Table of Contents

1. Overview

Why does the TEA Project exist?

- **1.1 Our vision to achieve fast, decentralized Web3 dApps.**
- **1.2 How TEA Project fulfills this role.** Multiple layers + trusted computing.

2. Different layers of the TEA Project

This section explains the technical framework of the TEA Project.

- **2.1 Layer-1** (built on Polkadot Substrate, WASM, IPFS etc.)
- **2.2 Layer-2 computing network** (built on top of layer-1's trust).

3. Trusted Computing

- **3.1 Proof of Trust/TPM**
- **3.2 Proof of Time/GPS**
- **3.3 Layer 2 consensus**

4. Token Economy

- **4.1 Dual tokens**
- **4.2 CML's role in ecosystem**
- **4.3 TEA's role in ecosystem**

5. Ecosystem Development

- 5.1 **The bonding curve**. Its impact on developers and miners and its value for fundraising.
- 5.2 **The 3-phase rollout.** How we onboard miners, developers, and finally consumers.

6. Conclusion

1. Overview

1.1 How the TEA Project Solves an Essential Need of Web3

The TEA Project provides a platform where rich, decentralized applications can run at native speeds across a decentralized network of computing nodes. It solves the Blockchain Trilemma by offering a scalable, decentralized, and secure blockchain without compromising any aspect.

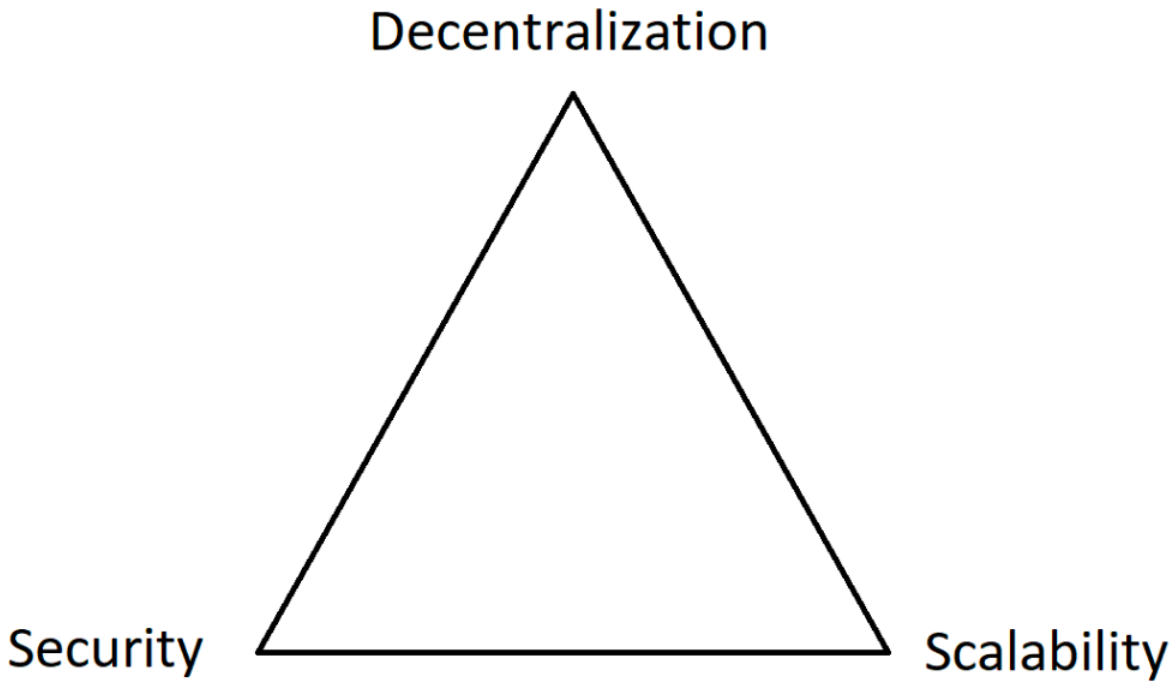

The TEA Project builds on many emerging projects and paradigms which in and of themselves are not sufficient to solve the Blockchain Trilemma. Blockchains like Ethereum provide a world computer where smart contracts can run in a decentralized manner. They are further augmented by Web3 innovations like IPFS, a peer-2-peer file system that stores files decentralized. But smart contracts by themselves cannot currently run

complex algorithms. Attempts to do so have shown smart contracts to be too slow or too expensive as they lack the processing power compared to modern cloud computers. A layer-2 solution would be needed to offload the computation tasks as long as it could provide a similar trust level as the layer-1 blockchain. By combining blockchain with IPFS and trusted hardware (TPM and GPS modules), the TEA Project achieves scalable speed without compromising security or decentralization.

TEA's Biggest Innovation is Trust

In the TEA Project, computing nodes become trusted through a remote attestation process. These trusted nodes can now store and transmit sensitive information and securely run WebAssembly code. The concept of trusted computing has so far been stuck in the limiting paradigm of one-computer = one-metal-box. This old thinking limits its expansion and possible uses cases in a distributed world. It needs the help of blockchain and other emerging technologies to expand beyond its current barriers. In particular, it needs a solution like the TEA Project to make a decentralized network of mining nodes trustable without sacrificing app execution speed.

Byzantine-fault Tolerance Leads to Slow Processing

No matter how many nodes are running in a typical blockchain system, the one-by-one ordered transaction queue of existing Byzantine fault tolerant (BFT) consensus algorithms leads to slow data processing times. Smart contracts currently cannot run complex algorithms. Attempts to do so have shown smart contracts to be too slow or too expensive as they lack the processing power of modern cloud computers. Existing blockchains are slow because they need BFT to reach consensus. They could host dApps that reach the speed of cloud computers if only they had trust. Trust is exactly what the TEA Project is able to provide through its 2-layer blockchain design. We don't yet have rich dApps running at cloud-speed on blockchains since these systems need consensus to guard against a trust-less environment. Cloud computing gives you rich and fast apps, but it's not decentralized and you must assume their ecosystem is trustable.

2.1 How the TEA Project Delivers Decentralized, Secure, and Scalable Apps

In the TEA Project, trust is already taken care of by our layer-1 blockchain. That allows our layer-2 nodes to all trust each other by default and no longer need to watch for Byzantine faults. The TEA Project’s decentralized apps can run as fast as cloud computers who are able to take their trustable (centralized) nodes for granted.

Decentralization

The TEA Project’s decentralized TApps run on decentralized miners that connect to each other through a peer-to-peer network. Every TEA node is protected and monitored by hardware TPM chips or the CPU's TEE (Trusted Execution Environment) which provides PoT (Proof of Trust) data.

Every TEA node can reference this PoT data for every node participant to verify any individual TEA node's integrity through a remote attestation process completed through smart contracts. Nodes in the TEA network are incentivized by the token economy to perform remote attestation verification.

Security

The hardware security modules allow the compute nodes to offer a hardware-protected enclave where code and data can securely run. Due to limited resources inside the protected enclave, developers must decide which functions need to run inside the enclave for security and which functions should run outside the enclave to gain performance.

Anyone outside the TEA module's hardware has no knowledge about what happens inside the module. Neither the TEA node's owner nor the server connected to the module will know, nor can anyone predict which node is running any particular task. This gives the TEA network security as no information can be breached during this secure workflow.

The TEA Project's goal is to build a platform that allows code and data to run inside trusted TEA modules without needing to trust anything else. The

technologies built into the TEA platform protects the data and the integrity of the result.

The decentralized app's computation is based on the expected code and the expected data. No one can alter the input code, the input data, or the output result. The result of the computation is correct given the computing environment is verified through Proof of Trust data that is shared in the network and verified via remote attestation. There's no chance of data breaches during the entire process.

Scalability

In the TEA Project, trust solves the scalability issue by opening up a second blockchain where every node is trusted.

Typically, only centralized cloud computing can take consensus for granted as a single company is able to oversee all of the computing nodes. That allows cloud computing companies to be scalable and secure (as long as you trust the company). But we can't totally trust our data with these companies as they are centralized and are typically not transparent with their data safeguards.

The TEA Project solves scalability by using non-BFT consensus to run TApps on its fast layer-2.

- The TEA Project uses its layer-1 blockchain to run BFT to filter out all non-trustable nodes.
- TEA's layer-2 can check the trust status of any of its nodes by querying the layer-1 blockchain. The TEA Project's layer-2 can now use a much faster non-BFT consensus because it can assume all of its nodes are trustable as reported by the layer-1 blockchain.

Developers and Code

In the current cloud computing model, trust is needed between developers, cloud service providers, and the end client.

Let's assume that a consumer has valuable and sensitive data such as health data output from wearable tech. It needs an algorithm to make sense of the data, and computing power to run the code. There are algorithms (code) created by developers that can do this job very well. The consumer wants to run the code with their data as input. Additionally, you don't trust the developers or service providers who host the code due to concerns about possible data breaches.

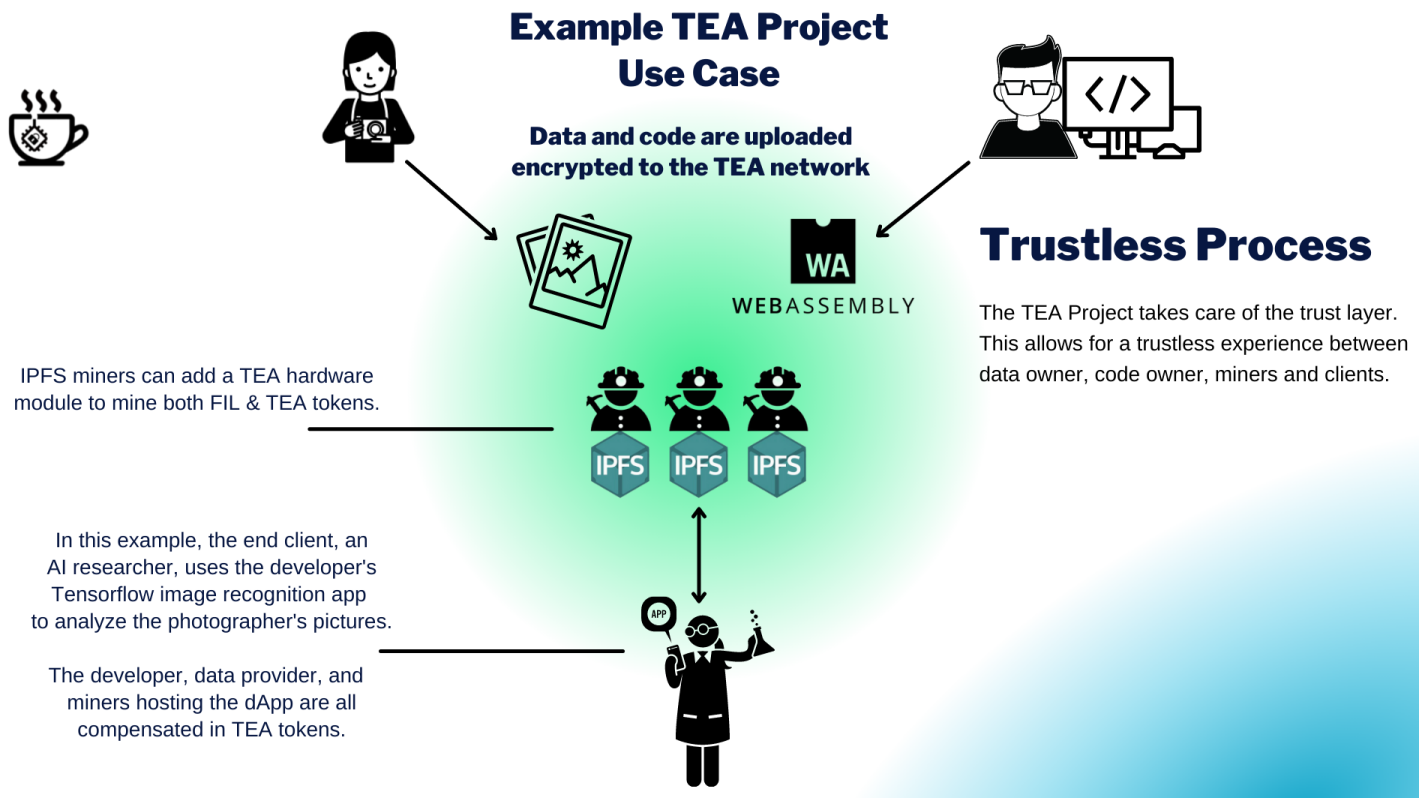

A tensorflow app like the one shown above can now be built in a way that respects the privacy of all participants.

TEA Provides Developers an Onramp to Web3

The TEA Project provides a familiar framework for developers to build decentralized apps on the emerging Web3 ecosystem. A big reason that Web3 has not yet become a reality yet is the lack of framework. To help developers bring their apps to Web3, TEA provides a familiar three-tier architecture that replicates what developers are used to developing for on cloud computing:

- Front-end (e.g. written in html, css, and javascript) that's hosted on IPFS.
- An execution layer that replicates the computer power of the server. This execution tier is decentralized among the mining nodes running on our layer-2.
- A database tier that allows developers to query the current state of their app. This functionality of TEA is also run on our layer-2 mining nodes and is made possible by their on-board GPS module that orders state changes using time as the source of trust.

The TEA Project will have an SDK available to help shorten development time on our platform. Creating dApps for the TEA Project will always be straightforward for developers as our native binary format is WebAssembly (WASM). WASM, the native language of TApps, is fast, secure, and easy to output to from a variety of different programming languages. This flexibility makes it easier for developers to take the leap into Web3 by using the TEA Project as their entry point.

Further Reading on Problems the TEA Project Solves

- Fixing a Broken Internet
- How TEA Fixes the Internet

Our TEA Project Whiteboard Explainer Videos are also helpful to understand the TEA Project ecosystem and how the various parts work together.

2. The TEA Project's Two Layer Setup for Faster Dapps

TEA stands for **T**rusted **E**xecution and **A**ttestation. These aspects of the TEA Project are achieved across two different layers: a layer-1 that uses blockchain and a layer-2 that uses hardware and time as roots of trust. The layer-1 has consensus as it uses blockchain, while the layer-2 does not. The TEA Project uses consensus not to confirm the result of the dApp's code; it instead uses consensus to verify the execution environment where

the code was run. The TEA Project's philosophy is that **if the execution environment, code, and input data are trusted, then the execution result can also be trusted.**

The TEA Project aims to build a platform that allows code and data to run inside TEA modules with complete trust. The technologies built into the TEA platform protects the data and makes sure:

- The computation is based on the expected code and expected data. No one can alter the input code, the input data, or the output result.
- The computation's result is correct given the computing environment is verified and there's no data breach during the whole process.

The TEA Project has made completely trustable nodes on its layer-2 possible through the following factors:

- TEA nodes connect to each other through a peer-to-peer network.
- Every TEA node is protected and monitored by a hardware TPM chip that provides PoT (Proof of Trust) data.
- Every TEA node can verify any other TEA node's integrity through a remote attestation process. The verification is controlled and determined by blockchain smart contracts.
- Anyone outside the TEA module's hardware has no access to what's inside the module. In addition, neither the TEA node's owner nor the TApp developer can predict which node is running or will run any particular task.
- After the task is done, anyone can verify the Proof of Trust data of how their code or data was processed.
- The layer-1 blockchain keeps the trust data on the layer-2 mining nodes. Its main job is to ensure that all layer-2 nodes are trustable.
- Because all layer-2 nodes are trustable (as determined by our layer-1 consensus), we no longer have to worry about Byzantine faults on layer-2. There's no consensus needed on layer-2, and the roots of trust on layer-2 are concerned with providing functionality to the

dApps that run on this layer. The hardware root of trust (TPM) ensures a trusted execution environment (enclave) where code and data can run decrypted without anyone being able to peak inside. The second root of trust for our layer-2, time, ensures that transactions are always properly ordered in our state machine.

The TEA Project's Layer-1

TEA's layer-1 blockchain is responsible for the security and overall trust of the TEA Project. Layer-1 keeps trust certificates on all the nodes running on layer-2. Each layer-2 node must be in good standing based on the information reported by the node's TPM chip under the watch of remote attestation.

TEA uses its layer-1 blockchain to provide:

- Economic incentives and penalties that form the basis of its token economy.
- Immutable trust information storage, such as credit history (incentives and penalties for nodes governed by smart contracts) as well as Proof of Trust information (key IDs, and hashes).
- Block verification and to maintain the block height as a universal clock between TEA nodes.

The immutable trust information stored in TEA's layer-1 blockchain is used in the remote attestation process. This process is used in conjunction with the TPM chips onboard the mining computers to ensure each node is trustable:

1. Multiple nodes are randomly called to conduct remote attestation on any particular node under inspection.
2. The results of these remote attestations are stored in the layer-1 blockchain as Proof of Trust (PoT) data.

Successful remote attestation on a candidate node ensures that the machine's TPM chip hasn't been tampered with. Therefore, the secure enclave overseen by the TPM chip is trustable and will protect the confidentiality of both the code and data processed within the enclave. This

is the fundamental premise of the TEA Project's concept of **trust**. If a node shows full integrity under remote attestations performed by multiple peer nodes, then its enclave is secure, and the resulting calculations (i.e. app output) can be trusted.

How TEA's Layer-1 Conducts Remote Attestation

In the TEA network, each trusted computing node has a Root of Trust (RoT): a key pair generated by the TPM manufacturer. The private key for this RoT is held inside the hardware and not known to anyone, including the manufacturer. The public key is disclosed and recorded on the blockchain. Periodically, the blockchain will run remote attestation by requesting a computational node to generate its hardware fingerprint and fingerprints for all the software modules running in the node. These will need to be signed by the TPM's private key before being returned as a response to the attestation request. Each attesting node would verify the signature and the Proof of Trust (PoT) data returned.

TEA's layer-1 consensus protocol randomly selects remote attestation nodes on layer-2 to challenge the randomly selected node (also on layer-2) being tested. If a two-thirds (2/3) majority of remote attestation nodes receive a positive result from verifying the target node's requested attestation material, the tested node will have passed the consensus verification and be marked as trustable to use. Note that trustable nodes are the only ones selectable to perform the remote attestation process. Therefore, as long as the network can keep a minimum ratio of 2/3 trustable nodes, we can trust the consensus result. To maintain this minimum threshold, the TEA Project team will initially bootstrap the network with clean nodes and control the birth rate of new nodes (i.e. control the issuance of Camellia NFTs needed to mine on the TEA network).

The results from all attesting nodes will run through the blockchain's consensus protocol before a final decision is made on whether the computing node still meets all the security requirements for it to remain on the network. Finally, the result of the attestation process is recorded on the layer-1 blockchain.

The different layers of software running in the computational node generate their own secrets, one after another, each using the parent layer's key pairs so that all of them can pass cryptographic verification. These nodes together form a trust chain, with the root of the chain being the hardware RoT. Because the public key of the RoT is stored in the blockchain publicly, anyone can easily verify any derived secrets by tracing back the chain of trust and confirming the RoT using data recorded on the blockchain.

The Technical Implementation of the TEA Project's Layer-1 Blockchain

We have developed our layer-1 blockchain using Substrate as the provider. But TEA can work with any layer-1 blockchain that has smart contract support will work with TEA, such as Ethereum. It can also run as a parachain in the Polkadot / Kusama ecosystem. As a parachain, the TEA Project would become a middleware layer for other projects needing to farm out secure computing tasks that would otherwise require centralized cloud computing.

- Chains such as Ethereum that are not based on Substrate will need some kind of smart contract (as in a bridge) to link messages between the chains at low-bandwidth and transfer funds (at the expense of gas fees from the connected chain).
- For Substate-based chains, the TEA Project's layer-2 can communicate with the Substrate-based layer-1 chain using XCMP. XCMP, or **Cross-Chain Message Passing**, is based on XCM, or **Cross-Consensus Message Format**, and is a method for Substrate-based chains to communicate with each other.

We chose Substrate because it's modern, fast, release-ready, and written in Rust and WASM (the same tech that the TEA Project is built on). Like Polkadot, the TEA Project uses the **GRANDPA** consensus algorithm for block finality. GRANDPA achieves consensus using asynchronous Byzantine fault tolerance (aBFT), allowing messages to be received and re-broadcast across all nodes before they are committed. This allows nodes to

recover from potential message loss and is a robust BFT consensus algorithm for asynchronous networks.

The Delegation Chain Builds the Trust Chain Between TEA's Layer-1 & Layer-2

The delegation chain is a network protocol between each TEA node's HSM to transfer trusted computing tasks securely. It uses multiple parties to conduct remote attestation to ensure that data is only transferred among verified trusted hardware modules.

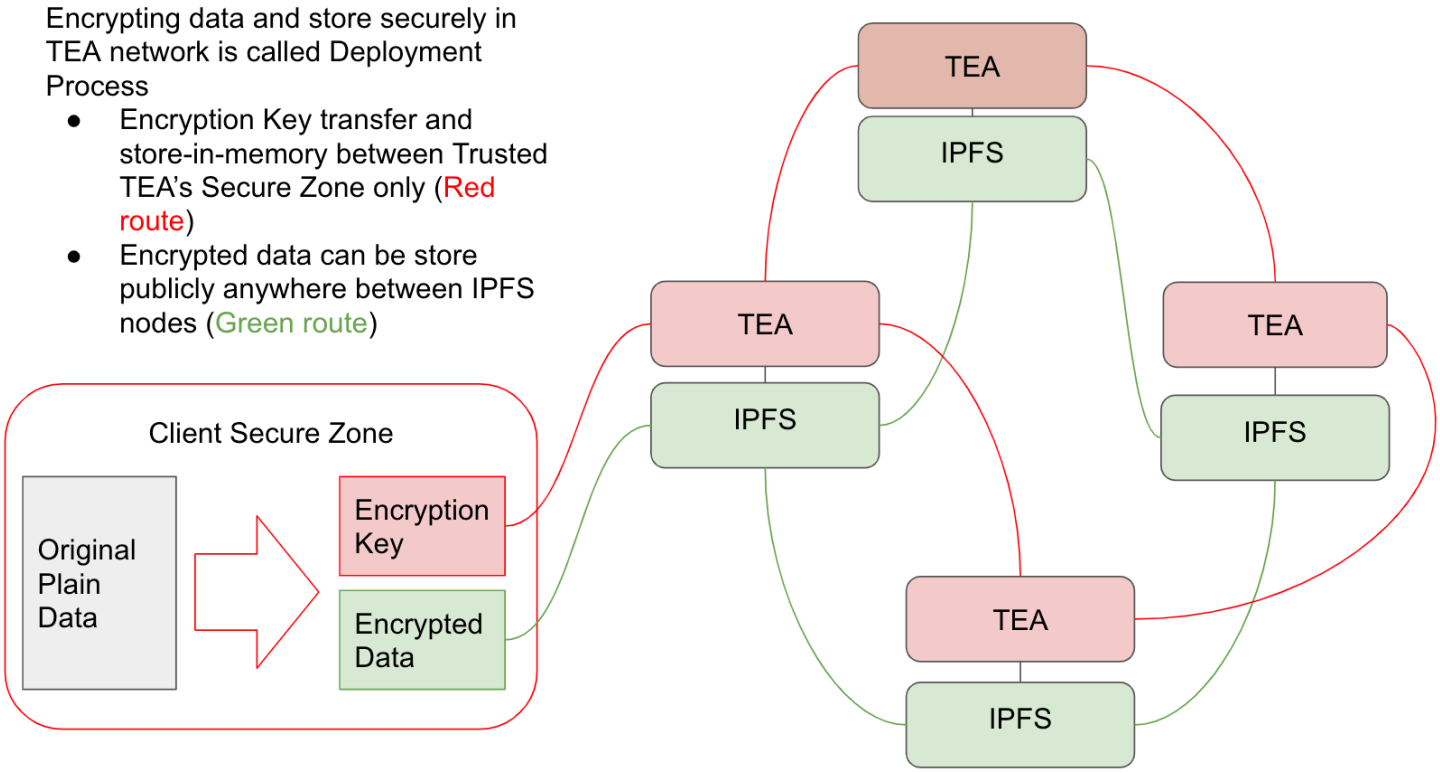

- A client sends secure data or code to a trusted TEA node as a delegator. If the client doesn't trust any other nodes, they would be best served to own a TEA node to act as a delegator for their own data or code.
- A delegator will be looking for qualified executors among all the TEA nodes in the IPFS p2p network. Remote attestation is done between each node before exchanging any sensitive information.

Data or code will be transferred via a **repin** to a new delegator (called a **pinner**) to host or an executor to run.

- No matter where the data or code goes, the proof of delegation data will be attached at each step, forming a delegation chain.
- Anyone can verify the delegation chain to ensure the chain is valid. Any hacks in the middle would be easily discovered via the blockchain.
- TEA's layer-1 blockchain will be used to do the verification. Any incentives or punishments is then applied to the participating nodes.

The trust chain starts from the hardware-based trust data (from the hardware security modules like TPM) that's stored in the layer-1 blockchain. The trust chain connects the Proof of Trust data up to the layer-2 nodes running the WASM code.

The TEA Project's Layer-2

TEA's layer-1 is rock-solid for both security and decentralization, but the Byzantine fault tolerance it must perform for its consensus means it's slow. It's an example of why you generally don't want to do computing inside the blockchain: complicated algorithms cannot run in smart contracts due to computational cost and time needed for completion. It's a common misconception to assume that Web3 will always require a layer to run smart contracts. The TEA Project is looking to run full-speed Web3 dApps, something that can't be done only using blockchain and smart contracts. That's why the TEA Project has evolved a layer-2 to run dApps at full speed while relying on TEA's layer-1 for trust.

The TEA Project's dApps are known as TApps. In each TApp running on the network, computation tasks are carried out by layer-2 components using a hardware-based Root of Trust to achieve efficient, trusted computing.

- Layer-2 nodes (CML nodes) only trust other CML with certificates issued by layer-1. Because the mining nodes have gained trust from the layer-1 blockchain, they can now reach cloud computing performance as layer-2 no longer needs to guard against Byzantine faults.

- Applications run on layer-2 nodes without any knowledge about blockchain and consensus, as if they were running on cloud computing. These apps are seamlessly decentralized by the TEA Project's 2-layer blockchain design.

The importance of the TEA Project's layer-1 is that it keeps track of the trust status of all nodes on TEA's layer-2. 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 while providing them a safe space to run at full speed. Once a node is certified trustable, it can run encrypted app code in its TPM-protected hardware enclave that interacts with encrypted customer data. The system is designed such that neither the miners nor the developers have access to the customer's data or app code once it enters the mining machine's protected enclave.

TApps Run Full Speed on a Familiar Architecture

Applications can run full speed on the TEA Project's layer-2 because it looks just like a normal three-tier cloud computing architecture from the app's point of view. 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 only 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.

The Technical Implementation of the TEA Project's Layer-2 Blockchain

The TPM module of the hosted mining nodes forms an integral part of the layer-2 hardware tech stack. The TPM modules allow a Root of Trust to be established between nodes running on layer-2 and the underlying layer-1 blockchain that keeps the nodes' Proof of Trust data. The TPM chip also provides the protected enclaves in the layer-2 mining nodes where

encrypted code and data are executed by the decentralized TApps of the TEA Project network.

TEA's data storage uses IPFS/libp2p to form a peer-to-peer network. Because IPFS is publicly accessible, storing data on IPFS means everyone can access it as long as the content identifier (CID) is known. In order to protect the data, everything TEA stores in IPFS needs to be encrypted. The encryption key will never be stored in IPFS or any persistent media. It resides only in the memory of the TEA modules.

TEA's runtime environment is WebAssembly (WASM) that is specially configured to run inside a secure enclave. WASM is the only code allowed to run in the TEA runtime, as WASM bytecode is considered one of the most secure form of executables.

Only minimized operating system is run (we currently use a modified version of a minimized NixOS). One important change from the stock distribution is that vSockets have replaced the TCP/IP networking stack. This means that the secure enclave cannot initiate or accept inbound/outbound internet traffic when processing a client's task.

Economics-wise, using a blockchain to record trust information of computation nodes and using VRF to randomly select which node will lead the execution of a client's computation job ensures that it would be very costly to attack the TEA network. We use token incentives to encourage many participants to run a node themselves or stake their tokens with trusted nodes. The token staking mechanism uses economics to reinforce well-behaving nodes.

Trust Allows WebAssembly Code to Run Full Speed

As the TEA Project's layer-1 verifies the trustable status of each node, developers can run their decentralized apps at native speeds on layer-2. At the application layer, it's only strictly required that developers use WebAssembly (WASM) to deploy code to the TEA runtime. The security built into the WebAssembly language itself also helps to secure the network. Using WASM is also conducive for developing applications in the TEA ecosystem. Most modern programming languages can be compiled into the

WebAssembly code format. Developers can stick with their preferred programming languages to write their apps which speeds up development on the TEA platform.

TEA Project's layer-2 currently uses the file storage and networking services provided by IPFS. Any current IPFS miner can begin participating in the TEA network by adding a TEA hardware security module (such as a Raspberry Pi with a TPM chip).

Further Technical Reading on TEA's Layer-1 Blockchain

- For more information on the TEA Project's proof-of-stake consensus algorithm that's built on Substrate (the consensus algorithm that Polkadot also uses): https://medium.com/polkadot-network/polkadot-consensus-part-1-introduction-3e3cd6237243
- For more technical details on how TEA's layer-2 can communicate with another Substate-based chain functioning as its layer-1: https://wiki.polkadot.network/docs/learn-crosschain

Further Technical Reading on TEA's Layer-2 Blockchain

- For more information on how code and data travels along IPFS while remaining secure in the TEA Project: https://teaproject.medium.com/turn-ipfs-file-system-into-ipfs-function-as-a-service-25dbcaff3975
- For an overview of how mining nodes interact and share code and data: https://teaproject.medium.com/tea-project-pinning-and-re-pinning-in-practice-9886e147e5ca
- To learn more about TEA's 3-tier decentralized architecture: https://teaproject.medium.com/the-tapps-3-tier-decentralized-tech-stack-43d2872f609b

3. Trusted Computing

3.1 Proof of Trust / TPM

The primary objective of the TEA Project's layer-1 is to maintain the trusted status of all nodes on layer-2. These nodes that gain trusted status through attestation on layer-1 no longer have to ensure Byzantine Fault Tolerance (BFT).

The apps can run full speed on layer-2 once trust among nodes can be taken for granted. This is in contrast to traditional blockchain apps which, because they run on the blockchain directly, need to perform BFT consensus. This is a roadblock for traditional blockchain apps looking to perform at scale while maintaining decentralization and security. Conventional blockchain is hamstrung by the two prevailing roots of trust, cryptography and consensus algorithms. The TEA Project creates a third possibility, a hardware root of trust, through cheap security modules like the TPM chip.

Sources of Truth: Blockchain, Hardware and Time

The TEA project relies on three sources of trust: the blockchain, hardware, and GPS satellites.

Every node stores its essential data to the blockchain. Based on the blockchain's secure nature, the data stored in the blockchain is considered trustable. When performing remote attestation, the verifier won't trust anything the attester claims; everything needs to come from either the historical data in blockchain or the hardware signed message. All other decisions are made based on those two sources of truth.

The TEA Project can use trusted execution environments like Amazon's Nitro or a hardware security module (HSM) like TPM as a source of trust.

- For Amazon Nitro, the validation is centralized and requires no further consensus.
- For hardware security modules like TPM, there are still known vulnerabilities. Since TPM alone is not secure enough, a delegation chain is created on the network for a remote attestation workflow.

Making Secure Mining Hardware Affordable Improves Decentralization

As opposed to crypto projects that require miners to have expensive CPU, GPU, ASIC, hard disk drives, or other expensive computing equipment, miners operating on the TEA Project network can use single-board computers as long as they have a TPM security chip and a GPS module. A typical mining machine capable enough to run on the TEA network can be setup for less than $100. TEA Project miners won't need to engage in high-energy complex cryptographic computing to run consensus algorithms. Making TEA Project mining affordable and accessible to the general public helps improve decentralization by making mining within reach for more regions of the world.

Leveraging Hardware Security Modules for Trust

The TEA Project works with existing secure hardware technologies like the Trusted Platform Module (TPM) to add a trusted layer to its blockchain. Each node on the TEA network must have a hardware security module (HSM). To be able to run decentralized apps and execute code, these decentralized nodes must pass remote attestation when queried by existing trusted nodes on the network.

The remote attestation process queries the embedded HSM on the candidate node, and its embedded HSM (such as a TPM chip) will respond to the query. If the result is as expected, then the candidate node gains trusted status. The result of the remote attestation process is stored on the blockchain.

Once a node is trusted, it can carry out computations for TEA Project apps within its secure enclave. The TEA Project's major premise regarding trust is that if we can trust the the node through remote attestation, we can trust the integrity of its secure enclave to run code and process data in a way that's secure and zero-knowledge.

TEA Project Hardware Support

The TEA Project's roadmap for supporting various Root of Trust (RoT) verification chains depends on the underlying hardware.

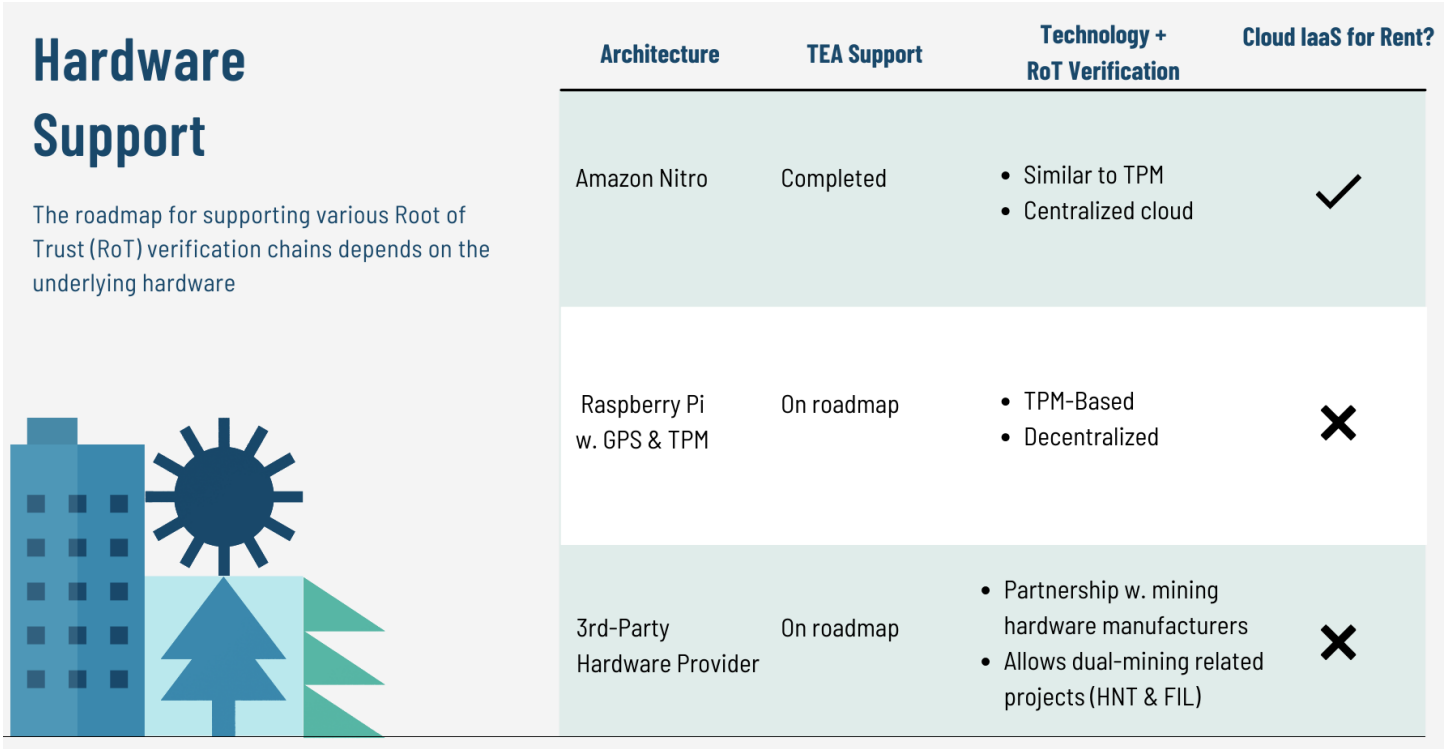

3.2 Proof of Time/GPS

Transaction Consensus Using Time

According to the needs of their specific business logic, TApp applications running on the TEA network have two state categories:

1. A strong-consistency state machine based on Proof of Time.
2. An eventually consistent CRDT database built on OrbitDB that can be used ad hoc by TApps.

The Strong Consistency State Machine

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. The TEA Project is continuous as it has no concept of blocks, and it’s able to sync nodes at very little cost as it has no PoW puzzles to solve and doesn't have to periodically stop and wait for all nodes to sync before continuing.

The complexity of the TEA Project's strong consistency state machine lies in synchronizing the current state between multiple nodes. Because we do not want to use the traditional blockchain consensus algorithms to achieve strong consistency, the 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.

The Eventually Consistent CRDT Database

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.

The TEA Project uses CRDT databases because they are cheap and fast.

- CRDT is fast as it has no time delay and doesn’t need to wait for others. Allowances are made for new transaction reports which are

added non-destructively. This is in contrast to the TEA Project's strong consistency state machine which must wait a minimum amount of time for confirmation.
- CRDT storage cost is cheap, relying on IPFS for decentralized hard drive storage instead of the more expensive RAM storage.

3.3 Layer-2 Consensus

How Layer-2 Maintains Consensus - CRDT Databases

The design of the TEA Project's layer-2 consensus ensures that even if transactions are missing (e.g. the transaction is delayed by network congestion), the ordering of transactions will still be correct. The TEA Project is able to achieve a reliable ordering of transactions through its use of replicas and CRDT databases.

- Delegators send **transaction hashes** and **transaction receipts** to multiple replicas. These followup receipts have timestamps that show when the original tx hash was sent out.
- Replicas match the hashes of transactions with their followup receipts. Replicas can re-order transactions as long as they remain within a period of time relative to the now time of the replica's clock.
- A decentralized, peer-to-peer sync function between replicas keeps a CRDT database of ordered transactions among all replicas. Because it's **conflict-free**, the sync function can only add a missing transaction hash to a replica when it's missing on their ledger compared to other replicas on the network.

Because it takes time for a delegator to package up a transaction and create its hash, this transaction doesn't have a timestamp of when it was sent out.

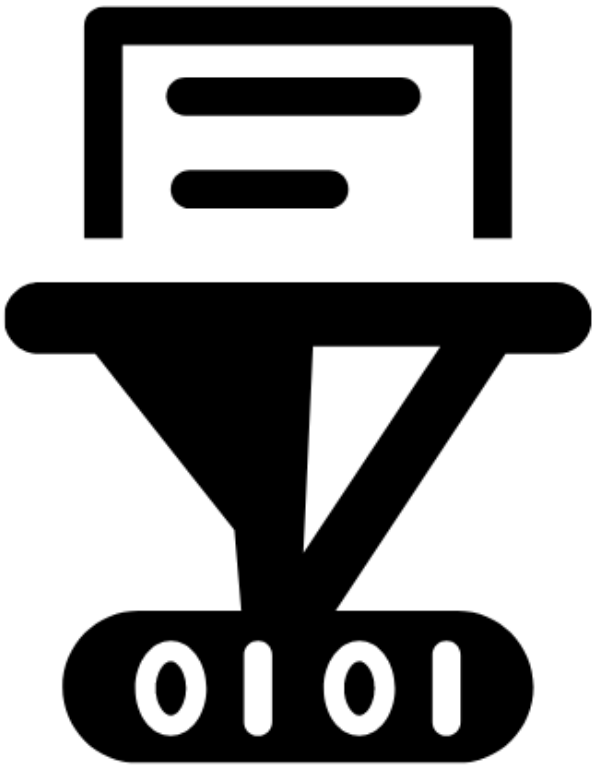

Transaction hash has no timestamp when sent out

You can imagine the delegator sealing the envelope of the transaction before sending it out to other replicas on the network. Since the envelope is sealed, there's no way to include the sent timestamp within the envelope.

How Layer-2 Maintains Consensus - Strong Consistency State Machine

The above is an explanation of how consensus occurs for our NoSQL CRDT database that runs through our B CML mining nodes. In contrast, our strong

consistency state machine that runs through A CML mining nodes doesn't need any consensus as there are no new blocks to wait for. Instead of blocks, new transactions land on a conveyor belt and eventually everyone ends up with the same state. This state machine functions just like a database, and its technological design gives the TEA Project a potentially limitless TPS.

The TEA Project's state machine only needs to store these state change in RAM and not the transaction itself. Since the on-board TPM chips of the mining nodes allows attestation to be run on them, they can be guaranteed to be trustable and don't need to sync up a historic ledger of all previous transactions. Once all nodes are trustable, any node can get the latest state from a nearby node's RAM. Syncing up to the latest state through reading a nearby node's RAM is a quick process, much quicker than having to reconstruct the entire ledger to get the latest block.

In the TEA Project:

- Transactions are processed with the resulting state change stored in the RAM memory residing in the enclave of A CML mining nodes.
- Only A nodes run the strong consistency state machine (including the sql database instance).
- Every A node will have the same copy of the state in their memory.
- There's also a persistent backup kept to IPFS but in an encrypted format.

Two State Machines = Better for Developers

Developers will have a choice when developing their decentralized apps on the TEA Project:

- The NoSQL CRDT database that runs through B CML and uses OrbitDB for storage on IPFS.
- The GlueSQL database that uses the on-board GPS modules of A CML to update the current state continuously. IPFS is only used to write backups of the current database state. The GlueSQL database can be used by developers the same as they use SQL databases in the cloud computing world.

The GlueSQL state machine is more expensive and appropriate for more complex app interactions. In an effort to keep costs down, application developers can allocate more complex database queries to A CML nodes while using the CRDT database running through B CML nodes for less important tasks.

References and Further Reading

- An interactive TPM simulator: https://google.github.io/tpm-js/
- The official website of the Trusted Computing Group who's responsible for maintaining the open standards of TPM chips: https://trustedcomputinggroup.org/
- Information on Google Spanner which also uses time as a root of trust for ordering transactions: https://en.wikipedia.org/wiki/Spanner_(database)
- A research article detailing how GPS can be used with TPM to create trustable time synchronization among nodes in a network: https://www.mdpi.com/2076-3417/11/18/8288/htm

4. Token Economy

4.1 TEA Project's Dual Token Design

The two tokens in the TEA ecosystem are the TEA utility token and the Camellia (CML) NFT. The two tokens play different roles towards incentivizing participation among TEA ecosystem participants. The TEA token is the token that users must pay when using the network's TApps. TEA is also used to incentivize miners who earn TEA revenue in exchange for the service their nodes provide the TEA network. CML are issued by the DAO according to demand, with the TEA that's exchanged for CML being burned by the DAO.

We envision demand for the two tokens to come from multiple factors:

- **CML** is required for mining on the network so miners will buy it to make their mining nodes active. There are three different types of

Camellia: A, B, and C CML. A CML maintain TEA's state machine, B CML host TApps, and C CML can only host TApps for private data.

- **TEA** is exchanged for CML for those who wish to earn mining revenue. But its primary demand will stem from those needing it to use TApps. This creates demand for TEA as users will need to bridge assets over from different blockchains in order to exchange for TEA. Every decentralized app running on the TEA network, called TApps, will each have their own TApp token as an investment and fundraising vehicle for the TApp. Users will need to use TEA to purchase these TApp tokens, another source of demand for TEA.

4.2 Camellia NFT's Role in the TEA Ecosystem

Camellia (CML) is a non-fungible token (NFT) that's used by TEA miners to activate their mining machines. The CML tokens essentially act as both an ID mechanism for their mining nodes as well as an access key that allows them to mine on the TEA network. The CML that is associated with a mining machine tracks that machine's trustworthiness (e.g. uptime). This information on each CML ID is stored on the TEA Project's layer-1 blockchain. In a scenario where the miner wishes to upgrade their hardware, moving the CML to their new machine also transfers its trust score.

Miners must pay 1000T as a deposit for mining. This amount is automatically withdrawn from the miner's wallet when they plant a CML. If a miner doesn't have enough TEA tokens for the 1000T initial staking slot, then they may take on debt. TEA token loans will only be for the beginning epochs and are designed to boost the supply of TEA in the early stages of the network.

Camellia start off as seeds and have different productive potentials and lifespans. Every CML has a life span that mimics organic plants - it grows from a seed, matures into a fully productive tree, and eventually dies. A mining CML mimics this same life cycle through its productive phases: it yields relatively less TEA near the beginning and end of its life and produces the most TEA at its middle-age peak. This natural life cycle CML helps keep the miner network decentralized as even early-adopters must compete with newcomers for new seeds when theirs die off.

The DAO is in charge of generating new CML tokens. The supply (birth rate of CMLs) is controlled by the DAO based on the level of demand for new CMLs. These new CML seeds are auctioned off to prospective miners; the TEA received in return for the CML seeds are burned by the DAO based on prevailing supply and demand. These selective burns help support the TEA token economy by limiting the supply of TEA.

4.3 The TEA Token's Role in the TEA Ecosystem

TEA is a utility token used by the network to pay miners and by clients to run decentralized TApps. As a utility token, it's also used to pay gas on the network. The unit cost of computing resources such as cpu utilization, ram, network traffic, and storage should be relatively stable when measured in terms of TEA tokens:

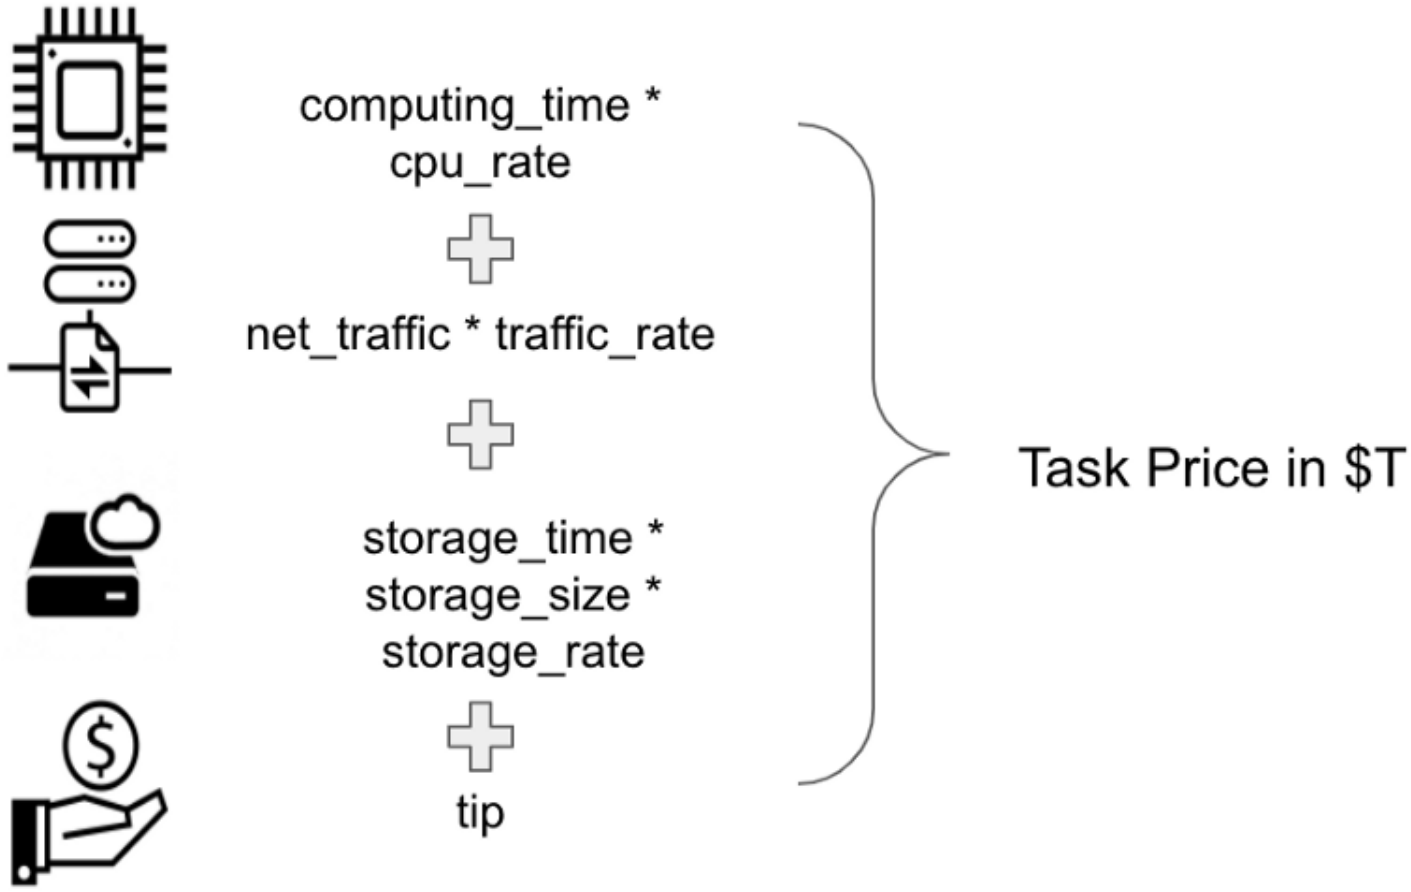

The TEA Project will generate initial tokens in its genesis block to distribute to investors and team members. The pre-mined amount is 100 million TEA

tokens. From that point on, TEA will be generated solely from Camellia mining.

Camellia mining gives miners a chance to earn block rewards from performing routine tasks such as running remote attestation of computing nodes, proposing new blocks, and running smart contracts. Decentralized applications (known as **TApps**) will need miners with computing nodes to host them (i.e. machines with a CML token planted). In return for hosting TApps, miners will either earn a share of that TApp's token or be paid out in TEA tokens. The TEA tokens that are rewarded to miners for hosting a TApp that's elected to payout in TEA will come from consume actions (i.e. usage of the app) or be paid out directly by cashing out the TApp's token for TEA to be paid by the miners.

The TEA token has many uses on the TEA Project 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 mining rewards. Staking is a way for investors to lock up their TEA tokens and earn revenue. Giving users the option to stake their TEA locks it up and helps support the token price.
- TEA tokens can be staked by users to vote on DAO governance issues.

Factors that Influence the Mining of TEA

Any node in the TEA network doesn't know what computing tasks will be assigned to their node. But there are many factors that a miner can take into account to help their nodes be chosen for high TEA revenue tasks. The TEA Project's Proof of Trust consensus will probabilistically favor nodes that:

- Are higher productivity machines, i.e. more powerful hardware.
- A higher productivity type of plant, i.e. their CML is at middle-age.
- Higher total stake value from investors staking into the machine.
- Some special features that others machines do not have, e.g. an onboard TPU.

CML Mining Machines and Its Bonding Curve Tokens

Miners are able to earn TEA tokens by running public service jobs (like remote attestation) and executing customer compute tasks as assigned by the TEA network. This income is shared with investors who decide to invest funds into the mining machine through its bonding curve token.

Each Camellia has a bonding curve token associated with it that investors can purchase. The stakers therefore have a claim on a percentage that the mining machine earns as mining revenue. This allows non-technical users (or uses who don't have the space or bandwidth to run a mining machine) to earn TEA revenue by participating in mining without having to setup and administer a mining machine.

The timing of when these bonding curve tokens are purchased matter when it comes time to distribute the mining revenue. The earlier the purchase time, the cheaper the price of the bonding curve tokens. And the earlier the purchase time, the more chances the token holder will have to earn dividend payouts according to the amount of tokens they hold. Dividends are paid out in the Camellia's bonding curve token. Because it's a bonding curve, token holders can cash out to TEA at any time at the price determined by the bonding curve.

The Supply of CML and TEA

CML Supply

CML supply is algorithmically determined based on demand for miners to host TApps on the network. Code monitors how busy / idle the miners are on the network. New CML are issued and put up for auction if miner utilization passes a set threshold in order to put some slack back into the system. TEA supply is reduced when someone buys CML through an open bidding process. After the miner or investor purchases CML, the TEA used in the purchase is burned by the DAO.

TEA Supply

The genesis block contains 100 million TEA tokens to be distributed as follows:

- Seed round investors: 10% (10 million TEA).
- A round investors: 10% (10 million TEA).
- B round investors: 10% (10 million TEA).
- Team tokens: 30% (30 million TEA).
- Operations and marketing: 8% (8 million TEA).
- Community: 6% (6 million TEA).
- Liquidity pool rewards: 6% (6 million TEA).
- Potential IDO round (i.e. parachain auction staker reward): 20% (20 million TEA).

Block Rewards

- TEA token reward reduction rate for every half-year: 0.7 (30% reduction in block reward rate after every 6 months during the first 2 years).
- How many blocks in half a year: 2,635,200
- Initial public service reward per block at Genesis Block 9.06944
- Initial mining rewards per block: 7T (A CML) + 2T (B CML) + .06944 (layer-1 staking)
- First reward reduction at 1/2 year: 6.348608
- 2nd reward reduction at 1 year: 4.4440256
- 3rd reward reduction at 1- 1/2 year: 3.11081792
- Last reward reduction at 2 year: 2.177572544

Links for More Information

- How the TEA Ecosystem Supports the TEA Token Economy
- The TEA Token Model: How Does the Business Support Value / Price?
- Camellia, an NFT With a Life. Section within longer article about the TEA token economy.

5. Ecosystem-Development

5.1 TApp Bonded Token Sales

The TEA Project offers decentralized app developers the ability to raise funding through the sale of **TApp tokens**. Each TApp can have an ERC20-compatible fungible token associated with it in order to raise funds and reward those who wish to invest in the project.

Whoever holds a TApp's tokens owns some percentage of the project. These tokens represent tokenized ownership of the project and are a claim to revenue dividends distributed to the project's stakeholders. Making a token available for a TApp allows its users to become investors in it.

TApp Tokens: Bonding Curves

TApp tokens follow a bonding curve where price programmatically increases as supply increases. A TApp's tokens can be traded into its bonding curve (buy or sell) at any time for TEA tokens. These tokens can also be traded outside of the Tea Project as a standard ERC20 token.

Every bonding curve has a **Theta** value, which is the percent of every token buy or consume event that's used to fund the development of the TApp. If **theta** is set at .2 (20%), then 1 - .20 = .8 is used to fund the sell side of the bonding curve (80%). The smaller the number for **theta**, the closer the buy and sell prices are on the bonding curve.

An important innovation, the bonding curve design makes sure that there's a reserve fund backing the TApp tokens. This ensures that everybody who buys into the bonding curve can sell their tokens back into it. The sell line will take time to move higher compared to the buy line for a token purchase at any point in time, which incentivizes investors to hold the token.

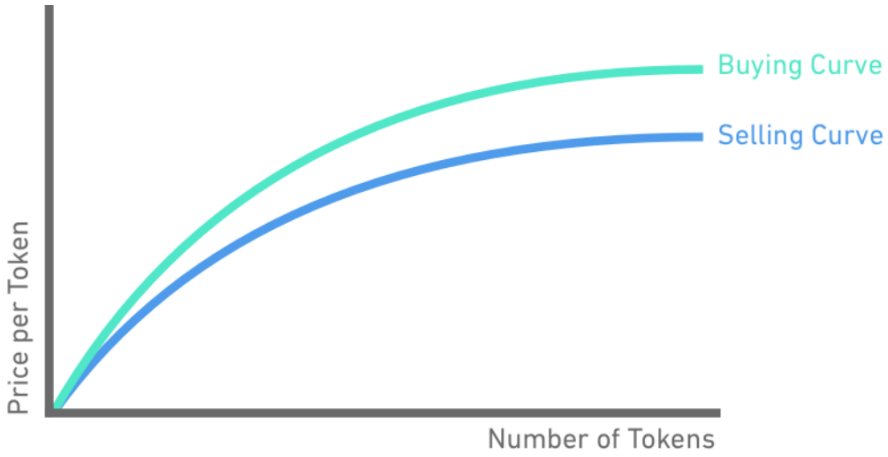

The area between the blue **sell** and the green **buy** curves in the graphic above is the portion that goes to developers from any funds going into the bonding curve. This is considered the developer's **funding pool**, a floating source of capital that is used to pay developers, licensing costs, or any other expenses related to the development of the TApp. As an example, a TApp's creators might have a bonding curve where the token's price = supply^2 with a theta of 0.03, meaning that 3% of every TApp token buy or consume event into the app for their own development expenses and salaries. The other portion, (1 - .03) = 97%, is put directly into the reserve pool and bonded directly into the bonding curve.

Bonding Curve Mechanics

The TApp token bonding curve has two different token payout algorithms depending on who supplies the funds into the bonding curve.

1. **A TApp user**: A TApp user pays TEA tokens in order to use a TApp. After the TApp owner takes their slice (TEA paid x theta), the rest of the TEA tokens are sent to the bonding curve to mint new TApp tokens. Since the user has already received utility from the TApp, they don't get any tokens in return. These newly minted tokens resulting from the user paying into the app are thus distributed to the existing token holders as a dividend. Therefore when a TApp gets used, the token holders get a dividend and the token price goes higher from the boost in token supply.
2. **A TApp token investor**: A TApp token investor is someone who simply buys tokens as an investment. The investor payment starts off a

process similar to when someone uses the TApp: the TApp owner is given their portion (TEA paid x theta), and the rest of the TEA tokens paid by the investor are sent to the bonding curve to mint new TApp tokens. The newly minted TApp tokens are sent to the investor's wallet, with the existing token holders still benefiting from the rise in the token price as no dividend is paid to them in this scenario.

Miners are also given token dividend payouts during app usage if the TApp creator elects to give miners staked tokens in exchange for their hosting. Staked tokens are unique to miners and cannot be sold. They function as rights to dividend payments when the TApp is used. For example, if the TApp creator has set aside 10 tokens for each miner hosting their TApp, then a miner will receive dividends equal to the amount as if they held 10 real TApp tokens.

Developers Funding

Developers can use TEA Project’s built-in bonding curve to generate investment in their TApp. This funding mechanism allows TApps to leverage expected future revenue into early development funding. The TAppStore is where the entire TEA ecosystem meets: developers to publish their TApps, miners can find the next TApp to invest their harvested TEA tokens, curators to publicize and invest in new trending TApps, and consumers to spend their TEA tokens on useful TApps they want to use.

TApp Tokens Help Onboard New Users

Because TApp token sales are a great vehicle for project teams to get funding, it benefits them to advertise their token sale to the public. Every time someone uses their TApp is another opportunity for the TApp to win over a possible investor.

Consumers of TApps are likewise incentivized to promote TApps they find useful. After buying into a TApp's tokens, they can then promote these TApps on social media in an effort to get others to invest in the TApp or use

it themselves. Either action will increase the supply of the token and therefore its price.

The TApp token economy acts as another path towards onboarding new users into the TEA Project's ecosystem. Consumers interested in TApps will exchange ETH to buy TEA which in turn will support the TEA token price and grow the TEA economy.

5.2 The 3-Phase Rollout

The TEA Project will use a 3-phase rollout that encompasses miners, developers, and consumers in exactly that order. A strong mining community must be developed to provide the infrastructure before developers are onboarded. Similarly, developers must develop compelling apps before consumers can be enticed to enter the ecosystem.

Each demographic will be encouraged to enter the TEA Project through a variety of methods.

1. Miners

The TEA Project aims to build a healthy ecosystem by starting with the miners. Miners harvest TEA tokens from hardware mining with CML.

- Miners "plant" CML into their mining hardware equipped with a TPM chip and a GPS module.
- Mining machines host Web3 applications and are rewarded in TEA token based on the app's consumed computing resources.

Miners are also able to exchange TEA in a liquid market with relative price stability as well as Miners can burn TEA to buy more CML. Through the self-interest of maximizing profit for their mining machines, the miners create the infrastructure necessary for the demographic in the next stage of our rollout, the developers.

2. Developers

After miners have built up the actual computing infrastructure, the focus shifts to onboarding developers. This segment of the rollout will include tech education & outreach on how to build on the TEA ecosystem.

- Hackathon events and grant program released.
- TEA SDK available helping developers build with the TEA dev framework.
- TApp store launched showcasing rich dApps running on the TEA platform.

Miners invest their TEA into TApp tokens which supports both early developers & TEA token price. Additionally, app revenue rewards the app developer, the hosting miners, and our next segment of users, the consumers.

3. Consumers

After miners and developers, next up is the consumer outreach phase. The rich TApps available in the TApp store are now marketed to consumers. We hope to see a positive feedback loop: as more consumers enter ecosystem, devs can see what apps consumers want. The devs then focus on making TApps that meet consumer demand, and popular TApps financially reward both miners and developers.

Additionally, consumers can invest in each TApp through its bonding curve and promote the TApp on social media. Becoming a curator for a burgeoning TApp helps the consumer (their TApp tokens they hold will increase in price along the bonding curve as more buyers push the supply higher). Their curation will also help the platform become better known on social media, which again shows how self-interested action within the TEA Project benefits the ecosystem on the whole.

Further Links on Ecosystem Development, Bonding Curves, and TApp Tokens

- How Devs & Miners Support Each Other in the TEA Project Economy. An explainer video about how the various participants work together on the TEA Project platform.
- Bonding Curve Theta determines what percentage of consume actions and TApp token purchases goes directly to the TApp developers.
- TApp Token Supply and Demand details how supply and price are correlated along the bonding curve.
- Staked TApp Tokens detail how miners are paid out in TApp tokens for hosting any particular TApp.
- TEAfluencer TApp gives a good overview on how potential users would invest in a TEAfluencer TApp.

6. Conclusion

Putting It All Together

Through the incentives mining for the TEA token as well as TApp token distributions, miners form the compute infrastructure ready to host decentralized applications known as **TApps**. This allows developers to deploy their apps to our platform including their existing web 2.0 apps because the TEA Project has the same three-tier architecture.

- The front end for TApps is IPFS.
- The server layer handles dev's code compiled to WASM. Dev code is encrypted and only decrypted to run in the TPM-protected enclaves of the network's layer-2 miners (B CML mining nodes).
- The database layer is handled according to the app requirements - NoSQL data is sent to IPFS as governed by B CML mining nodes, while relationship data is stored in GlueSQL as governed by A CML mining nodes.

The incentives provided by each TApp's bonding curve token acts as a source of fundraising for developers and an incentive to promote a TApp for its users and investors.

Although we have focused on many complicated blockchain and state-machine related topics in this whitepaper, we will not re-hash those points in the conclusion. It's important to us that developers can build on the TEA Project without having to be experts in blockchain design or having to learn a new language. Everything in the TEA Project is designed to look from the app's point of view like it's running on a normal cloud computing architecture, with the magical benefit that it's running fully decentralized. The TEA Project design handles these details so developers and decentralized app users don't have to think about them.