Supporting Document

Why the Targon Virtual Machine (TVM) Matters and how it solved confidential computing for decentralized systems

Information here: https://www.manifold.inc/releases/targon-v7

Targon from inception

Targon is the oldest subnet on Bittensor - registered on October 3rd, 2023. Subnets 1-3 have all changed owners since launch. Besides the longstanding experience this gives us within the Bittensor ecosystem (our co-founders are former OTF executives), it also means we have an alpha token with deeper liquidity (making it less volatile) and more emitted tokens meaning less dilution for new holders moving forward.

Given our strong reputation and experience operating within Bittensor, other teams are even soliciting advisory from Targon/Manifold on how to build great products and effectively compete in the ruthless competition that is Bittensor.

$10M Series A

Manifold Labs, the operator of Targon, raised a $10.5M Series A to provide runway, fund growth and expand our team. This shows long-term alignment and belief in our vision from some of the most well-renowned investors in the world, as well as provide widespread industry connections.

Announcement: https://www.manifold.inc/releases/series-a-announcement

---

How Targon incentivizes low prices.

For each compute market, there is a specific emission allocation which is up for grabs. The miners will place bids on what hourly price they are looking to get paid for their type of compute. The lowest bids get filled first until the allocation is gone for that interval, any bids above a price that cannot be filled in full, will be diluted. This incentivizes compute providers to bid prices down, which allows us to offer compute at low prices to our customers.

Revenue

Targon is already serving real customers and generating real revenue. Some of our existing customers include fellow subnets Ridges (SN62), 404Gen (SN17), Affine (SN120), Babelbit (SN59), Dippy (SN11), Quasar (SN24) and BrainPlay (SN117) to name a few.

We also have begun research projects powered by Targon compute with several universities including University of Michigan and UC Irvine.

With the global compute market expected to be well into the trillions by 2030, and compute being the currency that fuels all of AI, we are operating within one of the largest and fastest growing industry verticals in the history of mankind.

Targon Serverless SDK

The Targon Serverless SDK allows developers to easily access compute natively to their working environment using Python. It scales up and down with your usage, so you only pay for the exact amount of compute that you need.

Different types of compute available

CPU Servers - Used for lightweight applications, hosting websites, running chatbots, validating other blockchain networks, etc.

RTX4090s - Consumer GPU used for tasks such as running local LLMs, image generation, gaming, video editing, etc. Popular for home labs and small-scale inference.

H100s - Professional AI training and inference at scale, training LLMs, running production AI, large-scale simulations.

H200s - Similar to H100s but with significantly more memory, making them better for training/serving very large models, especially those with massive context windows or parameter counts that don’t fit on H100s.

B200s - Next-gen GPUs (Blackwell architecture) designed for even larger models and more efficient training/inference than the Hopper series (H100s/H200s). Used for frontier AI research, training next-gen models, and ultra high-performance inference.