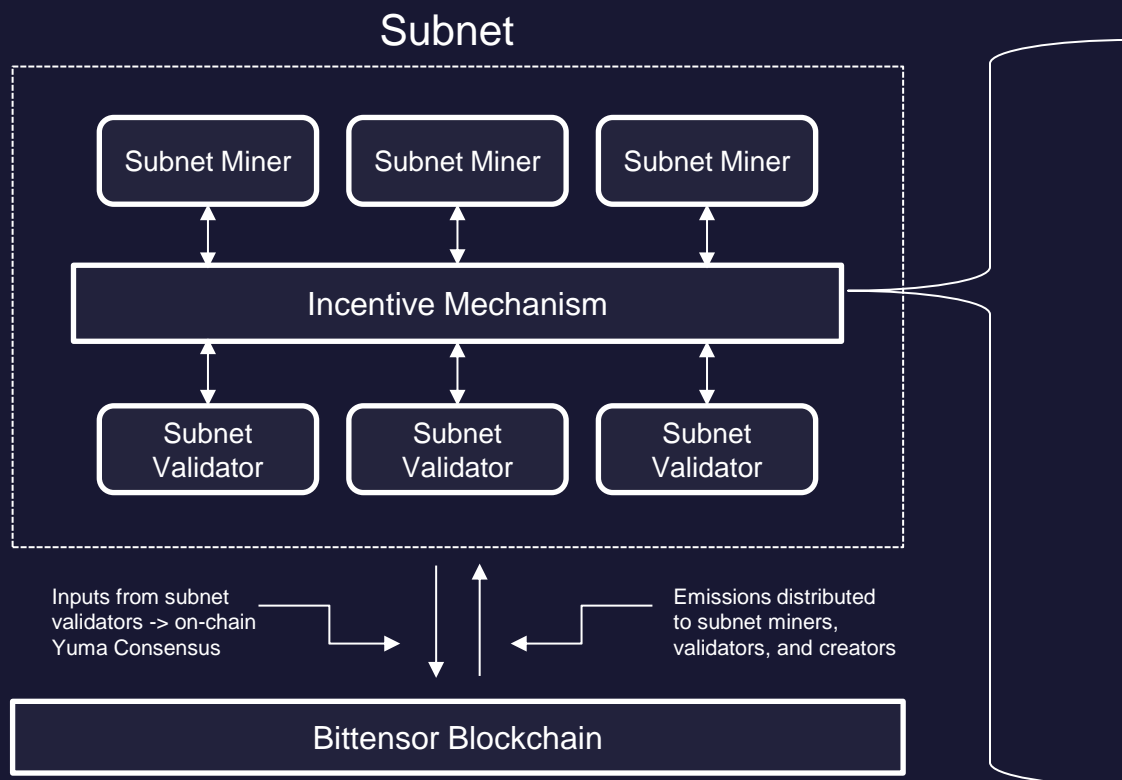
Users on Bittensor

- Validators are the gateway to accessing the Bittensor ecosystem, as they are responsible for routing jobs to the best subnets and relaying miner responses to Bittensor end-users. Bittensor is more oriented towards application developers than consumer apps.
- Users can purchase access to validator APIs using fiat currency in a more web2-like experience. If a validator chooses to validate on every subnet, a single API provides access to the entire Bittensor ecosystem.



Subnet Design

- Each subnet is designed with a unique incentive mechanism to drive what tasks miners complete and how validators score miners' work. It is up to the subnet creator to ensure the incentive mechanism's design aligns participants' goals.
- Validators evaluate miners' outputs and periodically compute a vector of scores for individual miners, representing an aggregate ranking system for performance. The ranking weight vectors are transmitted to the Bittensor blockchain and Yuma Consensus is leveraged to determine miner and validator emissions.



Design Considerations

For subnets to succeed, the incentive mechanisms should:

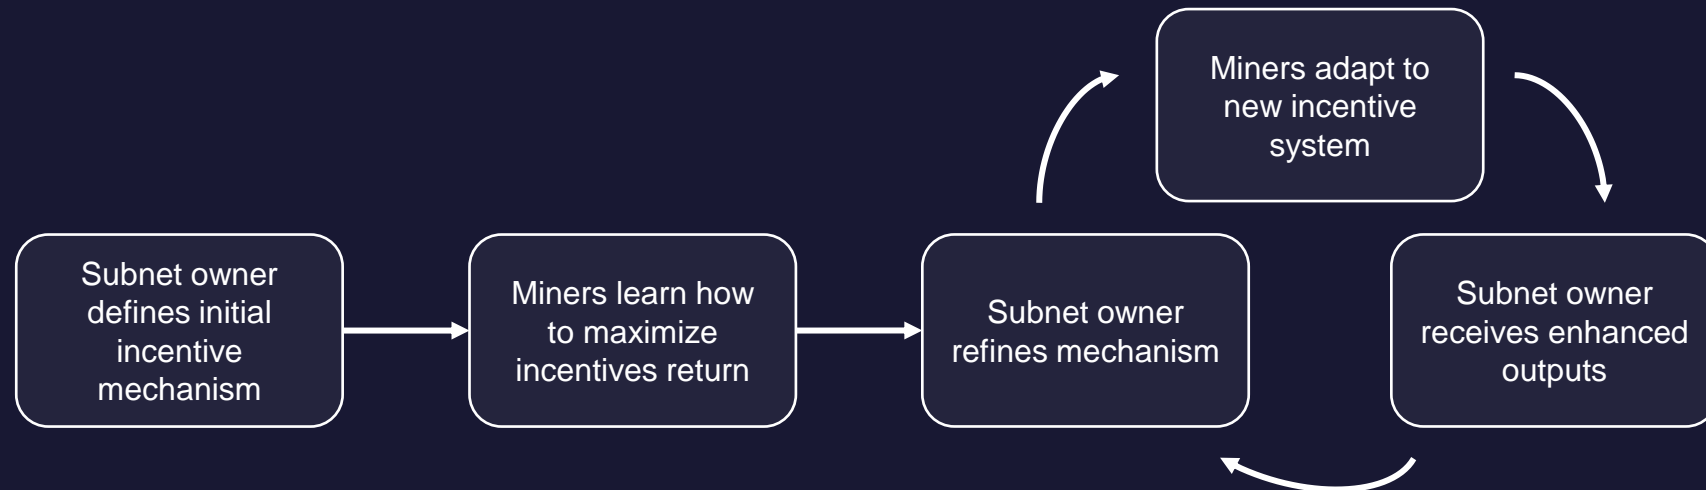
1. Motivate miners to optimize their validator scores (and hence emissions) to achieve the desired subnet outcome.
2. Facilitate miner and validator participation.
3. Prevent miner and validator exploits.

A well-designed incentive mechanism fosters a virtuous cycle, where competition drives continuous miner improvement, and a flawed mechanism encourages shortcuts that degrade the subnet's integrity and undermine participation.



Subnet Enhancement Cycle

- Subnets are designed for specific use cases, and miners are rewarded based on how effectively they fulfill the subnet's mission. As a result, miners are incentivized to find the most efficient path to the desired output, regardless of whether they use AI models.
- As miners exploit existing incentives, subnet owners will develop new mechanisms to ensure they continue receiving optimal outputs.
- This creates a dynamic feedback loop: the subnet enhancement mechanism functions like a perpetual bug bounty, where miners earn emissions for discovering the easiest path to a desired output—enabling continuous iteration and improvement by the subnet owner.





Validator Stake Weight

- If a validator wants to earn emissions from a subnet, they must “stake” TAO into that particular subnet’s liquidity pool in exchange for the subnet’s “alpha token”. Validators can also choose to stake to Subnet Zero (Root Network), which has no alpha token. Prior to the Dynamic TAO upgrade, the top 64 validators on Subnet Zero were responsible for determining emission allocations.
- Following the dTAO upgrade, staking to SN0 provides less stake weight than validating on subnets. The current Root Weight is set to 18%. Over time, the amount of circulating alpha tokens relative to TAO will increase, making subnet validation more central to the network.
- A validator’s total stake weight is the validator’s stake weight relative to the total network’s stake weight and is a quantifiable way to evaluate how credible a validator is. It feeds into Yuma Consensus when determining a validator’s voting power and emissions.

Validator Stake Weight = Alpha Stake + (Root Stake × Root Weight)

Alpha Stake = Amount staked to subnets

Root Stake = Amount staked to Subnet Zero

Root Weight = Adjustment Variable for determining Root:Alpha emissions

$$\text{Total Stake Weight} = \frac{\text{Validator Stake Weight}}{\text{Total Network Stake}}$$

Current Root vs. Alpha Stake

Root Stake

5.92 Million TAO

Alpha Stake

0.9 Million TAO



Yuma Consensus

- Yuma Consensus, also known as proof-of-intelligence, is an algorithmic process operating on Subtensor, which ingests validator weight vectors (miner scores) to evaluate the performance of miners, and in turn compute relative emissions.
- Validators assign a score from zero to one to evaluate the accuracy of a miner's outputs according to the subnet's designed purpose. A higher score indicates better accuracy and utility.
- The algorithm is designed to more heavily weight the inputs of trusted validators (higher stake weight), and disregard unreliable weight-settings. If a validator's assessment of a miner deviates significantly from the network's consensus, the validator's stake weight is "clipped", incentivizing validators to honestly and accurately evaluate miners.
- Yuma Consensus allows for fuzzy consensus around probabilistic truths like machine intelligence. It is not a rigid structure for evaluating miners' work, and it leaves an element of subjectivity in determining which miners are producing the best outputs, creating flexibility for developers to incentivize unique digital commodities or services.

Traditional Consensus

Bitcoin – Verify if a hash is valid

- Binary Outcome (Y/N)
- Same answer globally

Traditional Blockchain –

- Valid transaction (Y/N)
- Smart Contract Condition Met (Y/N)

Yuma Consensus

- How "good is this model output?
- How accurate is this prediction?
- How fast is this compute?

-
- Creates a framework for measuring quality
 - Measurable tasks can be commoditized
 - Computation can occur off-chain





TAO Demand

- The TAO token is integral to the Bittensor ecosystem and is used across all protocol components:
 - **Governance** – Holding the TAO token provides governance rights, allowing holders to help shape the future of the Bittensor network and potentially alter protocol parameters or mechanisms.
 - **Security (Validator Requirements)** – To qualify as a Validator, one must have a minimum total stake weight of 1,000 TAO. There are a maximum of 64 validator slots per subnet, allotted to the 64 largest validators by total stake weight, driving competition for increased stake weight, and therefore holding more TAO.
 - **Security (Miner Registration Fees)** – Costs per subnet vary but are generally under 2 TAO and are recycled into the emission pool. There are 192 miner slots per subnet so increased demand could drive cost upwards.
 - **Subnet Registration Fees** – Subnet registration cost varies, with a minimum cost of 100 TAO. The current cost is approximately 400 TAO. Registration fees are fully refundable if a subnet is delisted.
 - **Transaction Fees** - Transaction fees on Subtensor are paid in TAO. The introduction of EVM compatibility on Subtensor could potentially increase the significance of transaction fees as developers look to build out DeFi capabilities on Subtensor.
 - **Economic Scarcity** – By adhering to a fixed supply schedule similar to Bitcoin's, Bittensor's TAO could accrue a monetary premium beyond its fundamental utility-driven demand.



TAO Supply Dynamics

- TAO's supply-side tokenomics were inspired by Bitcoin, with a total supply of 21 million tokens, and a predetermined halving schedule which occurs approximately every four years. The first halving event is expected to occur in December this year.
- At a price of \$204 per TAO, Bittensor has a circulating market capitalization of \$1.7 billion and a fully diluted valuation of \$4.3 billion.
- Approximately 60% of total TAO supply has not been issued yet. Of the approximate 8.5 million circulating TAO, about 72% has been staked.





Total Emissions

- Bittensor allocates TAO and Alpha tokens to the network to incentivize productive AI development. Emissions are facilitated via two separate phases: Injection and Extraction.
- 7,200 TAO and 14,400 Alpha tokens (per subnet) are emitted daily into subnet TAO and Alpha Reserves (subnet liquidity pools). Alpha Reserve injection is capped at 1 Alpha per block to prevent excessive inflation.

Injection


There are three components to the Injection phase, where liquidity is added over every block:

1. Tao Reserve Injection – TAO is injected to a subnet's TAO reserve in proportion to the price of its alpha token over the sum of prices for all the subnets in Bittensor.
2. Alpha Reserve Injection – Alpha is injected in proportion to the price of the alpha token so that the growth of a subnet's liquidity pool does not change the price of the alpha token.
3. Alpha Outstanding Injection – Over every block, liquidity is set aside to the alpha outstanding pool to be emitted to network participants at a rate equal to the alpha emission rate.

Extraction

At the end of each tempo (360 blocks), the quantity of alpha accumulated over the tempo is extracted by network participants in the following proportions:

Allocations





Dynamic Emissions

- Although total emissions are fixed, the amount that subnet participants (creators, validators, miners, etc) receive relative to other subnets varies according to Alpha token prices.
- Alpha token prices are a function of how much TAO has been staked to a subnet, providing an easy way to compare subnets against each other.
- If more TAO is staked into a subnet's pool, it indicates that validators are assigning more value to the subnet's outputs and therefore Bittensor should allocate more incentives for enhanced contributions. Conversely, if capital is flowing out of a subnet, its indicative of less valuable outputs, and therefore should be less incentivized.

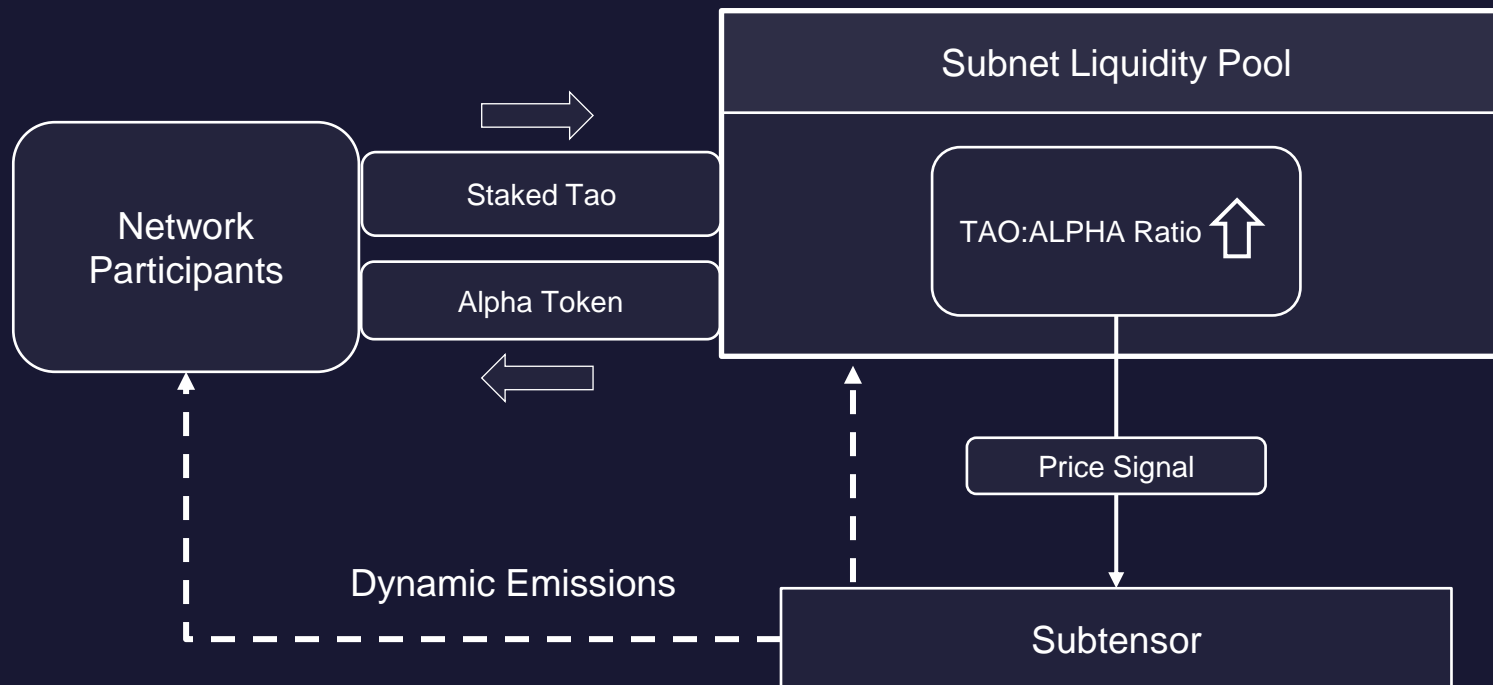




Table of Contents

1 Artificial Intelligence

2 What is Bittensor?

3 Network & Performance


4 Subnet Highlights

5 Looking Ahead



TAO Users


- Users across the Bittensor network have been steadily rising over the past two years, with total accounts surpassing 200,000. The continued development of new subnets, along with alpha token trading should help accelerate Bittensor adoption.





Top 10 Validators

- Bittensor has some notable large validators including the Opentensor Foundation, Yuma, and Polychain Capital. Multiple companies such as Yuma are fully dedicated to driving development and growth of the Bittensor ecosystem.






Crypto x AI

- Bittensor is the largest AI protocol by FDV (\$4.3B) and the second largest by circulating market cap (\$1.7B), only trailing NEAR's \$3.0B.
- Bittensor is uniquely positioned within the AI sector as its credibly neutral scalable design allows it to expand into a variety of subsectors within artificial intelligence, whereas many of the other protocols below are focused on singular use cases.

Top 10 Crypto x AI Protocols by FDV (\$B)






Agentic AI

- Bittensor is the only protocol on Cookie.fun, an AI agent analytics platform, to warrant its own category. Excluding memes and framework tokens (virtuals, ai16z, arc, etc.) Bittensor agents make over 5% of total agent market cap.
- Only seven out of Bittensor's eighty-two subnets having an AI-Agent-focus, displaying Bittensor's versatility and the scope of what the ecosystem can achieve across all different AI subsectors.

AI Agent Categories by Market Cap.





TAO Performance

- Since the Finney upgrade went live in March 2023, TAO has outperformed Ethereum by 167% but has trailed Bitcoin's performance by 26%.
- Investors have witnessed the proliferation of AI over the last two years, and Bittensor provides a unique way to gain broad AI and crypto exposure.

TAO vs Majors (Indexed to 100)



*Performance from the Finney upgrade (3.20.23) to 4.03.25

Source: Tradingview as of 4.03.25, FS Insight



TAO Performance

- Examining performance of some of the top AI names in crypto, TAO has outperformed them all over the previous two years, gaining 169% since March 2023.

TAO vs AI Basket (Indexed to 100)*



*Performance from the Finney upgrade (3.20.23) to 4.03.25

Source: Tradingview as of 4.03.25, FS Insight



Table of Contents

1 Artificial Intelligence

2 What is Bittensor?

3 Network & Performance

4 Subnet Highlights


5 Looking Ahead




Subnet Snapshot

- There are a total of 82 subnets built on Bittensor with a total market capitalization of \$1.66 billion. The Root subnet (SN0), which is used for staking TAO, is the largest subnet and comprises approximately 93% of the total subnet market cap.
- 2024 was a pivotal year for subnet growth, showing an increase in the pace of deployments, and subnets expanding from 32 to 82.

Total Subnets



Subnet Market Caps (\$B)





Subnet Ecosystem Map

AI Model Development		
SN	Name	Purpose
9	Pre-Training	Training Competition
29	Coldint	Collaborative Training
39	EdgeMaxxing	Model Optimization
11	Dippy	Roleplay LLMs
21	Omega	Multimodal
37	Fine-Tuning	Fine-tuning LLMs
5	Open Kaito	Text Embedding
35	LogicNet	Math & Data Analysis
56	Gradients	On-Demand Fine-tuning
58	Dippy Speech	Text to Speech

Distributed Training		
SN	Name	Purpose
38	Distributed Training	Distributed Training
3	Templar	Distributed Training

AI Data Pipeline		
SN	Name	Purpose
24	Omega Labs	Multimodal Data Collection
13	Dataverse	Scraping & Storage
33	ReadyAI	Data Structuring
40	Chunking	Chunking & RAG
42	Masa	Data Scraping
52	Tensorplex Dojo	Human-Generated Data
64	Chutes	AI Compute
51	Celium	GPU Compute

Generative AI		
SN	Name	Purpose
2	Omron	ZKP Inference
17	Three Gen	3D Content
18	Cortex.t	LLMs / Synthetics
23	SocialTensor	Image Generation
26	Storb	Image Generation
50	Synth	Audio Generation
46	Neural3D	3D Game Assets

Components & Tooling		
SN	Name	Purpose
47	Condense AI	Token Compression
31	NAS Chain	Neural Architecture Search
15	De-Val	LLM Evaluations
49	Hivetrain	Neural Network Search
81	Patrol	Decentralized Palantir

Scientific Research		
SN	Name	Purpose
25	Protein Folding	Protein Folding
43	Graphite	Algorithmic Optimization
36	Pyramid Scheme	Cellular Automata
76	Safe Scan	Cancer Detection
68	Nova	Drug Discovery

General Infrastructure		
SN	Name	Purpose
16	BitAds	Online Advertising
7	SubVortex	Bittensor Node Infrastructure

AI Powered Tools		
SN	Name	Purpose
34	BitMind	Deepfake Detection
61	RedTeam	Cybersecurity
32	It's AI	AI Text Detection
53	Efficient Frontier	Trading Strategies
8	PTN	General Trading Strategies
14	Paladn	Fraud Detection
22	Smart-Scrape	Twitter Data Analysis
54	WebGenieAI	Design-to-code
44	Score Vision	Sports Analysis
70	Vericore	Fact Checking
60	Bitsec	Code Auditing

Model Hosting		
SN	Name	Purpose
4	Targon	Rapid LLM Response
19	NineteenAI	Distributed Inference
75	Hippius	Cloud Infrastructure

Predictive Systems		
SN	Name	Purpose
6	Infinite Games	Events
30	Bettensor	Sports
41	Sportstensor	Sports
28	Foundry S&P	S&P 500
44	Precog	Bitcoin
48	NextPlace AI	Home Prices
57	Gaia	Geospatial
79	Taos	Trading Simulation

Agentic AI		
SN	Name	Purpose
1	Apex	Agentic Workflows
45	Gen42	Coding Assistant
62	AgentTAO	SWE Agents
59	Agent Arena	Agent Benchmark
20	BitAgent	Personal Assistance
10	Sturdy	DeFi Yield Optimizer
60	BitSec	Code Auditing


SN 63, 65-67, 69, 71-74, 77-78, 80-81 not depicted for lack of information or reserved for future use.

Source: Bittensor, TaoStats, Crucible Labs, FS Insight




SN19 - Nineteen AI

- SN19 provides optimized inference, offering decentralized access to the top open-source LLMs and image models. SN19 is developed by Rayon Labs, who is also behind SN64 and SN56.
- Miners earn rewards based on the speed and accuracy of inference requests, incentivizing them to host and serve high-demand LLMs and image models. Nineteen has outperformed the industry's top model-as-a-service providers, such as Together.ai and FireworksAI, achieving a 300 tokens/sec for the Llama 3.1 8B model at no cost to users.



SN64 - Chutes

- Chutes is a serverless AI compute subnet, allowing for deployment, running, or scaling AI models instantaneously without worrying about GPU infrastructure. Chutes enables enhanced efficiency, customization, and cost efficiency for AI developers through its platform which has over 4,000 GPU nodes.
- Chutes has integrated a large selection of AI models with a variety of code options from Python to cutting-edge LLM fine-tuning options. Chutes was one of the first to integrate DeepSeek's latest V3-0324 model, which boosted usage over the past few days.






SN42 – Masa Network

- Masa Network is a decentralized AI network incentivizing data contributions. Miners earn rewards by scraping data or providing compute services.
- Masa provides real-time and static, structured, annotated, and vectorized data from sources including Twitter, Discord, Telegram, TikTok, and more. All data sets can be used to optimize AI agents, LLMs, and AI applications, all accessible via API.
- Masa is also developing SN59 – Agent Arena , where AI agents battle for emissions based on engagement scores on X. Agents can integrate real-time data from Masa’s free Twitter API to become more contextually aware and boost their scores.

Masa Real-Time Data Performance



Twitter API Tiers				
	Free	Basic	Pro	Enterprise
Price	Free	\$100 / Mo	\$5,000 / Mo	\$42,000 / Mo
Post Limits	1,500	3,000	300,000	50+ Million

Masa API - \$0.00 / Mo

End-to-end enterprise grade API – Real time data scraping and all data in AI ready format.



SN41- Celium

- Celium is a peer-to-peer compute marketplace, delivering high-performance GPUs at prices far below those of traditional providers.
- Celium's platform provides real-time insights into GPU usage, helping developers monitor and optimize resource allocation for maximum efficiency and cost-effectiveness.

100+ Available GPUs

Hard-to-find GPUs

Access to the latest and highest-demand GPUs instantly, with an easy to manage platform for finding GPUs for client needs.

100% Uptime

Only GPUs that remain active and meet Celium's high performance standards remain available on the platform. Any GPUs with less than 100% uptime are automatically removed, ensuring extreme reliability.

GPU Prices

	Celium	Runpod	Paperspace	Coreweave
H100	\$1.35 / hr	\$3.29 / hr	\$5.95 / hr	\$4.76 / hr
A100	\$0.95 / hr	\$1.89 / hr	\$3.18 / hr	\$2.21 / hr
RTX A6000	\$0.38 / hr	\$0.69 / hr	\$1.89 / hr	\$1.28 / hr
RTX 4090	\$0.35 / hr	\$0.69 / hr	-	-
RTX 3090	\$0.22 / hr	\$0.43 / hr	-	-



SN8 – Proprietary Trading Network (PTN)

- Taoshi's PTN is at the forefront of innovation, merging cutting-edge decentralized technology with sophisticated trading strategies. By harnessing AI with the combined expertise of data scientists, traders, and machine learning teams, PTN creates a unique synergy between decentralized AI and quantitative strategies, empowering users to navigate financial markets with confidence.
- Miners contribute signals across financial markets, available for purchase on Taoshi's Request Network. PTN is a first of its kind decentralized prop firm, offering unparalleled benefits for both miners and users. Miners receive incentives immediately after signals are provided, and users receive highly performant institutional grade insights, reduce overhead, enhance trading performance.

Top 10 Profitable Miners				
Miner Rank	Miner Address	Annualized 90D Return	Sharpe	Sortino
1	5GZRzpXo97cZVshHAUCpBGxavG7PCJXti13vjQhX77vjrXhH	14.4%	6.22	10.18
2	5CyzSyxxCn1CFiSAn7rfYsRoiqHTR2623137if4fQZnwYJ7f	10.6%	3.28	6.43
3	5CcNVDt7YLa8YbyUJxS7TZ9y5gsR3qNq3rbKGU6B4A4H541W	22.8%	3.77	7.20
4	5DswH2LoRwivzv37tGvo27XJdVfPff2fhu5T8LHsfXmFS5u	17.0%	6.41	12.83
5	5DCaKgSaNxvTdy3kRBRA1BSpdRLs4YDstUsZmmS6HbDGt3eK	11.1%	3.10	4.77
6	5CktpCybceQN8BZCX9ahyHMBsMYsDEB7F3RrMHZ4LK6nPwG9	34.0%	2.21	4.97
7	5Cz4ETkXaZfjTxDLJuZiP9eyZbr7SrEjRwQCT3hLNyjCZLV1	19.7%	2.33	3.29
8	5FRpp1he5F1Z5wc7CVNfepoajXinwNAZAxh1RC4pEbQXj8n9	8.6%	2.00	2.60
9	5HVnqF3ibfVedMPgVUDxJ2CTabu9metpXv6PtbttohBoD7gCb	36.3%	2.18	6.17
10	5EsqVEf6zZ1zs3A27HUjmVWh4Yd2KrFQKFndUFEGAN9u48t	18.1%	2.75	3.64

PTN Flywheel





SN41- Sportstensor

- Sportstensor is an institutional-grade sports predictions platform operating through the world's first decentralized intelligence network. The Sportstensor Meta Model (STMM) is Sportstensor's advanced ensemble architecture that aggregates predictive outputs from top-performing miners across each league on SN41.
- Utilizing a dynamic weighting algorithm, the system evaluates historical accuracy and calibration metrics to synthesize consensus win probabilities. Sportstensor's STMM models has significantly outperformed the NBA market over the last three months.






Table of Contents

1 Artificial Intelligence

2 What is Bittensor?

3 Network & Performance

4 Subnet Highlights


5 Looking Ahead




Developer Talent & Bittensor DeFi

- Traditionally, the problem with open-source development has been the inability for developers to fund their work and coordinate contributions. Bittensor's economic model aligns incentives between developers and the overall ecosystem, leading to a coordinated effort in creating useful AI services.
- As subnet development continues to grow and services or products improve, incentives should continue flowing to developers who are constantly innovating, potentially igniting a flywheel of attracting top-tier developer talent who want their share of incentives.
- Full EVM compatibility is now available on Subtensor, opening the door to a larger crypto-native developer base and for out-of-the-box application deployments. A Bittensor defi ecosystem could develop with unique innovations such as revenue share agreements between miners facilitated via smart contracts or TAO/Alpha liquid staking protocols.

Dev. Talent Flywheel



Bittensor DeFi






Alpha Tokens Provide Venture-Like Risk Profiles

- Excluding SN0, the average market cap across alpha tokens is \$1.85 million, with Chutes being the largest at \$24.3 million. Alpha tokens are extremely cheap compared to public or private centralized AI companies.
- Venture-like investors may focus their attention on the Bittensor ecosystem, fundamentally valuing subnets to find cheap exposure to innovative AI products.
- Bittensor also has a competitive advantage in quickly launching specialized subnets to meet new use cases, potentially providing a liquid and accessible way to be early to cutting edge technology.

*Top 10 Subnets by Market Cap (\$M)



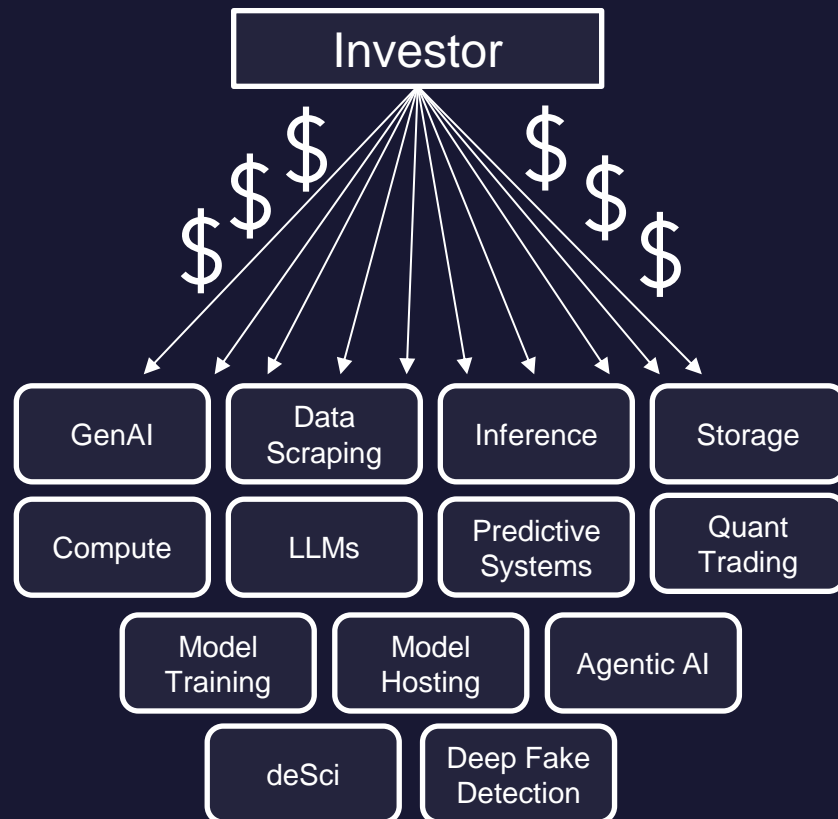
*Excludes Root Network

Source: Taostats 3.26.25, FS Insight



TAO Encompasses the Whole AI Stack

- The gateway to the Bittensor ecosystem flows through the TAO token with participants needing TAO before being able to access alpha tokens. For those not interested in making individual subnet bets, TAO may become the best way to gain broad AI exposure.
- Most subnets have yet to introduce monetization methods, but if products can find meaningful traction, profits can be distributed to TAO and Alpha holders, initiating an entirely new incentive stream that would only stimulate the Bittensor ecosystem further.





Investment Thesis








- TAO addresses the coordination and monetization challenges inherent in open-source development. As a novel mechanism for incentivizing AI contributions, Bittensor is well-positioned to attract a vibrant developer base, accelerate AI innovation, and expand growth across its subnet ecosystem.
- As the ecosystem expands, Bittensor's protocol mechanics generate structural demand for the TAO token. Increased subnet deployments fuel demand through registration fees and competition for validator and miner slots, creating buying pressure for TAO.
- From an investment standpoint, TAO serves as the gateway to the Bittensor ecosystem, providing unique exposure to a diversified AI ecosystem through a single investment vehicle. For more venture-oriented investors, purchasing TAO is a prerequisite to investing in individual subnet tokens.
- As AI becomes increasingly politicized and central to geopolitical competition, accessibility restrictions across borders may emerge for both users and developers, underscoring the importance of a permissionless global alternative. While Bittensor still possesses elements of centralization, its long-term vision is to evolve into a decentralized, credibly neutral network for AI development.
- **Bottom Line:** TAO offers investors a unique opportunity to gain diversified exposure to open-source AI development through a single asset, with built-in demand drivers, exposure to nascent subnet projects, and alignment with a long-term vision of creating a permissionless, credibly neutral AI network.



TAO Illustrative Valuations



Price: \$204
FDV: \$4.3B

Company	Funding	Valuation	Equivalent TAO Price	Multiple From Current Price
 Mistral AI	\$1.2B	\$6.5B	\$310	1.5x
 Perplexity*	\$0.5-\$1B	\$18B	\$857	4.2x
 Coreweave	\$13.5B	\$23B	\$1,095	5.4x
 xAI	\$12B	\$50B	\$2,381	11.7x
 Anthropic	\$16B	\$62B	\$2,952	14.5x
 Databricks	\$18.6B	\$62B	\$2,952	14.5x
 OpenAI	\$24B	\$300B	\$14,286	70.0x

*Estimated funding and valuation



Grayscale Bittensor Trust

- Grayscale Bittensor Trust (the "Trust") is one of the first investment vehicles that enables investors to gain exposure TAO, the underlying token of the Bittensor Network, in the form of a security while avoiding the challenges of buying, storing, and safekeeping Bittensor directly.
- The Grayscale Bittensor Trust is designed to track the Bittensor market price, less fees and expenses. To invest or learn more, please direct inquiries to info@grayscale.com.



Ticker	N/A
Investment Objective	Shares reflect the value of the TAO held by the Trust, less expenses and other liabilities
Inception Date	6/10/2024
Subscriptions	Daily
Sponsor	Grayscale Investments Sponsors, LLC
Legal Counsel	David Polk & Wardwell LLP
Auditor	KPMG LLP
CUSIP	38963B 100
ISIN	US38963B1008
AUM*	\$4.44M
NAV Per Share*	\$3.87
Minimum Investment	\$25,000
Sponsor Fee	2.5% Annually
Secondary Trading	Following a one-year holding period, Grayscale intends to attempt to have Shares of new products quoted on a secondary market. However, there is no guarantee that we will be successful.
Private Placement	Available for Accredited Investors
IRA Eligibility	Shares are eligible to be held in certain IRA, Roth IRA, and other brokerage and investor accounts.
Redemptions	Redemptions of Shares are not currently authorized.
Contact	info@grayscale.com



Risks

- The Bittensor ecosystem is primarily driven by emissions of TAO and Alpha tokens. Roughly \$2 million worth of TAO is distributed daily, creating persistent sell pressure and weighing on price in the short term. A halving event scheduled for December should help reduce this pressure, but meaningful emissions will remain a structural headwind for the foreseeable future.
- While Bittensor's economic model addresses key pain points in open-source development, large and well-capitalized tech companies continue to benefit from superior hardware, deeper developer pools, and access to proprietary datasets, potentially giving them an edge in scaling AI.
- Subnets have shown early signs of technological promise, but without meaningful monetization, confidence in the broader ecosystem and the TAO token could erode. We believe the TAO ecosystem must establish a clear path to revenue generation, with mechanisms to pass fees on to miners and validators. Without this, TAO may face continued downward pressure.
- Bittensor's long-term strategic goal is to become a credibly neutral AI development tool. While Bittensor may move to further decentralize the network in the future, there are centralization risks in its current state. The OpenTensor foundation is the sole party responsible for validating blocks on Subtensor. Additionally, the top 10 largest subnet validators comprise about 67% of total network stake weight, with OpenTensor also making up 11%.
- Like many early-stage crypto projects, Bittensor still faces significant technological risk. While diversification can mitigate the impact of isolated incidents, the risk to investors remains. Crypto assets are inherently high-risk and continue to exhibit strong correlation to the broader macroeconomic backdrop. Any shifts in global growth or inflation that could weigh on risk assets may negatively affect TAO price.



Disclosures

This research is for the clients of FS Insight only. For important disclosures and rating histories regarding sectors or companies that are the subject of this report, please contact your sales representative or FS Insight at 150 East 52nd Street, New York, NY, 10022 USA.

Analyst Certification (Reg AC)

Tom Couture, the research analyst denoted by an “AC” on the cover of this report, hereby certifies that all of the views expressed in this report accurately reflect my personal views, which have not been influenced by considerations of the firm’s business or client relationships.

Neither I (Tom Couture), nor a member of my household is an officer, director, or advisory board member of the issuer(s) or has another significant affiliation with the issuer(s) that is/are the subject of this research report. There is a possibility that we will from time to time have long or short positions in and buy or sell, the securities or derivatives, if any, referred to in this research.

Conflicts of Interest

This research contains the views, opinions, and recommendations of FS Insight. This report is intended for research and educational purposes. Grayscale is a client of FS Insight and receives FS Insight’s published research reports. Grayscale has commissioned/paid for this specific research project and has permissions to distribute. This report is intended for research and educational purposes.

Analyst Industry/Sector Views

Positive (+): The analyst expects the performance of his industry/sector coverage universe over the next 6-18 months to be attractive vs. the relevant broad market benchmark, being the S&P 500 for North America.

Neutral (N): The analyst expects the performance of his or her industry/sector coverage universe over the next 6-18 months to be in line with the relevant broad market benchmark, being the S&P 500 for North America.

Negative (-): The analyst expects his or her industry coverage universe over the next 6-18 months to underperform vs. the relevant broad market benchmark, being the S&P 500 for North America.

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 June 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 June 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 research is for the clients of FS Insight only. Additional information is available upon request. Information has been obtained from sources believed to be reliable but FS Insight does not warrant its completeness or accuracy except with respect to any disclosures relative to FS Insight and the analyst’s involvement (if any) with any of the subject companies of the research. All pricing is as of the market close for the securities discussed, unless otherwise stated. Opinions and estimates constitute our judgment as of the date of this material and are subject to change without notice. Past performance is not indicative of future results. This material is not intended as an offer or solicitation for the purchase or sale of any financial instrument. The opinions and recommendations herein do not take into account individual client circumstances, risk tolerance, objectives, or needs and are not intended as recommendations of particular securities, financial instruments or strategies. The recipient of this report must make its own independent decision regarding any securities or financial instruments mentioned herein. Except in circumstances where FS Insight expressly agrees otherwise in writing, FS Insight is not acting as a municipal advisor and the opinions or views contained herein are not intended to be, and do not constitute, advice, including within the meaning of Section 15B of the Securities Exchange Act of 1934. All research reports are disseminated and available to all clients simultaneously through electronic publication to our internal client website, FSInsight.com. Not all research content is redistributed to our clients or made available to third-party aggregators or the media. Please contact your sales representative if you would like to receive any of our research publications.

Copyright 2025 Fundstrat Global Advisors LLC. All rights reserved. No part of this material may be reprinted, sold or redistributed without the prior written consent of Fundstrat Global Advisors LLC.