How the TEA Project Solves a Central Blockchain Problem

As blockchain gains popularity for developers seeking decentralization and security, the trade-off for them coming from the centralized cloud is that they must sacrifice speed. Traditionally, blockchain is slow, which prevents apps from running at the same speed as traditional cloud computing apps. Many developers have accepted slow app response as a tradeoff for entering the decentralized blockchain sphere. They figure some slowness is worth it to avoid the centralization of power and security breaches that expose customer data that are inherent with existing cloud computing.

While traditional cloud apps achieve their speed at the high cost of centralization, the TEA Project's dApps (known as TApps) run full speed while still remaining decentralized. The TEA Project is able to thread the needle of the blockchain trilema and achieve all three facets - decentralization, speed, and security - through its two-layer blockchain design.

- The layer-1 blockchain keeps trust data on the mining nodes running on layer-2. This layer deals with Byzantine fault and ensures that all validated layer-2 nodes are trusted.
- TEA Project's layer-2 can ignore Byzantine faults as the decentralized mining nodes on layer-2 have already gained trust from layer-1. It runs encrypted app code in the TPM-protected enclaves of the mining nodes that interact with encrypted customer data. The system is designed such that neither the miners nor the developers have access to the customer's data.

The trust certification kept by layer-1 on all layer-2 nodes is how the TEA Project solves the quandary of how to keep apps decentralized without sacrificing speed. Rich applications can run at cloud computing performance and scale because the TEA Project's layer-2 no longer has to deal with Byzantine faults.

Applications can run full speed on the TEA Project's layer-2 because, from the app's point of view, it looks just like a normal three-tier cloud computing architecture. And developers themselves don't need any special knowledge about blockchain or consensus algorithms to get the decentralization and security benefits inherent in the TEA Project's 2-layer blockchain. They just need to set WebAssembly as their code compilation target, allowing them to use their preferred programming languages. TApps follow the same three-tier architecture prevalent in cloud computing (presentation tier, app tier, and database tier), but TEA Project TApps run fully decentralized.

How The TEA Project Looks in Use

The TEA Project is looking to disrupt the world of cloud computing by bringing decentralized apps running on blockchain at cloud computing speeds. Because we no longer have the cloud infrastructure, we must incentivize miners to provide the underlying infrastructure. IPFS is the actual storage infrastructure for storing both encrypted app and client data, which miners will host within their protected enclaves when their node is selected for a task. Let's look at the economic incentives for a hypothetical TApp running on the TEA network:

- The **developer** uploads their Tensorflow image recognition analysis TApp to IPFS.
- A **photographer** uploads their pictures to IPFS.
- A **researcher** wants to use the **developer's** TApp with the **photographer's** pictures.
- A random **miner** is selected to host the Tensorflow TApp and the pictures within the protected enclave of their mining machine.

In economic terms, the **researcher** pays to use the TApp. The payment flows through the **miner** who provided the node to perform the secure and private calculation, the **photographer** who is paid for the use of their pictures, and the **developer** who is paid for the usage of their code. Every transaction enriches a wide variety of participants in the TEA ecosystem and helps sustain and incentivize further usage.

The TEA Project is able to provide all participants a decentralized computation network that is both secure and private. In practical terms, that means the developer who wrote the code cannot breach the photographer's image. And the miner who is randomly selected to host it in their enclave cannot access either the photographer's data or the developer's code. And the researcher is only able to receive the result of the app's calculation and isn't able to access the photographer's data outside of the enclave. The security and privacy that the TEA Project's decentralized computing network is able to provide has wide range of emerging use cases like residential IoT data and patient health records.

Because the TEA Project is able to provide the underlying trust, there's no longer any trust required between all four parties.

The TEA Project's Two Tokens: TEA & CML

The TEA Project uses a dual token setup to incentivize participation among miners, developers, investors, and consumers in the TEA ecosystem.

TEA

The first token, TEA, is a stablecoin utility token that must be mined - there's no TEA token in the genesis block (except for a very small amount for the initial gas fee), nor will there be any airdrops. TEA is generated as the reward given to miners for running security validation services on the TEA network, and TEA is burnt when users buy new CML tokens. Its supply is therefore influenced by the supply and demand of CML tokens which miners need in order to mine TEA.

TEA has a variety of uses on the platform:

- TEA pays for gas when performing any transaction on the network.
- TEA is used by consumers to pay for TApps they want to use.
- For miners, TEA can be part of the reward for mining on the network.
- Users can also stake TEA to mining nodes to share in a percentage of that node's mining rewards.

TEA is a stable coin that's pegged to the cost of computing resources. That means the same amount of TEA tokens should be able to pay for the same amount of computing resources at any point in time.

CML

The TEA Project's second token, CML, is an NFT with a life cycle determined randomly via algorithm. CML, also known as Camellia, has some unique properties and uses on the TEA Project platform:

- A TEA mining node can only be activated by associating a CML NFT with it. CML functions as a mining license and a credit record. Note that mining TEA doesn’t require a GPU or an ASIC. It only requires the Camellia NFT and cheap secure hardware such as an RPi with a GPS module and a TPM chip. The TEA Project will offer a TEA Box Raspberry Pi mining machine in the future complete with TPM chip, GPS module, and preinstalled with a type B CML seed necessary for mining.
- Miners buy new Camellia seeds through open bidding, with the TEA exchanged for them burned by the TEA Project DAO.
- Investors can stake their defrosted CML to mining machines for revenue sharing.
- CML is issued by the DAO at a rate to regulate supply and ensure enough uptake by miners and investors to avoid flooding the market. The DAO also burns CML when it when it dies.

Although CML has a lifespan, it can remain frozen past its defrost date for an indefinite amount of time. Investors who have participated in funding rounds in exchange for CML seeds can therefore keep their seeds frozen if they don't wish to mine with them.

Each TApp Can Have Its Own Token

The TEA Project is designed to allow for TApp token bonding curves for TApp developers and other creators to monetize their ideas. Investors back projects by purchasing TApp tokens on its bonding curve. TApp revenue, if set by the developer, can also go directly to the TApp token bonding curve. The gains from consumer app usage are therefore shared by the app developers, the hosting miners, and the investors. Additionally, any investors or enthusiastic users of the TApp can share in its profits by investing directly in the TApp's token.

TEA Ecosystem Development

Besides early investors, we envision the TEA Project ecosystem being built through successive outreach to the following three demographic groups: miners, developers, retail investors, and end-users.

Miners Prior to the maturity of the TEA Project's Web3 Rich dApps ecosystem, a mining economy is necessary to keep the TEA economy running. TEA's carefully designed token economy creates NFT scarcity during mining. The scarcity encourages miners to reinvest their harvested TEA back into CML instead of selling TEA. Miners are of course free to host whatever TApps they wish with free market principles guiding them towards hosting TApps that are popular and able to reward miners.

Developers After attracting miners to the platform, the TEA Project will begin outreach to developers to build TApps using the TEA dev framework. Developers upload their apps to the TApp market and wait for miners to host them. The outreach phase will feature developer tutorials and emphasize the ease of compiling to the WebAssembly format used by TApps.

Retail Investors Retail investors want to enjoy capital appreciation through investment. We expect retail investors to be interested in CML just like institutional investors, but there's also another way of participating in the TEA ecosystem as a retail investor: TApp tokens. Each TApp will have its own TApp token that investors can purchase as if investing in a stock. Each TApp token represents a share in the TApp and gives the holder dividends in the form of more TApp tokens issued to its holders whenever someone uses the TApp. As these tokens are issued on a bonding curve, the earliest TApp token investors will enjoy the largest amount of price appreciation.

End-users Concurrent to attracting developers to the platform, the TEA Project will actively market emerging TApps to consumers. We imagine the TEA Project being an ideal ecosystem for existing cloud apps looking to migrate and benefit from decentralization and data privacy. We look forward to the types of decentralized apps that will flourish on the TEA Project network. We also want to play an active part in welcoming consumers looking to make the leap from the centralized web 2.0 world to the TEA Project and Web3.

The TEA Project's Technical Details

TEA = Trusted Execution and Attestation

Any system with a high degree of anonymity and decentralization needs trust. The TEA Project leverages existing hardware security modules to achieve a secure computing environment that's both trustable and scalable. The TEA Project ecosystem achieves trust through the self-interested actions of the miners on its network. TEA miners are economically rewarded for running nodes that, through remote attestation, ensures that any TEA mining nodes hosting apps on the network haven't been tampered with. These attestations query the information stored inside each mining computer's hardware security module, and the reports are stored directly on the blockchain.

Multiple nodes will conduct remote attestations on any specific node before deciding if it's trustable. Once a node has successfully passed attestation, then these mining nodes gain two significant trusted computing capabilities:

1. All computations carried out within its hardware-protected enclave are trustable. The TEA Project's premise relative to TPM-protected mining hardware is simple: if we can trust the hardware integrity of the node, then we can trust the computational result. These hardware enclaves also ensure privacy for both the developer's code and the user's data.
2. The trustability of nodes is kept as Proof of Trust (PoT) data on the layer-1 blockchain. Because these nodes are trustable, they can run on the TEA Project's layer-2 without needing BFT (because they're already trustable.)

Let's summarize the above two points as follows:

1. **Trusted Execution** comes from protected enclaves where app logic can run while being protected by hardware security modules. All TEA Project TApps run inside of these protected enclaves. Nobody (which includes the app developer and the miners) has any control of the apps nor can they extract any data from the running enclaves.
2. **Trusted Attestation** refers to the process whereby the network nodes run reports on each other to ensure mining hardware running on the TEA network hasn't been tampered with. Once trusted status is attained, these nodes can run on TEA's layer-2 at cloud computing speeds without having to worry about (slow) BFT-consensus.

Transaction Consensus Using Time

We say that the TEA Project has three roots of trust: hardware, blockchain, and Proof of Time. We've already seen how hardware security modules and blockchain can create a trustable computing environment through remote attesation and trusted enclaves. This provides a scalable and trustable application execution tier for TApps, but we also need a data tier to track TApp state changes and transaction accounting.

The TEA Project has two state categories:

1. A strong-consistency state machine based on Proof of Time.

The first state category based on Proof of Time is for transactions requiring strong consistency, which would govern transactions involving funds and accounting. Using the time stamps from navigation satellites under watch of hardware attestation, our strong consistency state machine can achieve continuous state updates at a small synchronization cost. Its most crucial task is ensuring that the transaction sequence is consistent across all replicas. We do this by relying on the accurate time provided by the atomic clock of GPS satellites.

The reported time is recognized under the supervision of trusted TPM chips and used as the basis for the final ranking of all replicas. It's not necessary for all nodes to periodically reach a consensus on the latest block. But to ensure that most replicas can be synchronized to a consistent state, the TEA Project’s state machine requires a short waiting queue due to network latency. Since time is stable in our universe, it follows that each replica can achieve strong consistency using non-BFT consensus as the possibility of Byzantine faults has already been handled by the layer-1 blockchain.

2. An eventually consistent CRDT database built on OrbitDB that can be used ad hoc by TApps.

The other state category is a CRDT database that allows for short-term inconsistencies in the business logic of TApps. The TEA Project uses OrbitDB databases built on top of IPFS for these transactions. CRDT stands for conflict-free replication data type, which allows conflict-free mergers between different replications before ultimately achieving network-wide consistency. In fact, the business logic of most apps can tolerate short-term inconsistencies to achieve both decentralization and efficiency. A typical example of a traditional cloud app that deals with this issue gracefully is Google Docs.