

Liquid Opps: AI x Crypto (Apr'23)

Akash (\$AKT)

- Akash is a decentralized compute network that is adding GPU support for AI/ML workloads. Since launching Oct'20, the Akash network has grown to 36 miners with 2.6k CPUs (40% utilization), 5TB of RAM (25% utilization), and 80TB of storage (10% utilization). \$AKT is currently valued at \$65m (\$115m FDV) and generates \$25k of annualized protocol revenues, a premium multiple of token burn (2600x) relative to its peers Arweave/Render/Filecoin (150-400x). **Nov'23 Update: \$AKT is up ~5.5x and Akash revenue is up ~33x since the time of writing, compressing the multiple down to ~450x.**
- (+) Margin-of-safety at the current token price (\$0.30). The pre-mine is fully-vested, and go-forward inflation is low-teens and only to miners. Investor allocation (35m tokens with a cost basis of \$0.02-0.04) has been fully-unlocked for six months and price has been range-bound at \$0.20-\$0.40. Akash is the only token on our list trading below the public sale price of \$0.35, and there are no big-name crypto VCs on the cap table. All of this leads us to believe \$0.30 is an attractive entry price, plus \$AKT is already listed on several major exchanges. **Nov'23 Update: several liquid funds have built \$AKT positions this year on the way up (including Modular, Pangea, Tribe) but we still consider it to be an underweight holding across the industry.**
- (+) Leadership team is fully-focused on GPU support (see: [founder tweets](#)). That said, none of the leadership team has background in AI/ML - they mostly worked on cloud infra and dev tooling - and they seem to be relying on the community to bring AI/ML expertise. Can we trust a "non-native" team to be ahead of the curve in such a fast-moving space? Just to give an example: the Google Form sent out to potential testnet miners a few weeks forgot to specify type of GPU. **Nov'23 Update: the space has gotten markedly more competitive since April, with VCs funding new private competitors including IONet (Multicoi) and Hyperbolic (Faction).**
- (-) Progress on AI/ML has been slower than expected. Since adding GPUs to the roadmap Q1'21 and originally guiding towards Q1'22 launch, Akash's private GPU testnet is just now about to launch in Q2'23. We understand the delay is primarily due to lack of technical infrastructure - there is no secure containerization platform like Gvisor for GPUs - which the Akash team ultimately resolved by working with Kontain. The team also had competing roadmap priorities from the "legacy" (non-GPU) network, namely persistent storage and IP leases. Conceptually, it feels a lot like the path to launching Helium 5G/Mobile vs Solana/IoT. **Nov'23 Update: Akash successfully launched GPUs, which is now contributing a major chunk (on some days a majority) of network revenues.**
- (-) Upcoming catalysts are more hype than substance. Overclock Labs launched an internal GPU testnet two weeks ago and expects to launch a public testnet imminently. The testnet will prove that Akash has the technical infrastructure to stitch together GPUs, but it doesn't prove there's a working consensus/incentive mechanism to bootstrap the supply-side nor a plan for how to sell to AI/ML end-users. It could take years before the GPU network is a meaningfully contributor to \$AKT token burn, and there's no guarantee that shipping a major roadmap item will move prices regardless of how hyped it is (e.g., \$FIL is roughly flat since launching FVM). **Nov'23 Update: Akash's has grown into its revenue multiple somewhat, currently trading around ~450x (roughly double Filecoin's multiple) with stronger growth.**
- (-) AI/ML could look a lot different by the time Akash GPUs launch. Akash's go-to-market strategy on the supply side involves partnering with large datacenter operators' excess GPU capacity (similar to their current network which has only 36 miners) and focus on inference vs training/tuning. On the supply side, demand for Nvidia GPUs is off the charts right now, so excess capacity is less of a problem now than ever before. On the demand side, LLMs with dramatically lower computing requirements have launched over the past few weeks, like GPT4All and MiniLLM, which can run locally

on a macbook, [iPhone](#), or even a [TI-84](#). In short, Akash is building a two-sided network from scratch, in a space where supply/demand are both rapidly-evolving, and they have no advantaged access to capture either side.

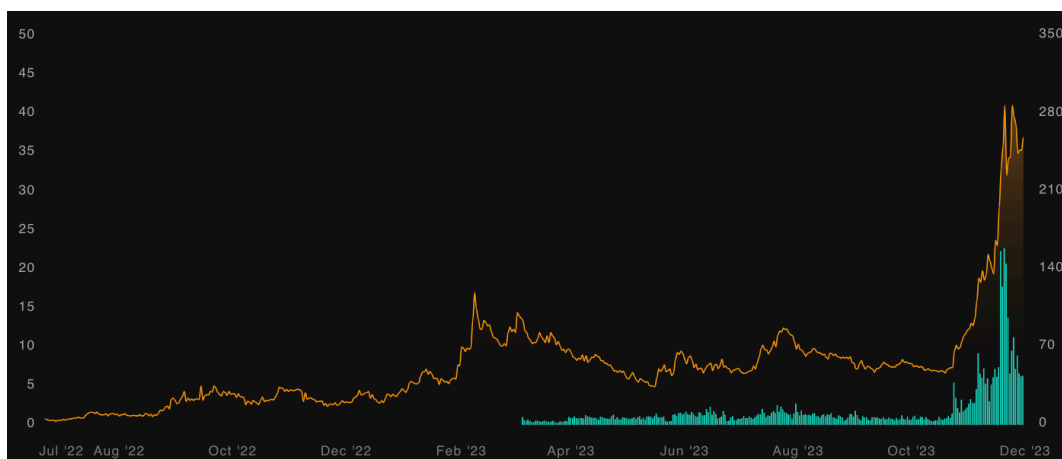
- Where are the tokens? 210m \$AKT have been minted - 110m to miners, 36m to investors, 28m to the foundation/ecosystem, 27m to the team, and 9m for testnets/vendors - of which 143m (68%) are staked with a 3-week cooldown; of the remaining 67m tokens, average daily CEX trading volume is 4m tokens (6%). There's 4.5m AKT tokens in liquidity pools against ATOM/OSMO (\$2m+ TVL) that should allow us to get in and out of a small position entirely on-chain.

BitTensor (\$TAO)

- BitTensor is a decentralized neural network that incentivizes machine intelligence. Miners (GPUs) run inference on ML models and validators (CPUs) query models and rank responses. Validators can run apps, like [ChatTensor](#), where users must delegate \$TAO to the validator in order to use the app. When a user enters a prompt, ChatTensor's validator queries the network of miners (via [ensemble learning](#)), combines and scores responses relative to their utility (via [proof-of-intelligence](#)), and presents the results to end-users. The protocol splits block rewards 50/50 between validators and miners: the former weighted by delegated \$TAO (i.e., usage) and the latter weighted by validator scores in proof-of-intelligence consensus (i.e., utility). ELI5: BitTensor is a two-sided network of ML apps and models, where the most used apps and the most useful models drive value to and from each other. **Nov'23 Update: BitTensor has become a "network of networks" after decentralizing its consensus mechanism and enabling permissionless creation of AI subnets.**
- (+) Unlike other "AI x crypto" teams, BitTensor's founding team are true AI/ML builders. Out of 15 people on the core team, 5 are machine learning engineers including from [Instacart](#), [Workday](#), and various research [universities](#), and the founders were ML engineers at [Google](#) and IBM. **Nov'23 Update: we've spent meaningful time with the founders over the past few month - including hosting them on [our podcast](#) - and believe they are intellectually and ideologically within the top 0.1% of crypto founders.**
- (+) The BitTensor network is already at breakout scale. The network has 100+ active validators and 3400+ active miners. Based on conversations with miners, we estimate \$10m+ in GPU spend since its launch in Nov'21 (comparable to an estimated \$12m spent training GPT3) and \$30m+ worth of GPUs mining the network. Mining power is growing rapidly: at the start of the year a 2.7B parameter LLM earned competitive rewards; today anything below 6B+ parameters is left behind. The mining community is intensively competitive and the network has already survived a number of attacks, including an attack from one of the core developers that led to hard fork in 2021 and another (potential DDOS) attack in Jan'23. **Nov'23 Update: BitTensor's momentum has only accelerated. There are now 32 different subnets running AI services, and the largest individual subnets are generating >\$30m in annualized incentives rewards (the size of the entire network back in April).**
- (+) BitTensor's community is fiercely committed to fair token distribution. BitTensor's founder is a self-proclaimed Bitcoin maximalist. BitTensor's issuance schedule mirrors Bitcoin, with a 4-year halving schedule, no pre-mine, and 21m max supply. The core team and investors have mined their tokens vs receiving zero cost-basis insider allocations. That said, inflation is significant at a level similar to BTC in 2012, i.e. 45% annualized as circulating supply goes from 4.3m today to 10.5m at the first halving in 2025. Given the certainty in the supply curve, we think it's reasonable to underwrite \$TAO based on a 10.5m supply in Q3'25 (vs 21m cap). **Nov'23 Update: with the increase in registration fees post-upgrade, the first token halving has been pushed out marginally to Q4'25 or early Q1'26.**
- (-) On the other hand, demand is still far away. The only real application is [ChatTensor](#), which today feels like a clunky token-gated UI on top of GPT-2 (new version coming end of April). The core team is only 15 people, compared to 85 for Akash and 55 for Fetch, and there are only a handful of professional third-party developers on the network. We have seen crazy active AI developer

communities on, and Bittensor is not that—yet. However, the network is limited in functionality to date with only a GPT-2 style LLM: the recent launch of subnets allows developers to experiment with prompting, text-to-image, spacial recognition, and other multi-modal models that all learn from each other. **Nov'23 Update: BitTensor is now undoubtedly an emerging ecosystem that is attracting top-tier AI developers. In the first month since the upgrade, developers have built a NumerAI competitor, multi-modal search, natural language translation, image generation, model pre-training, and more.**

- (-) Unproven path to decentralization. Although it takes inspiration from Bitcoin, BitTensor is a delegated proof-of-stake protocol at its core. Tokenholders delegate to validators who provide the most value for the ecosystem, first and foremost the BitTensor foundation (currently 21% of stake). This is the only funding mechanism for the BitTensor Foundation (core team), which at current prices generates \$3m of 'ARR' from staking commissions. Other major delegates, all with \$0.1-1m 'ARR', include TaoStats analytics platform, NeuralInternet developer DAO, Tensor.Exchange OTC, TaoBridge cross-chain bridge, and Tao-Validator and TaoStation validators. This structure could slow the network down if the core team can't get requisite funding to keep developing the protocol. **Nov'23: BitTensor decentralizing its consensus mechanism in October; now third party developers can spin up subnets to compete directly with the Opentensor Foundation.**
- (-) Lack of near-term catalysts. The protocol just completed the Finney hardfork, which is analogous to HIP-51, ie going from one protocol (text) to multi-protocol (text/image/video). This was originally planned to launch alongside a move to Polkadot, which would have had numerous benefits - e.g., smart contracts, DeFi ecosystem, shared security - but the core team chose against it at the last second and decided to continue as their own substrate-based L1. The roadmap for 2023 is primarily about incremental protocol features (i.e., launching models to reach parity with ChatGPT-3.5 and Stable Diffusion). Also, the token only listed on tiny CEXs (Bitget/MXC) and the team has been publicly adamant about not pursuing exchange listings themselves. Like Bitcoin, we don't expect BitTensor to have many "catalysts"... just the slow, steady grind of hashpower growth. **Nov'23: Bittensor revolution was an inflection point; since then, developers have built dozens of AI applications on top of the network. At least one of these applications will achieve 1m+ users in 2024.**
- (-) Today, mining is more attractive than buying tokens. BitTensor mining is like horse racing: first you need a horse (in this case, a dataset); then you train the horse (in this case, training a model on GPUs); then you register for races and/or participate in prelims (in this case, pay a registration fee or participate in PoW); and finally the horse can compete in races for cash prizes (becoming a miner and earn \$TAO rewards). Today, datasets can be downloaded for free on HuggingFace or ScaleAI; registration fees are \$250-500 at current \$TAO prices; training a competitive model takes \$3-5k on platforms like Runpod (7-10 days of training on 8x A100s at \$2/hr). The average miner earns 1.1 \$TAO daily, which represents a cash payback of less than 2 months at current prices with no physical/operational component (so long as you can rent GPUs). We estimate new miners today are mining at a \$25 cost basis vs the current market price of \$60. **Nov'23 Update: mining is still highly attractive IF you have differentiated datasets, models, or inferencing capabilities. It has become extremely difficult for 'retail' miners to earn outsized rewards.**
- The vision is exciting but the network is early so we need to be mindful of entry prices. We need to understand: 1) does the consensus mechanism (proof of intelligence) work from technical perspective? 2) are there signs of life in the developer community? 3) who are the key players in this community? I suggest adding a position if token price approaches 1-year miner breakeven and as we get conviction on these two points. \$TAO is up +110% YTD vs 70% for \$BTC. **Nov'23 Update: \$TAO is up 4x since the time of writing. Given the pace of fundamental progress, \$TAO supply dynamics (89% staked), and increasing attention to the project as the leading AI token, we believe there will be opportunities to buy \$TAO at extremely attractive risk/returns over the next few months.**



Not Investable

- Petals (4.4k stars) is a BitTorrent-style protocol that for inference and fine-tuning of LLMs, backed by HuggingFace. Currently there's 7 nodes on the network and a rudimentary chatbot app. They don't have a token per se but last week they do have an incentive mechanism which they explicitly state is not crypto related: "We are a centralized incentive system, even though Petals is a fully decentralized system in all other aspects. We do not plan to provide a service to exchange these points for money, so you should see these incentives as 'game' points designed to be spent inside our system."
- KoboldAI (1.7k stars) is another BitTorrent-style protocol, but with a token incentive mechanism. Nodes can contribute GPU capacity or request it for inference on an LLM, with points determining priority in the case of network congestion. There's currently 21 nodes on the network and a rudimentary chatbot app. The founder is also involved in the BitTorrent community.
- Learning@home (1.4k stars) is a SETI-style protocol for training large neural networks, different from other projects' focus on inference/tuning. It's particularly useful for training LLMs in other languages (e.g., Arabic) where there are fewer geographically concentrated pools of GPUs. Development began in 2020 and it now has a handful of projects building on top of it.