

Pensament

Version: 1.0.5

By: Jordan Fischer

Table of Contents

Table of Contents	2
Pensamento Blockchain	5
Consensus	5
What is Proof of Stake?	6
Stake Pools	7
Validation Process	8
Fees, Rewards, and Incentives	12
Block Rewards	12
Slashing	12
Transaction Fees	13
Ecosystem Fees	15
The Burn	15
The Treasury	16
Stake Pool Rewards	17
Saturation Point	18
Privacy	20
The Privacy Chain	21
How will it Work?	22
Regulation	23
Network Protocol	23
Programming Language	24
P2P Layer/Networking Stack	24
Serialization Format	24
Smart Contracts	25
Data Model (Account Type)	25
Data Storage	26
Governance	26
Native Cryptocurrency (\$PMTO)	26
Anti-Whale Measures	27
Anti-Dump Measures	27
Asset Allocation	28
Team Allocations:	29
Project Allocations:	29
Total Allocations:	29
Founders Asset Vesting Schedule	29
Pensamento ID	31
Pro Plans	33

Pensamento Points	34
Pensamento Cloud	35
Network Protocol	35
Underlay P2P	36
Overlay Immutable Storage	38
Addresses	38
The DHT	38
Chunks	39
Structure	40
Storage	41
Retrieval	42
Subsystems	43
Economics	44
Bandwidth Incentives	44
Storage Incentives	46
Lottery	47
Competitive Insurance	49
Blockchain Incentives	49
PPSF and PPSR	50
Staking	52
Network Storage	52
High-Level API Access	53
Application	54
Storage Tiers	55
Pensamento Swap	56
Regulation	56
pSwap	56
pSwap Plus	57
pSwap Pro	57

“At its core, Pensamento is a blockchain innovations company dedicated to realizing the original spirit of the Internet— focused on innovation, collaboration, and universal freedom. We understand that the Internet’s centralized approach has encroached upon this very spirit, with compromises to security, privacy, and individual autonomy becoming ever more common.

Blockchain technology has emerged as the key to this solution, providing a decentralized foundation for the Internet. However, decentralized tools are often complex, leaving users to choose between convenience and control. We view this choice as unjust and believe everyone deserves a great user experience without sacrificing their digital autonomy. The Pensamento ecosystem is designed to bridge this gap and usher in a new digital era we refer to as ‘Web4.’

Our mission is to redefine the Internet and its global impact. We envision Web4 as the embodiment of this mission and the Internet’s original spirit. Pensamento is unequivocally devoted to setting the new standard for user experience while safeguarding our users’ right to security, privacy, autonomy, and freedom above all else.”



Pensamento Blockchain

Consensus

As with any blockchain network, our first priority was to establish how the network would reach and maintain consensus. When designing the Pensamento Blockchain, we dedicated the first few months to researching and analyzing several different consensus mechanisms ranging from Proof of Work and Proof of Stake to diverse alternatives like Proof of Time, Proof of Space, etc. Throughout this process, we consistently found ourselves gravitating toward one protocol in particular: [Ouroboros](#). Ouroboros is the consensus mechanism utilized by the [Cardano](#) network, and it offers several advantages over traditional models, such as a unique Proof of Stake consensus mechanism, streamlined user participation through the use of stake pools, incredibly low energy consumption, and provable security backed by peer-reviewed research.

Ultimately, we found the [Ouroboros Chronos](#) framework to be the perfect fit. Ouroboros Chronos is a newer variation of the Ouroboros Protocol and incorporates a native time synchronization protocol to ensure that nodes remain synchronized to a central source of time—mitigating the chances of time-based vulnerabilities like [Front Running](#), [Eclipse Attacks](#), and [Double Spending](#). Simply put, Chronos offers increased security and network resilience to an already robust and secure protocol, making it the perfect foundation for the Pensamento Blockchain. As we discuss the design of the Pensamento Blockchain in more detail, it's important to note that we plan to implement several new additions to this protocol, such as Pensamento Smart Contracts and Native Smart Accounts, a first for Ouroboros.

As mentioned, we began our research by comparing several different consensus mechanisms to find one that was resilient, secure, and efficient. While there are several differences between consensus mechanisms like Proof of Work or Proof of Stake, the most prominent distinction lies in the mining process used during the validation process in Proof of Work networks. The mining process in proof-of-work networks like Bitcoin provides an additional layer of security through randomization, as the time to mine a new block is variable depending on who is actively participating in the mining process and the difficulty of the mining puzzle itself. (Learn more [here](#)) When a new block is created, validators on the network will race one another to see who can solve the puzzle first. The first person to find the correct answer wins the race and adds the block to the blockchain, earning the entire reward. This increased randomness actually serves as an advantage to these networks by reducing the likelihood of centralization or manipulation.

However, this mining process also requires excessive energy to power increasingly powerful computing systems as validators continuously upgrade and expand their validator setups to gain a slight advantage over their competitors. For example, the mining process used by the Bitcoin network alone consumed more energy than several countries, such as Argentina, Netherlands, and Ukraine. ([Source](#)) Furthermore, a majority of this energy is ultimately wasted by participants who do not win the race as they still consume vast amounts of energy in their attempt at solving the puzzle. To account for this, Proof of Stake networks eliminate the need for a mining process by individually selecting a validator to create and validate new blocks and earn the reward.

What is Proof of Stake?

As the name implies, Proof of Stake consensus mechanisms rely on staking to secure the network. In Proof of Stake, validators must allocate, or '**stake**,' a specific amount of the network's native cryptocurrency to become eligible to participate in the validation process. Once a user stakes cryptocurrency into the network, the assets are temporarily locked and cannot be traded until they are removed as stake. Most Proof of Stake networks enforce a 'lock-in' period where staked assets are inaccessible for a set duration, such as three months, six months, or even a year or more. This lock-in period is designed to enhance network security and stability as validators within a 'lock-in'

period are more likely to remain trustworthy, reliable, and fully committed to the network—discouraging short-term network manipulation.

Once a validator is set up and has staked a minimum amount of the network's native cryptocurrency, they are enrolled in the validation process to be chosen to create and validate new blocks for a reward. In most Proof of Stake networks, the chances of being selected as the validator for a new block are proportional to the amount of cryptocurrency staked in the network. The more a validator stakes, the more often they are selected to create and validate new blocks. While Proof of Stake has proven to be a fantastic alternative to traditional Proof of Work networks like Bitcoin, it also presents some shortcomings.

First, Proof of Stake networks still require participants to establish and maintain individual validator nodes, resulting in a high cost of entry. Expenses related to hardware, software, energy consumption, and the overall complexity of the process act as barriers to new users, leading to lower participation and a less secure network overall. Second, these networks are inherently less secure than Proof of Work networks due to the loss of the mining process, which adds an additional layer of randomization to the validation process. Lastly, this approach can lead to centralization, as validators with substantial amounts of staked cryptocurrency can grow to dominate the network, posing a risk of network manipulation. The Pensamento Blockchain aims to address these concerns to establish a provably secure, scalable, and efficient network.

Stake Pools

Our goal with the Pensamento Blockchain is to create an accessible network for everyone. One of the key reasons we've chosen to utilize Ouroboros is its unique implementation of *stake pools*. Stake pools are designed to encourage greater network participation and enhance network security, scalability, and decentralization by allowing users to participate in the network and earn rewards without maintaining their own validator node. Users on the network can earn rewards in two ways: by running a validator/stake pool or by delegating their stake to a stake pool run by someone else.

Stake pools are run by individual validators, also known as *'Stake Pool Operators'* or *'Operators'*, and can be set to either public or private. A private stake pool functions similarly to a traditional validator node, retaining all rewards. A public stake pool, however, is open to users who wish to participate in the network without running their own node. These users, also known as *'Delegators'*, can delegate their stake to these pools for a share of the total rewards. Furthermore, delegators will not be required to lock their stake for a predetermined duration and can trade these assets while participating in the network, regardless of how they are delegated.

Just like traditional Proof of Stake networks, validators are selected in proportion to their stake in the network. However, rather than relying on a node's *individual* stake, the Pensamento Blockchain will examine the *total* stake of the stake pool. The total amount of cryptocurrency staked by delegators and *'pledged'* by the validator/operator are combined to find the total stake of the stake pool. For example, let's assume a stake pool has five delegators. If each of the five delegators has staked 1,000 coins into the pool, and the stake pool operator has pledged an additional 1,500 coins, then the total stake of the pool would be 6,500 coins, which is the number used during the selection process. Please refer to this document's [Stake Pool Rewards](#) section to see how these rewards are distributed.

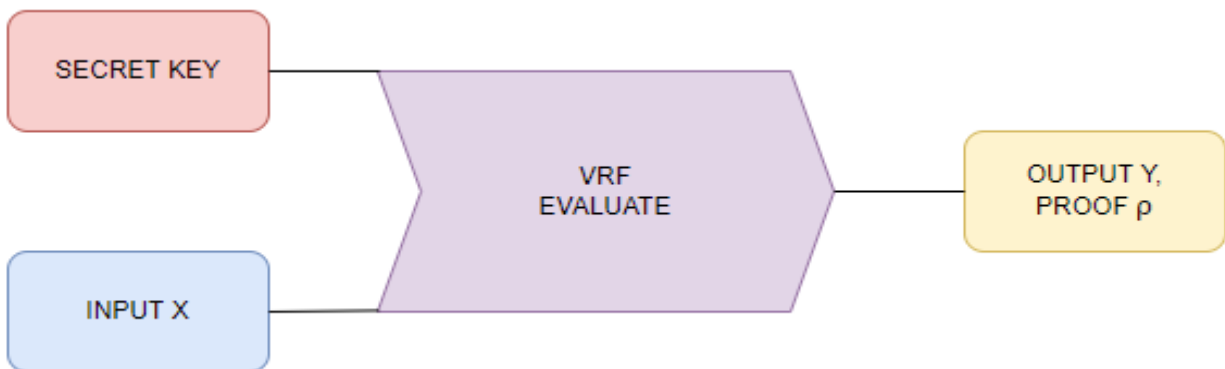
Validation Process

Now that we understand how stake pools work and what their role is within the network, let's discuss the validation process. One of the key reasons we've adopted a Proof of Stake consensus mechanism over Proof of Work is the elimination of a mining process. While Proof of Stake has been proven to be a scalable and efficient solution to Proof of Work, the removal of a mining process can lead to compromises in network security as the mining process brings an inherent randomness to the validation process. To overcome this, the Pensamento Blockchain will generate its own randomness and achieve the same security guarantees as Proof of Work networks without the need for an energy-intensive mining process.

First, the network is to be segregated into units of time known as *'slots'* and *'epochs.'* Each slot will last for one second and each epoch for five days, meaning there are 432,000 slots per epoch. As


new validators, or *'nodes,'* join the network, they must first synchronize their global clock and random oracle with the rest of the network before they can participate in the validator lottery mechanism, which is used to determine the node responsible for creating and validating the next block on the chain, also known as the *'slot leader.'*

This mechanism, which runs once per slot (i.e., once a second), functions by employing a Verifiable Random Function (VRF). At the start of each slot, nodes will run the VRF in an attempt to generate a number that falls below their individual threshold. This threshold is determined proportionally to the total stake of the nodes' corresponding stake pool, meaning that stake pools with a larger total stake have a greater chance of being selected to validate the next block and earn the reward. When a node runs the VRF, it will first generate a verification and secret key based on the random input. From there, it will run an *'evaluate'* algorithm to produce a pseudorandom output and proof based on the secret key and a message.



Public input x is evaluated in the VRF along with the node's unique secret key. Source: [Cexplorer.io](https://cexplorer.io)

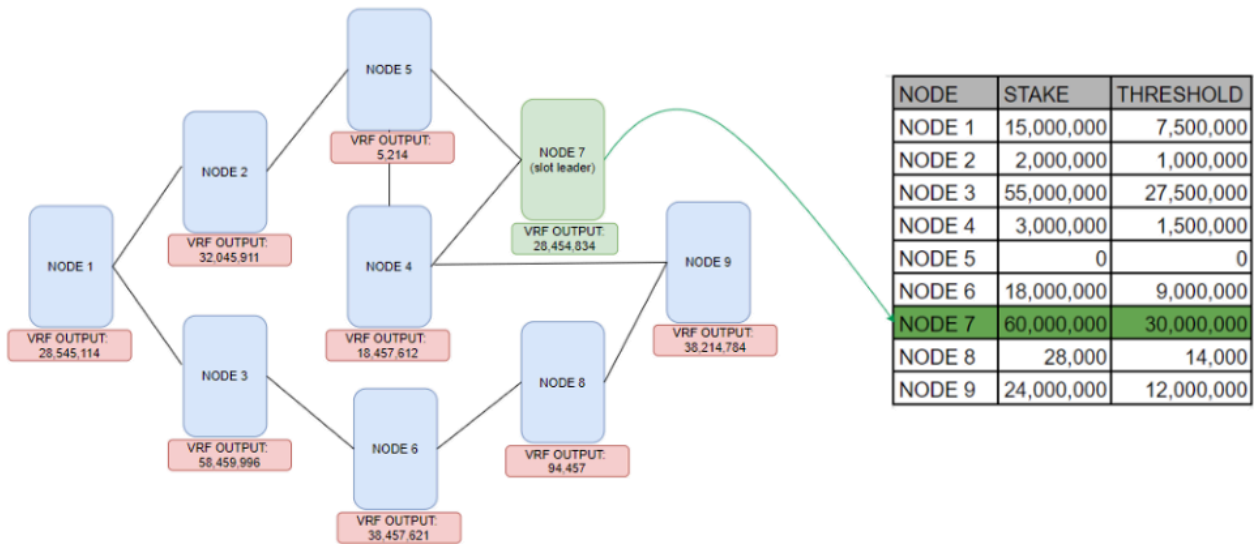
After the evaluation algorithm has run, the VRF will verify the output as either true (1) or false (0). This is what makes the function *'verifiable,'* as anyone on the network can take the nodes' public inputs and the public VRF key to validate that the node's number is true without revealing the node's private key, ensuring network integrity without sacrificing user privacy.



The results are then verified through the VRF with the corresponding public key. Source: [Cexplorer.io](https://cexplorer.io)


As mentioned, the VRF utilizes public inputs in its generation process. One crucial input employed during this process is the current Slot ID at the time the VRF is executed, helping to ensure that slot leaders are only selected in approximately 5% of new slots. In other words, despite validators being able to run the VRF once per slot, a slot leader will only be found approximately 5% of the time. This brings the missing layer of randomization found in proof of work networks, as malicious actors cannot pre-determine when a new slot leader is elected.

Once each node has run the VRF, the network will check to see if any nodes have generated a number below their given threshold. In the image below, we can see that node seven has successfully generated a number below its given threshold and has been selected as the next slot leader. Now that a slot leader has been selected, it can begin to generate and validate the new block. The slot leader will include the results from the lottery, also known as the ***VRF Proofs***, within the block headers to allow all other network participants the ability to verify its legitimacy.



Node Selection Process. Node 7 was elected the slot leader as it produced a number below its given threshold. Source: explorer.io

Once the new block has been added to the blockchain, all other nodes can verify the slot leaders' legitimacy through the VRF proofs. Each new block must be signed by a KES key (Key Evolving Signature cryptography) to add an additional layer of security and prevent attackers from being able to sign blocks, even if they somehow break the VRFs operations. (explorer.io)



At the end of each epoch, a Global Random Oracle will generate new randomness for the VRF in the upcoming epoch. This is accomplished by computing the hash of 2/3 of the prior epoch's random values and employing its value as the random seed for the staking procedure. This is where the name Ouroboros originates, as the term refers to a snake eating its tail, symbolizing an infinite loop. In the context of the Pensamento Blockchain, Ouroboros takes form as the network utilizing its own data as the source of randomness, establishing an infinite loop. (medium.com)

Fees, Rewards, and Incentives

Now that we understand how the protocol will function let's discuss how validators are rewarded and the overall structure of network fees and incentives. The Pensamento Blockchain is designed to offer a competitive fee and rewards system to offer network participants high rewards while ensuring that user fees remain low enough to achieve high adoption and scalability.

Block Rewards

As with any decentralized network, we first need to ensure that validators are properly rewarded for their contributions to the network. The Pensamento Blockchain will reward slot leaders that correctly generate, validate, and add new blocks to the blockchain by minting new \$PMTO coins, the native cryptocurrency of the Pensamento Blockchain. These rewards are to be determined algorithmically by the network as an incentive to participate in the network and behave honestly. These rewards are given to the slot leader and distributed throughout their corresponding stake pool in proportion to each delegator's individual stake within the pool.

Slashing

To prevent bad actors from participating or sabotaging the network, the Pensamento Blockchain will also incorporate a new slashing mechanism. If a validator is found to be acting maliciously, some or all of their individual stake can be automatically seized. Furthermore, bad actors also risk being either temporarily or permanently barred from future participation within the network if they are found to be dishonest or act maliciously.

Slashed validators will not directly impact their fellow stake pool participants, aside from the subsequent decrease in the pool's total stake in the network, which may result in smaller rewards. All fees collected from the slashing process are to be collected by the network and allocated to the [Pensamento Burn Pool](#), along with a 10% burn fee, at the end of each epoch to help offset the minting of new coins used as block rewards.

Transaction Fees

When a user initiates a new transaction on the network, they will be required to pay transaction fees to execute the transaction. These fees are determined by several factors and redistributed accordingly at the end of each epoch. The Pensamento Blockchain will implement a unique transaction fee structure consisting of four layers:

- **Base Fee** - The base transaction fee is determined by current network demand. This mechanism is designed to prevent network congestion and offer users a reliable and cost-effective experience.
- **Privacy Fee** - An *optional* fee for users opting to anonymize the data within a transaction. These fees are designed to offset the additional computational resources needed to encrypt data through cryptographic mechanisms like zk-STARKs. They will be paid directly to the validator responsible for processing the transaction.
- **Expedite Fee** - An *optional* fee for users who wish to prioritize their transactions. These fees are determined by a unique fee oracle in response to current network demand and the number of existing pending transactions. They will be paid directly to the validator responsible for processing the transaction.
- **Foundation Fee** - An *optional* fee that allows users to round up the total cost of the transaction to the nearest dollar and allocate the remaining funds to the Pensamento Foundations Fund.

When a user goes to submit a new transaction, they will be given the option to set the transaction as either *'Public'* or *'Private.'* Users will then see an option to *'Expedite'* the transaction, prioritizing it over other transactions, as well as an option to round to the nearest dollar and donate these funds to the Pensamento Foundation to help fund initiatives like Pensamento Academy, Pensamento Tech, and Pensamento Hubs. All transactions will be set as *'Public'* by default, as public transactions only require the base fee unless the user chooses to expedite it.

Transactions set as *'Private'* will face an additional fee to offset the computational resources required to process and validate anonymized data. Private transactions use additional cryptographic mechanisms like [zk-SNARKS](#) to anonymize data securely. These mechanisms rely on advanced mathematical techniques, such as elliptic curve pairings and polynomial commitments, and require additional resources to facilitate. These fees are to offset the costs to validators to run these mechanisms and will be determined algorithmically based on network demand, just as they are in the public network. Private transactions will include both the base and privacy fees unless the user also chooses to expedite it. Please refer to the [privacy](#) section of this document to learn more about this process.

Lastly, *'Expedite'* fees will allow users to prioritize their transactions over others waiting in the queue. These fees will also be determined algorithmically in response to current network demand to maintain stability and reliance on the network while incentivizing validators to prioritize transactions with higher fees. As of now, any transaction can be expedited.

Now that we understand how fees are collected let's discuss how they are allocated and distributed across the network. At the end of each epoch, the network will allocate the total *'Base fees,' 'Privacy fees,'* and *'Expedite fees'* into the *'Rewards Pool'* and all *'Foundation'* fees into the *'Foundations Pool'* which will be sent directly to the Pensamento Foundation.

Ecosystem Fees

Following the distribution of foundation fees to the Pensamento Foundation, the Rewards Pool will face additional fees before being redistributed to stake pools across the network. The **'Treasury Fee'** and **'Burn Fee'** are referred to as **'Ecosystem Fees'** and are designed to support the network and fund community-driven growth. These fees will be collected at a fixed rate of 10% each. Treasury Fees are automatically sent to the [Treasury](#), whereas Burn Fees are to be allocated directly to the Burn Pool. In comparison, the Cardano network currently allocates 20% of this pie solely to its treasury. ([Cardano.org](#))

The Burn

As mentioned, 10% of the ecosystem fees are to be redistributed to the **'Burn Pool'** along with any slashing fees collected throughout the epoch. Once these fees have been allocated to the Burn Pool, a unique **'Burn Oracle'** will send \$PMTO coins to the **'Burn Wallet,'** removing them from circulation indefinitely. The Burn Wallet is designed to hold these assets indefinitely and will have its keys destroyed to ensure that they remain inaccessible to anyone, even the Pensamento team. The Burn Oracle is designed to send only up to an equal amount of \$PMTO coins minted as rewards during the previous epoch to protect the \$PMTO cryptocurrency from inflation and keep the total circulating supply as close to 150,000,000 as possible.

Once the burn process is complete, a unique **'Validator Rewards Oracle'** (VRO) will be tasked with adjusting validator rewards for the upcoming epoch in response to the current state of the Burn Pool, as well as the total circulating supply of the \$PMTO cryptocurrency. The goal of the VRO is to adjust validator rewards to remain within a threshold of +/- 5% of the total circulating supply to remain as close to 150,000,000 as possible. This way, we can ensure fair compensation for validators while preventing the \$PMTO currency from inflating or deflating out of control.

For example, let's assume the network minted 1,000 \$PMTO coins as validator rewards and collected 1,100 \$PMTO coins from slashing and burn fees during the previous epoch. In this instance, the Burn Oracle will send 1,000 \$PMTO coins to the Burn Wallet to offset the 1,000 \$PMTO coins minted as rewards and retain the remaining 100 \$PMTO coins in the Burn Pool for use after the next epoch. Since the Burn Pool was left with a surplus, the VRO will slightly increase rewards for the upcoming epoch. However, if these numbers were reversed, meaning the Burn Pool was 100 coins short of offsetting the newly minted rewards, the VRO would slightly decrease validator rewards to bring the total circulating supply back down following the next epoch.

The Treasury

The remaining 10% of the Ecosystem Fees will be allocated to the Pensamento Treasury. Based on the Cardano Treasury, the Pensamento Treasury is designed as a fully decentralized, self-sustaining mechanism to fund community-driven projects, research, and initiatives that directly contribute to the growth and enhancement of the Pensamento Ecosystem. The platform will rely solely on community engagement and feedback with projects requiring the support of the Pensamento community. Anyone with a valid Pensamento ID can pitch ideas and vote on new proposals so long as they hold any amount of \$PMTO cryptocurrency in their wallet.

Community members, planners, developers, and even organizations can submit proposals to request funds from the Treasury to research, design, and develop new projects or initiatives that benefit the Pensamento Ecosystem. These proposals are subject to a thorough review process by the community and must receive support from the community to receive funding. Projects that gain sufficient community support will be allocated a predetermined amount of funding directly from the Treasury to support the research and development of the proposal.

The total funding amount for each project is split into phases, with each phase being allocated a percentage of the total funding amount. Project teams will be required to submit valid and

transparent updates prior to the end of each phase, as well as maintain community support and approval to continue receiving funds from the Treasury. Projects that do not meet these requirements will forfeit the remaining funding allocations, which will be reallocated to the treasury for future proposals. This self-sustaining model is vital to Pensamento's commitment to long-term growth, scalability, sustainability, transparency, and decentralization.

Stake Pool Rewards

Once all ecosystem fees have been distributed, the remaining 80% of total rewards are to be allocated to stake pools throughout the network in proportion to their total stake. The default setup for pool rewards is to first allocate 0.00000075% of the total circulating supply as a '*Pool Operator Fee.*' The network sets this fee and cannot be changed by pool operators.

For example, if the network has a total circulating supply of 150,000,000, each pool operator would be allocated 1.125 \$PMTO coins. This fee is intended to help offset the costs required to host and maintain a stake pool. The network will then allocate a '*Variable Margin Fee*' set by the operator, which can be anywhere from 0-10%. This fee is paid directly to the operator and must be visible to all network participants to ensure that delegators can choose a stake pool with fair fees, incentivizing transparency, and fairness across the network. Once these fees have been allocated to the operators, the remaining rewards are split amongst all pool participants, including the operator, in proportion to their *individual* stake in the pool.

For example, let's assume a stake pool with 100 participants has earned a total reward of 10 \$PMTO coins. With a total circulating supply of 150,000,000 tokens, the operator fee equates to 1.125 \$PMTO. The operator set the variable margin fee to 3%, meaning that an additional 0.26625 \$PMTO coins are also allocated to the operator. After all operator fees have been allocated, the total remaining rewards are 8.60875 \$PMTO coins, which will then be dispersed amongst each


participant in proportion to their individual stake in the pool. Assuming each participant holds an equal stake, this would equate to a total of 0.0860875 \$PMTO coins per pool participant.

Total Stake Pool Rewards

Total Rewards	
Total Rewards	10

Rewards Allocation	
Operator Fee	1.125
Variable Margin Fee	0.26625
Total Participant Allocation	8.60875

Rewards Per Participant	
Total Number of Participants	100
Stake Proportion	N/A
Rewards Per Participant	0.0860875



The above example shows that 14% of the total rewards are allocated to the operator, and the remaining 86% are split amongst each participant, including the operator, in proportion to their stake. This approach will create a more secure network by incentivizing users to increase their stake. However, this incentive can also lead to centralization as larger stake pools can accumulate disproportional network dominance and receive a majority of rewards.


Saturation Point

To mitigate this risk, stake pool rewards are designed to increase proportionally alongside its total stake, but only up to a certain point. This point is referred to as the *'saturation point.'* As stake pools reach the saturation point, rewards will stop increasing, even if the total stake of the pool

continues to grow. The saturation point encourages decentralization and ensures that smaller stake pools still have a fair opportunity to earn rewards.

The saturation point will be determined by dividing the total circulating supply by a factor of k . At launch, k will equal 150, meaning that with a total circulating supply of 150,000,000 coins, stake pools can continue to increase rewards until they reach the saturation point of 150,000 coins. Once a pool becomes oversaturated, any additional stake will no longer equate to additional rewards. In other words, stake pools will continue to earn higher rewards until the total stake of the pool reaches 1/150th of the total circulating supply, at which point rewards stagnate. The saturation point aims to incentivize pools to grow their total stake while disincentivizing centralization and network dominance.

The value of k may be updated by the protocol at any time to account for network demands, but the starting rate for k will be 150. The graph below demonstrates the saturation point in relation to total stake and subsequent rewards up to the saturation point. Please note that the numbers used in the graph are for illustrative purposes only and do not represent real rewards.



Privacy

As part of our commitment to ensure security, privacy, and autonomy, the Pensamento Blockchain will implement additional cryptographic mechanisms, such as zk-SNARKs, to safeguard user data. zk-SNARKs are a type of [zero-knowledge proof](#) that utilizes advanced cryptographic algorithms, such as elliptic curve cryptography (ECC), to anonymize data and ensure it is verifiable without revealing the data itself.

For example, imagine you were going into a bar and needed to prove to the bouncer that you were over 21 but didn't want them to know your age. With zero-knowledge proofs, you can securely prove to the bouncer that you are, in fact, over the age of 21 without revealing any other information about you. If the bouncer needs you to prove your age, you can provide the proof hash to verify your age. The bouncer can then use this hash to check against the original information on the blockchain to confirm your age and ensure your privacy while taking advantage of the immutability and transparency of the Pensamento Blockchain.

While this is just one example of how zero-knowledge proofs can be incorporated into the real world, the list of possibilities is seemingly endless. Whether privatizing transactions or storing sensitive data like banking, health, or tax information, zero-knowledge proofs offer users and developers a new level of anonymity when they need it most, ensuring privacy and anonymity throughout the Pensamento Ecosystem.

While there are several types of zero-knowledge proofs, the Pensamento Blockchain will feature zk-SNARKs specifically due to their small-proof sizes and fast verification times. While zk-SNARKs are a fantastic option for anonymizing data, they don't come without limitations. For example, zk-SNARKs present a lack of plausible quantum-resistance and a reliance on a trusted setup. Thankfully, there are already a few new iterations of the zk-SNARK algorithm to address these concerns, such as [Redshift SNARKs](#) by zk-Sync, which we plan to adopt in the Pensamento Blockchain.

The Privacy Chain

However, for users looking for unparalleled privacy, we are also developing a dedicated privacy sidechain we refer to as the Pensamento Privacy Chain. The privacy chain, or *'pChain Plus,'* is designed to offer users, businesses, and developers a dedicated platform to take advantage of the latest cryptographic algorithms to build and launch private dApps, achieve quantum-resistance, etc. This offers increased user data security and opens the door for businesses in highly regulated sectors, such as healthcare, finance, government, etc., to adopt blockchain technology while remaining compliant with certain regulations through the Pensamento Ecosystem.

The privacy chain will integrate advanced cryptographic algorithms, including [zk-STARKs](#), [stealth addresses](#), and [ring signatures](#), to further enhance privacy and security. Unlike traditional zk-SNARKs, zk-STARKs employ more advanced cryptographic techniques, such as low-degree polynomials and error-correcting codes, to generate and validate proofs. This approach offers enhanced security and presents plausible quantum resistance, meaning the data can remain secure even against potential attacks from quantum computers. However, the proof sizes for a zk-STARK are much larger than those of traditional zk-SNARKs. For example, the average proof size of a zk-STARK is currently around 142x larger than a zk-SNARK proof, resulting in higher processing and storage requirements and slower verification times.

Trusted setup		
zk-SNARKs		
Prover	Verifier	Size
2.3s	10ms	288B
Very fast	Fastest	Smallest

Bulletproofs		
Prover	Verifier	Size
30s	1100ms	1,3KB
Slowest	Slowest	Middle

zk-STARKs		
Prover	Verifier	Size
1.6s	16ms	>40KB
Fastest	Very fast	Big

Despite the larger proof sizes found with zk-STARKs, we believe this zero-knowledge proof will cement itself as the next industry standard, and we are keen on implementing it within the Pensamento Privacy Chain.

How will it Work?

When interacting with the Pensamento Ecosystem, users will be presented with the option to set data as either public or private. If a user chooses to set a transaction or another form of data as private, they will be presented with two additional options: *Private* and *Private Plus*.

- **Private:** Transactions labeled as 'Private' will utilize zk-SNARKs directly on the main Pensamento Blockchain. Their smaller proof sizes of 515 KB and 80-bit security are a great option for those looking for simple privacy without a substantial change in cost or validation times.
- **Private Plus:** Transactions labeled as 'Private Plus' will take advantage of the more advanced cryptography on the privacy chain, such as zk-STARKs. These transactions will incur larger

fees due to the larger proof sizes (40KB) and slower validation times of zk-STARK cryptography.

- **Private Pro:** There are also plans to add even more secure cryptographic algorithms, such as [homomorphic encryption](#), in future updates to the privacy chain; however, these techniques are impractical to use in their current form.

Regulation

While our main goal is to provide our users with the most secure and private platform possible, we also recognize the importance and necessity of adhering to regulations, such as Know Your Customer and Anti-Money Laundering laws. With this in mind, these robust privacy features will be rolled out gradually, following a comprehensive vetting process, to ensure compliance with regulators without compromising the integrity of our privacy objectives.

Network Protocol

Now that we have established Consensus, Stake Pools, Validation, Fees and Rewards, and Privacy, let's dive into the Network Protocol itself, including:

- Programming Language
- P2P Layer/Networking Stack
- Serialization Format
- Smart Contracts
- Data Model (Account-Based)
- Data Storage
- Governance
- Native Cryptocurrency (\$PMTO)

Programming Language

The Pensamento Blockchain will be written in the Haskell programming language due to its high security, performance, assurance, concurrency control, formal verification, and interoperability.

P2P Layer/Networking Stack

The Pensamento Blockchain will utilize [LibP2P](#), a modular networking framework, to facilitate efficient and secure peer-to-peer communications within the network. As a modular approach, LibP2P offers a collection of open-source protocols, specifications, and libraries that Pensamento can utilize across other products. It offers support for multiple transport options like [WebSockets](#), [TCP/IP](#), and [QUIC](#), making it flexible and adaptable to several different network conditions, as well as built-in features for peer discovery and routing, which are crucial for effective communication between network nodes in the Pensamento Ecosystem.

LibP2P also brings an established presence in the blockchain space and is adopted by several other networks, such as [Ethereum](#) and [IPFS](#). This not only streamlines our development but fosters compatibility with other blockchain ecosystems, offering Pensamento users and developers even more flexibility without sacrificing security. This flexibility also benefits the Pensamento Ecosystem, allowing other Pensamento products, such as the [Pensamento Cloud](#), to integrate LibP2P as its peer-to-peer protocol and establish a more secure and seamless connection between the two services. In short, LibP2P aligns perfectly with the core vision of the Pensamento Ecosystem, offering a secure, private, and autonomous future for both users and developers.

Serialization Format

The Pensamento Blockchain will utilize the [Concise Binary Object Representation](#) (CBOR) serialization format. CBOR is a highly efficient binary data format that encodes and decodes structured data in decentralized networks. CBOR is also an [Internet Engineering Task Force Request for Comments](#) (IETF RFC), signifying its status as a well-established and widely accepted standard for compactly representing structured data in binary form. By pairing CBOR with LibP2P, the Pensamento

Ecosystem can ensure secure and efficient data exchange between network peers, maintaining the ability to operate independently and harmoniously with other networks.

Smart Contracts

[Smart contracts](#) are a key component of the Pensamento Ecosystem. Simply put, smart contracts are scripts of code designed to execute when pre-determined conditions are met. First introduced in the Ethereum network, smart contracts have proven to be incredibly powerful tools. They can automatically process payments, transfer ownership of digital assets, manage supply chains, and more, all in a safe, secure, and trustless manner.

Regarding the Pensamento Ecosystem, smart contracts are required to establish our innovative [Pensamento IDs](#) that serve as a user's digital ID and wallet. Due to its heightened security and reliability, Pensamento Smart Contracts will be written in the [Plutus](#) programming language to ensure that the Pensamento Blockchain maintains a high level of trust and security for all its users.

However, we also recognize the importance of developer flexibility and autonomy and don't feel it's fair to require all Pensamento developers to learn the Plutus language if they are already comfortable with other scripting languages like [Solidity](#) or [C++](#). To accommodate this, there are future plans to create a native translator within the Pensamento Developers Portal to translate code in real-time without sacrificing native support within the ecosystem, ensuring that the network remains adaptable and developer-friendly.

Data Model (Account Type)

We plan to utilize an Account-Based model approach for the Pensamento Blockchain. This is in contrast to the UTXO model used by Cardano. However, an account-based approach aligns better with our vision than a UTXO model. (See more [here](#)) This will simplify our approach to implementing Native Account Abstraction to power Pensamento IDa and help streamline the user and developer

experience throughout the Pensamento Ecosystem. For example, a big part of the Pensamento Ecosystem lies in our innovative [Pro Plans](#), an innovative subscription model that allows users to pay a monthly fee for additional discounts and benefits throughout the ecosystem.

Data Storage

Blocks on the Pensamento Blockchain are currently designed to offer a maximum size of roughly 3,500,000 bytes, or 0.0035 GB per block. Validators on the Pensamento Blockchain will use random access memory (RAM) as temporary storage during validation. This is a standard practice among blockchain networks; however, we do have future plans to possibly utilize [Pensamento Cloud](#) to store an easily accessible, dynamic state of the Blockchain ledger that is updated in real-time.

Within the Ouroboros protocol, each network node retains a copy of the transaction mempool, where transactions are added if they are consistent with existing transactions, and the blockchain. The locally stored blockchain is replaced when the node becomes aware of a newer, longer valid chain. ([Source](#))

Governance

We currently plan on incorporating an on-chain governance mechanism for proposals and upgrades to the Pensamento Blockchain. For example, each verified account holding at least 1 \$PMTO coin will receive one vote and can either vote themselves on new proposals or delegate their vote to another user. If a proposal meets a simple majority, it will be added to the network without the need for a hard fork. While innovation moves quickly, and a future hard fork may be inevitable, this governance system will help us limit the number of hard forks that may be required down the line. Our plans are similar to [CIP-1694](#) for reference.

Native Cryptocurrency (\$PMTO)

As with any blockchain network, the Pensamento Blockchain will launch with a native cryptocurrency, \$PMTO. This asset will launch with a total circulating supply of **150,000,000**. Our

main goal when designing our native cryptocurrency is to ensure user safety and compliance with regulations. The \$PMTO coin will launch with innovative Anti-Whale and Anti-Dump measures, a dynamic sales tax for larger sales, and a mandatory vesting period for founder allocations.

Anti-Whale Measures

- No single wallet can purchase or hold more than **4.99%** of the total circulating supply at any given time, with the exception of the Pensamento Growth and Maintenance Fund (Admin Wallet) and the Founder/CEO (Jordan Fischer). This is designed to ensure the project remains decentralized and abides by the current disclosure requirements in the United States.

Anti-Dump Measures

- No individual wallet can sell more than **1.0%** of the circulating supply (**1,500,000**) within a 24-hour period.
- All sales exceeding **0.1%** of the total circulating supply within a 24-hour period will incur a sales tax in proportion to the size of the sale for all tokens that fall between this threshold.

Total Circulating Supply Percentage Within 24 hours	Tax Rate on Sale(s)
0.0% - 0.1%	0%
0.1% - 0.2%	1%
0.2% - 0.3%	2%
0.3% - 0.4%	3%
0.4% - 0.5%	4%
0.5% - 0.6%	5%
0.6% - 0.7%	6%
0.7% - 0.8%	7%
0.8% - 0.9%	8%
0.9% - 1.0%	9%

For example, if a wallet sells **0.5%** of a total circulating supply of **150,000,000 (750,000)** within 24 hours, the following taxes would be charged to the sale:

- **0% – 0.1% = 0%** Sales Tax on first **150,000** tokens (**0** tokens)
- **0.1% – 0.2% = 1%** Sales Tax on second **150,000** tokens (**1,500** tokens)
- **0.2 % - 0.3% = 2%** Sales Tax on third **150,000** tokens (**3,000** tokens)
- **0.3 % - 0.4% = 3%** Sales Tax on fourth **150,000** tokens (**4,500** tokens)
- **0.4 % - 0.5% = 4%** Sales Tax on fifth **150,000** tokens (**6,000** tokens)
- **Total Sale = 750,000** tokens (**0.5%** of total circulating supply)
- **Total Sales Tax = 15,000** tokens

All collected sales tax is to be redistributed to the network at the time of sale:

- **2.5%** to every holder in proportion to their total percentage of tokens held.
- **2.5%** to stake pool participants in proportion to the individual stake amount
- **2.5%** to liquidity pools for price stabilization
- **2.5%** sent to burn pool

Note: Sales tax is not included in the total **1.0%** sale limit. The sale limit is **1.0%** of the total circulating supply within a 24-hour period, excluding any accrued sales tax. In other words, a holder can sell **0.5%** of the total circulating supply (**750,000**) but it will cost them a total of **765,000** after tax.

Asset Allocation

Lastly, we want to ensure we are transparent with our community and wish to share our current plans regarding asset allocations for our founders, as well as our planned vesting schedule, to highlight our long-term commitment to this project.

Team Allocations:

- **Jordan Fischer** (Founder and CEO) - **9,000,000 (6%** of the total circulating supply)
- **Joshua Jackson** (COO) - **6,000,000 (4%** of the total circulating supply)
- **Elijah Fullam** (CIO) - **4,500,000 (3%** of the total circulating supply)
- **CTO** (TBD) - **4,500,000 (3%** of the total circulating supply)
- **Roberto Piraino** (Marketing Strategist) - **4,500,000 (3%** of the total circulating supply)
- **Jayson Edwards** (Design Affiliate) - **1,500,000 (1%** of the total circulating supply)

Project Allocations:

- **Administrator Wallet** - **4,500,000 (3%** of the total circulating supply)
- **Pensamento Growth and Maintenance Fund**- **10,500,000 (7%** of the total circulating supply)
- **Pensamento Foundations Fund** - **3,000,000 (2%** of the total circulating supply)
- **Pensamento Treasury** - **1,500,000 (1%** of the total circulating supply)

Total Allocations:

- **Total Pre-allocation** - **49,500,000 (33%** of the total circulating supply)
- **Total available for public trade** – **100,500,000 (67%** of the total circulating supply)

The purpose of this allocation is to ensure that the public maintains a Super Majority of holdings by making **67%** of the total circulating supply available at launch. One thing to note is that the holdings for the Pensamento Growth and Maintenance Fund, Foundations Fund, and Treasury are designed to be used to fund different aspects of the project and will decrease over time, further increasing decentralization.

Founders Asset Vesting Schedule

- **100%** of the founder's allocations will vest at a rate of **2.77%** per month for **36** months (3 years), subject to continued employment with Pensamento. In the event of employment termination

for any reason, any remaining funds are to be transferred to the Pensamento Growth and Maintenance Fund.

Sale or Transfer of Funds (Founders):

\$PMTO tokens cannot be purchased, sold, or traded in any capacity by Pensamento employees during “blackout” periods. These periods will be determined on a quarterly basis and are designed to ensure there is no insider trading or market manipulation done by Pensamento employees.

Sale or Transfer of Funds (The Company):

Funds allocated to the Pensamento Growth and Maintenance Fund and Foundations Fund are available to use on day one. These sales will require multi-sig from two or more project leaders and must be documented and publicized with the following details:

- Reason for sale
- Who approved the sale (Must be a Super Majority of C-Suite and Directors)
- Who signed the sale (Must be C-Suite and/or Director)
- Total sale amount

For sales larger than **X**, a community vote will also be needed to ensure most holders support the sale's purpose. Each wallet holding at least **1** \$PMTO token is eligible to vote, and a simple majority is needed to pass.



Pensamento ID

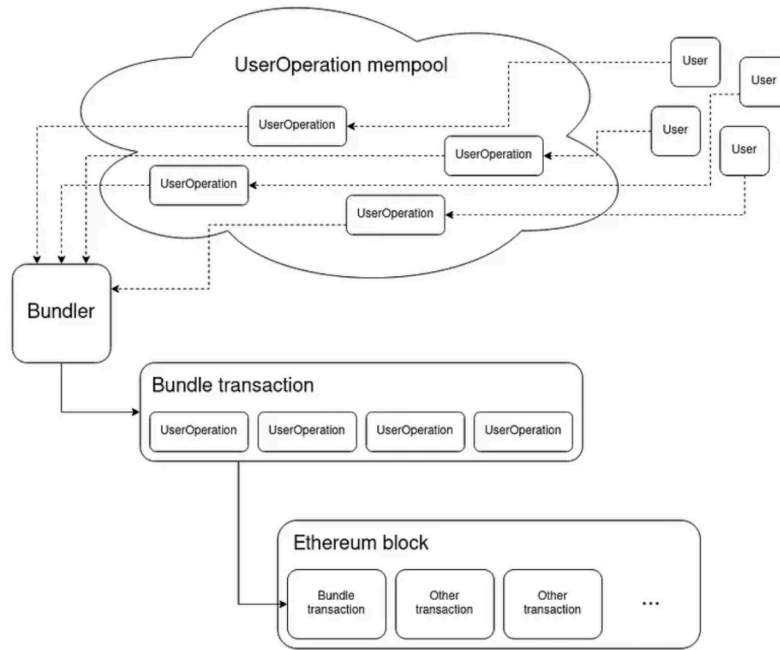
A core component of the Pensamento Ecosystem is the Pensamento ID. Pensamento IDs are native smart accounts built directly into the blockchain layer of the Pensamento Ecosystem. They are designed to offer users and developers unparalleled flexibility and control over their data while also providing a seamless user experience. Pensamento ID was designed with three characteristics in mind:

- A Native Smart Account/Wallet
- A Decentralized Identification Protocol
- The Key to Pensamento

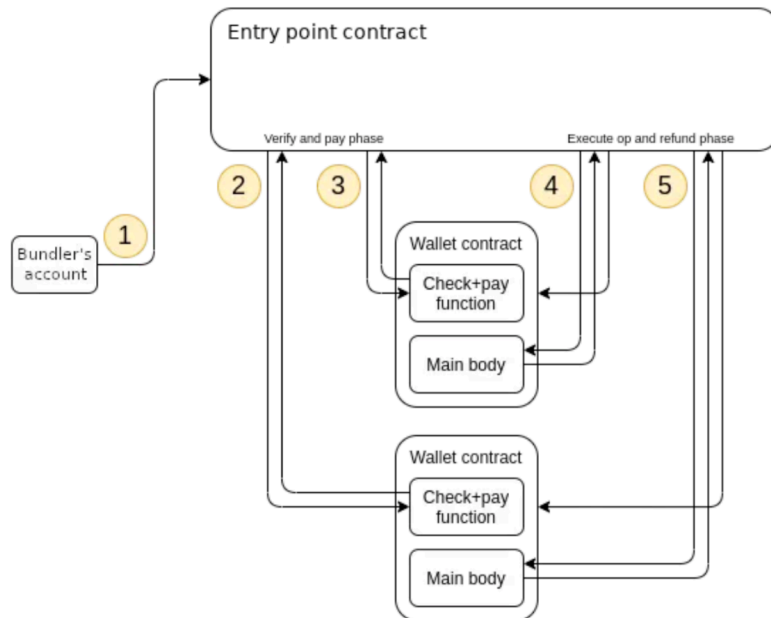
Simply put, these characteristics aren't possible with traditional crypto wallets or accounts. Externally Owned Accounts (EOAs) have been the industry standard for years and are used by almost every network to send and receive transactions. However, these accounts are inherently limited in their design due to their basic functionality, reliance on private keys, and lack of features such as account recovery, delegation, and biometric verification. For example, the reliance on private keys is common practice in the crypto space even though it presents a single point of failure for user accounts, acting as the single largest inhibitor to mainstream adoption.

We are determined to right this wrong within the Pensamento Ecosystem by shifting away from EOAs in favor of Smart Accounts. Smart Accounts leverage **'Account Abstraction'** to transform EOAs into individual smart contracts, empowering seemingly limitless possibilities. Account abstraction was first introduced to the Ethereum network with the [ERC-4337](#) upgrade. This upgrade

allows Ethereum accounts to take advantage of the flexibility and customization of smart accounts. However, the ERC-4337 upgrade required a number of key technical components to be added to the network, increasing complexity. (See below)



Account abstraction architecture. Source: [Ethereum Improvement Proposals](#)



Account abstraction control flow. Source: [Ethereum Improvement Proposals](#)




We believe this inherent complexity to be unnecessary and plan to implement Native Smart Accounts directly into the blockchain layer of the Pensamento network. This not only simplifies the processes within the network but also offers a few crucial benefits to the ecosystem for both users and developers.

- **Smart** - Pensamento is designed to be a ‘Smart’ blockchain network. As a result, all Pensamento accounts (Pensamento ID) are smart accounts. EOAs will not exist within the ecosystem.
- **Simple** - Account abstraction also allows for more complex smart contract logic while keeping the underlying blockchain protocol relatively simple. All infrastructure within the ecosystem is designed to take full advantage of Pensamento smart accounts and also simplify the development process, as developers no longer need to take EOAs into consideration.
- **Security and Privacy** - Pensamento smart accounts also create a more secure and private user experience with multi-factor and biometric authentication, account recovery processes, social recovery processes, account privileges, and delegation.
- **Subsidization** - Smart accounts will allow users to pay transaction fees using any currency by default, further simplifying the user experience. They also open the door for subscription models for platforms and dApps.
- **Account Automation** - Automate execution of transactions based on smart contract criteria, such as pre-approving transactions while gaming or when day-trading.
- **Customization and User Experience** - Unlock the full potential of a smart account to unlock a decentralized user experience unlike anything before.

Pro Plans

Pensamento ID also opens the door to our revolutionary Pensamento Pro Plans. Pro Plans are designed to offer Pensamento users a standard monthly fee for benefits that extend across the ecosystem. These can include discounted transaction fees, discounts in the Pensamento store, set monthly storage limit within Pensamento Cloud, and discounted trading fees or additional perks

within Pensamento Swap. We believe that Pro Plans will offer a streamlined approach for Pensamento users around the world and, when paired with Pensamento ID and the rest of the Pensamento Ecosystem, pave the way to onboarding the next billion users to a truly decentralized future with Web4. See our first example of how we envision Pensamento Pro Plans below.

 <p>PRO</p> <p>\$9.99</p> <ul style="list-style-type: none"> • Pro Badge Next to PID • 25% Discount on Transaction Fees • 2X PENS Points on Every Trade/Purchase • 10% Discount Across the Store • Free Yearly Docs • Access to Pro Bonuses, Airdrops, and Partner Discounts • Invitation to Pro - Only Comms Channels • 250GB pCloud Storage • 2X Cache Storage • Pro Trading Fees (25% \$PENS Trades, lower fees with higher monthly trading value) 	 <p>PRO PLUS</p> <p>\$14.99</p> <ul style="list-style-type: none"> • Pro Plus Badge Next to PID • 50% Discount on Transaction Fees • 3X PENS Points on Every Trade/Purchase • 15% Discount Across the Store • Free Yearly Docs • Access to Pro Plus Bonuses, Airdrops, and Partner Discounts • Invitation to Pro Plus - Only Comms Channels • 1TB pCloud Storage • 2X Cache Storage • 2X Data Availability • Pro Plus Trading Fees (25% \$PENS Trades, lower fees with higher monthly trading value) 	 <p>VIP</p> <p>\$19.99</p> <ul style="list-style-type: none"> • VIP Badge Next to PID • 75% Discount on Transaction Fees • 5X PENS Points on Every Trade/Purchase • 25% Discount Across the Store • Free Yearly Docs • Access to VIP Bonuses, Airdrops, and Partner Discounts • Invitation to VIP - Only Comms Channels • BETA Access and Participation • 2TB pCloud Storage • 2X Cache Storage • 2X Data Availability • Proximity Storage (When Applicable) • VIP Trading Fees (25% \$PENS Trades, lower fees with higher monthly trading value) • pSwap Pro • Trade Prioritization
--	--	--

**The above example is for demonstration purposes only and does not reflect the final product.*

Pensamento Points

Now, you may have noticed ‘PENS Points’ as a benefit to Pensamento Pro Plans, and those are another benefit we plan to offer Pensamento users thanks to our native smart account and Pensamento ID. Pensamento Points are designed as incentives and rewards for Pensamento Pro users for using the ecosystem and participating in the network itself. Users can then redeem these points for discounts across Pensamento or any of its partners or even cash them out for cryptocurrency. We are incredibly excited about Pensamento Pro Plans and the PENS Points system and believe these will prove to be monumental in the world of decentralization.

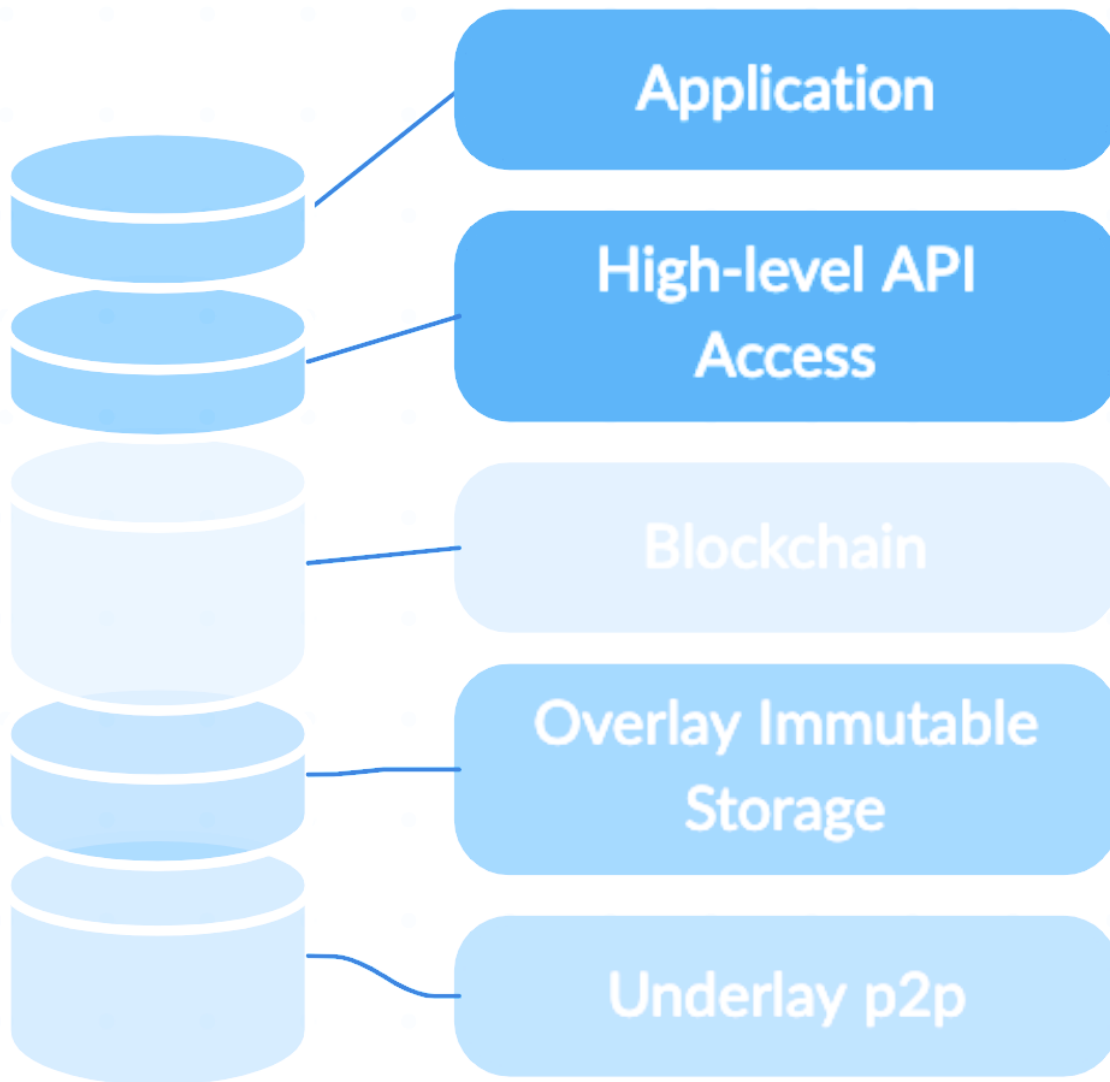


Pensamento Cloud

Pensamento Cloud, or 'pCloud,' is a fully decentralized cloud storage protocol that also serves as the computational core of the Pensamento Ecosystem. It provides users and developers with a safe, secure, and privacy-preserving cloud network, which can be used to store data, host content, or leverage to build uniquely powerful dApps. The Pensamento Cloud is inspired by other decentralized networks like [Ethereum Swarm](#) and [IPFS](#). However, Pensamento Cloud includes its own set of adjustments to weave itself tightly within the Pensamento Ecosystem.

Network Protocol

Pensamento Cloud consists of five layers that form the infrastructure of the network, of which the Overlay Network, Blockchain, and API layers form the core of the Pensamento Cloud protocol. See a breakdown of the network protocol below.



Underlay P2P

The Underlay P2P layer of Pensamento Cloud will utilize LibP2P, just like the Pensamento Blockchain, and will serve as the low-level communications network and foundation of the Overlay Network. This design enables greater flexibility and interoperability with other networks, reduces complexity and costs related to development, and aids in establishing a deep integration with other Pensamento products. By adopting LibP2P within Pensamento Cloud, we can establish a stronger,

more secure, and native integration between Pensamento Cloud and the Pensamento Blockchain. The underlay network for Pensamento Cloud must satisfy the following:

- **Addressing** - Nodes can be successfully identified by their underlay address.
- **Dialing** - Nodes can establish safe and secure connections with their peers by dialing their corresponding underlay address.
- **Listening** - Nodes on the network can listen to other nodes dialing them and accept new connections made on the network.
- **Sustained Live Connection** - Nodes can connect and maintain a channel of communication in the network until explicitly disconnected so that the existence of a connection means the remote peer is online and accepting messages.
- **Channel Security** - The channel can verify identity and implement encrypted and authenticated transport, resisting “man in the middle” attacks.
- **Protocol Multiplexing** - A node can establish common protocols with new connections by sharing the name and version of the protocols they have implemented to find a match with a new connection, so long as the underlay network supports their protocol. Once common protocols are established, peer connections are generated for each match.
- **Delivery Guarantees** - Protocol messages within the system ensure guaranteed delivery, meaning that if there are any network issues causing delivery failures, an immediate error response is generated. The order of message delivery within each protocol is guaranteed, and the underlay protocol should provide prioritization. If protocol multiplexing is over the same transport channel, this implies framing so that long messages do not block higher-priority messages.
- **Serialization** - The protocol message construction supports arbitrary data structure serialization conventions.

(Source: [The Book of Swarm - 2.1.1](#))

Overlay Immutable Storage

Pensamento Clouds Overlay Immutable Storage Layer will determine how the network allocates, distributes, stores, locates, and retrieves data stored throughout the network. The Pensamento Cloud storage model utilizes a Distributed Immutable Store of Chunks (DISC) model that was developed and introduced by the Ethereum team with the launch of Swarm. The DISC forms the basis of how nodes communicate with each other. The DISC will utilize a Kademlia DHT to distribute and organize data directly within the network.

Addresses

Storage nodes joining the network will first be assigned a *Network Address* and a *Cloud Address*. The Cloud Address is to be used for data storage and allocation, whereas the Network Address is to be used to facilitate communications between storage nodes directly. Another way to think of this is the Cloud Address is for *data*, and the Network Address is for *communication*.


- **Network Address** - Based on the storage providers associated Pensamento ID using the Pensamento Naming Service (PNS) protocol to facilitate communications between other storage nodes within the network.
- **Cloud Address** - Assigned to storage nodes for data storage and allocation within the Pensamento Cloud network. These addresses are used by the Kademlia DHT for proximity-based allocation and retrieval of data. The network will allocate data chunks to storage nodes with the closest proximity to the corresponding Unique Content Address (CID) in order of XOR distance.

The DHT

Communication between storage nodes is a crucial component of the network as Pensamento Cloud functions as one giant [Kademlia Distributed Hash Table \(DHT\)](#). Storage nodes are expected to prioritize new connections and communications with other nodes in relation to their Cloud Address. Nodes with similar Cloud Addresses are deemed closer in proximity and are

expected to establish strong connections with one another. This approach ensures that local connection decisions contribute to globally optimized message routing, a concept known as Kademlia connectivity.

Determining and prioritizing proximity is vital to a fully decentralized network like Pensamento Cloud. Storage nodes with the closest proximity can form what are referred to as **Neighborhoods** within the network, which is to be comprised of anywhere from 8 to 64 individual nodes. Nodes within these neighborhoods will then join other neighborhoods with nodes sharing the closest proximity to their individual Cloud Address, and so on, establishing a global network. This design ensures that messages intended for nodes with very low proximity to one another can always reach their destination, even if the nodes are not directly connected.



(Source: [Swarm Whitepaper](#))

Chunks

There are two types of data chunks within the Pensamento Cloud network: **Content-Addressed** chunks and **Single Owner** chunks. In short, content-addressed chunks are identified and addressed based on the content they hold. The **chunk address (CID)** is derived from

the hash of the content itself using the SHA-256 hashing algorithm. Single Owner chunks, on the other hand, are identified and addressed based on the public key of the data owner, with the CID being derived from the user's Pensamento ID itself.

Structure

How does this all work together? When users upload data to Pensamento Cloud, the network starts by splitting the data into several smaller chunks. Each chunk must be less than or equal to a total of 4KB and will be assigned a unique Content Address (CID) based on the content itself and the owner's corresponding [Pensamento ID](#). This CID follows the same address formatting principles as Cloud Addresses, enabling the calculation of the proximity of nodes and chunks.

The data chunks will then undergo a process known as Cauchy-Reed-Solomon (CRS) erasure coding. Each 4KB is to be divided into several sub-chunks, which then utilize CRS to generate 'parity bits.' Each parity bit will contain redundant information about the data within the original data chunk and will be tied directly to it. This process is intended to not only enhance fault tolerance but also ensure data availability. For example, in the event a storage node becomes unavailable, or the data chunk itself becomes lost or damaged, the network can use these parity bits to reconstruct the original data chunk.

After these parity bits are established, each data chunk is hashed together to establish a Merkle Tree and subsequent Merkle Root of the overall file. The CID of the Merkle Root can then be used to enable random access and file integrity verification of the data after the data chunks have been allocated and distributed throughout the network. Pensamento Cloud can also utilize [Manifests](#) to represent collections, allowing it to model a directory tree in a fully decentralized model. This way, Pensamento Cloud can function as a file system, a database, and even host content like websites and dApps, all from one secure platform.

Storage

Following the formation of the Merkle Tree, each data chunk is to be encrypted and distributed to storage nodes within the network. Pensamento Cloud will allocate data chunks to the storage node with the closest proximity to its CID and will require this node to store that data chunk locally. This is referred to as a *Host Node*. For example, let's say we have three storage nodes on the network: **1023**, **2023**, **3023**, and three data chunks ready to be stored: **1011**, **2011**, and **3011**. Using both address types, the network will allocate each data chunk to the storage node with an address closest to the chunk's address.

- Chunk: **1011** will be allocated to Provider: **1023**
- Chunk: **2011** will be allocated to Provider: **2023**
- Chunk: **3011** will be allocated to Provider: **3023**

Once stored, the host node will then share this data chunk with the four nodes in closest proximity to ensure redundancy of the data through a push-pull-sync protocol. This way, even if a node were to go offline or leave the network entirely, the other three nodes would retain a copy of the data. Furthermore, if a data chunk were ever to become corrupt, the network can utilize any of the three copies to ensure data availability and integrity. Figure 4 below illustrates this process where data is uploaded and pushed to the closest neighborhood, where it is continuously synched between each node in the neighborhood itself.




Figure 4: Push, Pull and Retrieve Protocols

(Source: [Swarm Whitepaper](#))

Retrieval

Once the data has been allocated, encrypted, and stored, it can be accessed anytime by the network. When a request is made, a client communicates with a node that is in close proximity to itself. From here, the Merkle Tree of the data is pulled, and requests for each data chunk are sent to the storage nodes with the closest proximity to each chunk. When a request is initialized, the initiating node will send the request to the neighborhood housing the host node. Once the message reaches the host node, it will send the data chunk along this same path back to the initiating node. This way, the network can quickly retrieve data while the initiator and data remain completely anonymous.

However, if any node along this path is also housing the requested data chunk, they can send back the data themselves to reduce bandwidth and computational waste to increase network

performance. Storage nodes are incentivized to store additional data chunks through *opportunistic caching*, where they can receive payment for completing requests for data chunks themselves.

Subsystems

Each storage node within the network is required to create three subsystems within their storage allocation. These are referred to as the *Reserve*, *Cache*, and *Network* subsets.

- **Reserve** - The primary storage subset for storage providers. This is where data chunks that are currently being funded are stored.
- **Cache** - Storage subsystem for data chunks *not* protected by the Reserve. This could be due to data no longer being funded by the owner or the storage node opting to store data chunks not assigned to them as part of an *opportunistic caching* scheme. Cached data is ranked by the latest retrieval time as a means to determine popularity and whether it is worth it to store the chunk. The cache is regularly cleared of unpopular chunks, ensuring that content accessed more commonly is permeated across the network and easily retrievable. For example, if a node is hosting a data chunk belonging to a popular website in its cache, this chunk is less likely to be cleared by the network than a chunk related to a cat photo that hasn't been accessed in a few weeks, maximizing rewards to storage nodes and incentivizing them to store popular data within their cache to reduce network bandwidth and maximize efficiency and performance.
- **Network** - A mandatory allocation of storage to be used by the Pensamento Ecosystem to fulfill needs such as temporary storage for blockchain data, chat data, Pensamento ID data, etc.

With this approach, funded user data is protected by the reserve of its corresponding host node, whereas data not being funded can be added to the node's cache subset to be cleared from the network. Furthermore, nodes can opt to store additional data chunks within their cache to

potentially earn additional rewards by providing data chunks with few message relays. So, the question now is, how is data funded, and how do nodes earn rewards within the network?

Economics

Pensamento Cloud will adopt the *Pensamento Account Protocol (PAP)*, which is an economic model based on the principles of the Swarm Account Protocol (SWAP) introduced in Ethereum Swarm. The goal of PAP is to create a self-sustaining economic model within the network that incentivizes storage nodes to collaborate with one another and decrease computational or bandwidth waste. Opportunistic caching is just one example of the PAP in action, incentivizing nodes to store popular chunks of data for a reward by reducing the length of the message relay. In fact, storage nodes on Pensamento Cloud are incentivized and rewarded in three different ways to offer maximum rewards to storage providers. These include:

- [Storage Incentives](#)
- [Bandwidth Incentives](#)
- [Blockchain Incentives](#)

Bandwidth Incentives

Bandwidth incentives are what make things like opportunistic caching possible. As we know, storage nodes within the network may need to forward requests to other nodes when a request for data is made. The initiating node will send the request to the host node, which then routes the data chunk back through the chain. However, we need a way to incentivize storage nodes within the network to participate in these relays and even work to shorten them through opportunistic caching. This is where Bandwidth incentives come in!

By rewarding nodes for participating and distributing content, the network can incentivize consistent availability and contributions from storage nodes to the overall reliability of the network.

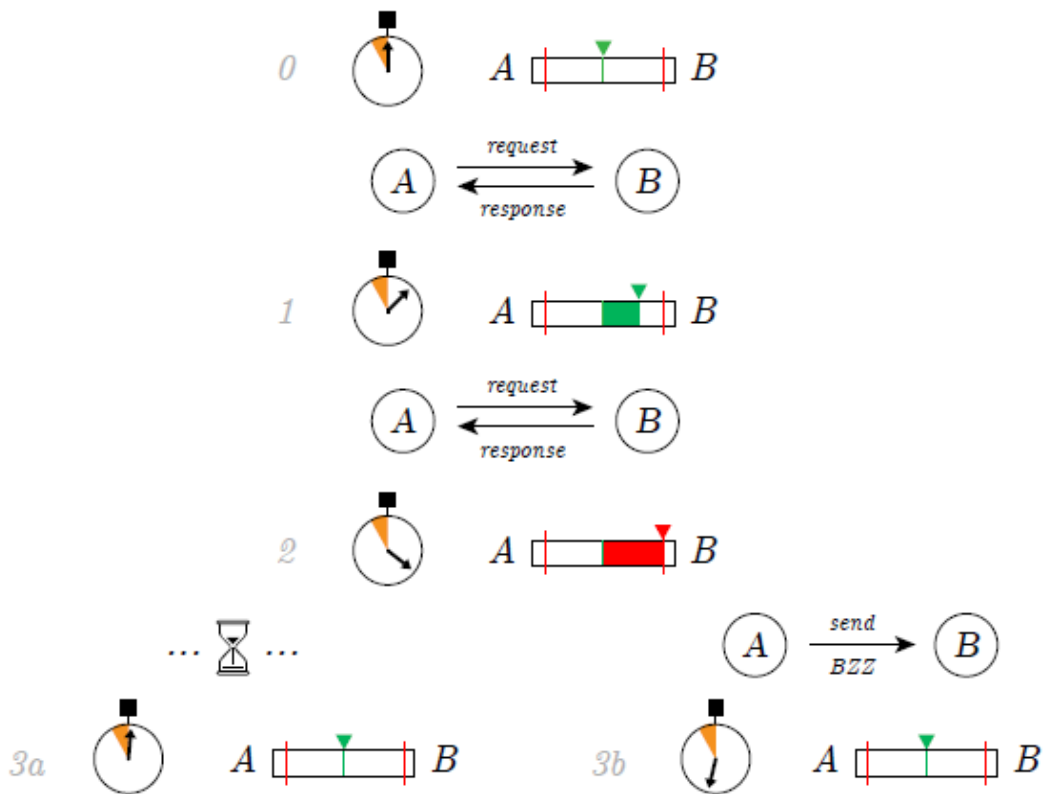
We believe that these incentives will further enhance network scalability by attracting more participants to store and distribute content, completely eliminating the need for a centralized entity.

Another way to think of bandwidth incentives is to imagine a tiny ledger between each storage node within the network. Storage nodes track the relative bandwidth consumption of their connected peers, establishing a decentralized ledger of debts and credits between any two nodes at any given moment. For example, when Node A requests data from Node B and receives a response, Node B accumulates a credit surplus, while Node A incurs debt liabilities. This process continues until a predefined threshold, known as the *debt ceiling*, is reached. Once the debt ceiling has been reached, Node B will refuse any further requests from Node A until it pays its debts to Node B.

To restore balance to the ledger and/or remain within the debt ceiling threshold, Node A can either fulfill new requests from Node B, reducing its debts, or simply pay Node B its debts through a check. If a node is paid with a check, they have the option to either cash the check or hold it in their Pensamento Wallet for use at a later time. These checks are smart contracts run on the Pensamento Blockchain and can be used to pay down the node's own debts, lower transaction fees within the ecosystem, or cash them out for cryptocurrency through Pensamento Swap. However, if a node chooses to hold a check instead of cashing it right away, they do run the risk of the check bouncing if the check provider has insufficient funds in their Pensamento Wallet.

To combat this, bounced checks play a factor in a Pensamento user's reputation within the ecosystem, just as failed validations on the blockchain, slashed funds, etc. If a node writes a check that later bounces, other nodes within the network will see this immutable hit on their reputation and may choose not to communicate with that node in the future. It's worth noting that Bandwidth incentives are meant to be balanced automatically by the system, as Node A will likely make just as many requests as Node B over time. However, these incentives double as a deterrent from nodes overwhelming the network, as nodes who do not participate will end up owing debts to their connected peers, who will ultimately discontinue communications with the node once it reaches its

debt ceiling. This process establishes a reciprocal "service-for-service" relationship between the nodes and aids in creating a self-sustaining economic model within the network. See example below from how this process works in Ethereum Swarm.



(Source: [Swarm Whitepaper](#))

Storage Incentives

As a decentralized storage protocol, Pensamento Cloud also incentivizes storage nodes to, of course, store data! Once again, Pensamento Cloud takes inspiration from Ethereum Swarm and adopts the RACE lottery system for paying storage providers and ensuring compliance within the network. As we know, storage providers in the Pensamento Cloud network have three subsystems: *Reserve*, *Cache*, and *Network*.

The network is designed to automatically remove unfunded and unpopular data chunks to optimize the system and maintain free storage. Bandwidth incentives aim to reward nodes that successfully supply data chunks. However, this process doesn't guarantee the availability of unpopular chunks, meaning that users can't take advantage of the network for long-term data storage. This is where the Reserve subsystem comes in.

Pensamento Cloud users can utilize the network for long-term storage through the use of [tickets](#). When a user initiates a data upload, the selected files will be added together to form what's known as a [batch](#). From here, the network examines the batch's size, compares it to the total available space on the network, and provides a fair market price determined via a [rent oracle](#). Once a price has been established, the user can customize the size of their [ticket roll](#) and estimate how long the batch can be stored based on the ticket price and size of the roll. The goal of the rent oracle is to provide a fair ticket price where one ticket lasts roughly 24 hours at the time of purchase. Once the user selects a price, they can submit the batch to the network. The ticket roll is attached to the batch via its [Ticket Roll ID](#) and signed by the user via their Pensamento ID.

(Please note: Ticket prices change periodically in response to network demand, and price estimates are not real commitments.)


Tickets serve as a way to pre-pay for storage on the network, protecting unpopular data from being removed from the network while incentivizing storage providers to maintain its availability. Storage nodes can also use the price associated with a ticket as a signal to decide which content to store and which to disregard, enabling them to allocate storage space that prioritizes essential chunks. Tickets will gradually depreciate over time, and users can either refill them to retain their data on the network or let them run out and have their data moved to cache storage.

Lottery

These tickets are not used to pay the storage provider directly. Instead, funds acquired from the ticketing system are pooled together to form the [lottery](#). The lottery is used to issue probabilistic

payments to storage nodes, simulating actual payments for each data chunk without the need for individual transactions. The lottery operates on a VM blockchain and follows a protocol consisting of three phases: *pre-committal*, *submission*, and *refutation*.

- **Pre-committal:** During the *pre-committal* phase, storage nodes engage in pre-committal by submitting current price offers for storing specific data chunks. These offers serve to establish a pool of applicants for the subsequent lottery process.
- **Submission:** During the *submission* phase, storage nodes compile lists of data chunks along with proofs of custody corresponding to the data requested by the lottery mechanism. This requested data is referred to as a *witness batch* and is chosen for each applicant, enabling nodes to demonstrate active storage of the assigned data. Nodes failing to prove data availability for the requested data in the witness batch are disqualified from participating in the lottery, incentivizing nodes to retain data assigned to them.
- **Refutation:** Lastly, storage nodes address challenges by submitting refutations during the *refutation* phase. Following this, a set of winners is determined and rewarded. The lottery serves as a way to reward storage nodes for actively participating in the network by proving data availability of randomly requested data allocated to them. This timeline is known as *RACE* (Raffle-Apply-Claim-Earn) and is used throughout each lottery round.



(Source: [The Book of Swarm](#))

The reward payout depends on a few factors, such as the total number of winners, data chunks, ticket prices (lower prices are better), and the proportion of a storage node's personal [stake](#) in the network. To calculate the expected reward amount, simply multiply the total storage rent paid between rounds by the storage provider's stake and divide by the total stake.

Competitive Insurance

The lottery also employs a negative incentive system called 'competitive insurance.' Competitive insurance within the network mandates every storage node to uphold every storage commitment. For example, if a host node is found not storing a data chunk that was allocated to it, the competitive insurance system imposes consequences, possibly leading to the slashing of the node's stake, ensuring accountability and reliability in long-term data storage commitments.

In the competitive insurance system, nodes engaging in long-term storage are required to provide a verified [stake](#) locked by a smart contract on the Pensamento Blockchain. If a node fails to prove ownership of allocated data during the commitment phase of the lottery, it risks the slashing of its entire stake in the network and being barred from future participation. Users or nodes encountering inaccessible data from a host node can also initiate a challenge through a smart contract, triggering the verification process and possible seizure of stake if a host node is found to be inactive.

Blockchain Incentives

Pensamento Cloud also utilizes the Pensamento Blockchain to enable some unique, industry-first features. One of these features includes the option for permanent data storage, or data pinning, through the Pensamento Blockchain. Users can pay a one-time fee to take advantage of permanent storage to have the data saved directly to the blockchain, ensuring immutability. This method does not require any effort by storage providers and is solely powered by the Pensamento Blockchain. However, storage providers will be issued a reward for passing the request along to the blockchain for validation. When a user chooses to upload data permanently, they will first select the

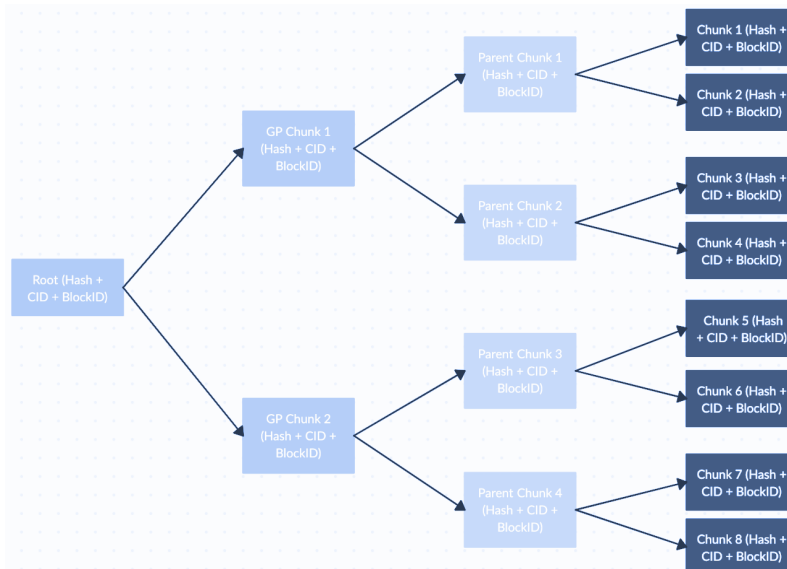
data they wish to store, and the system will calculate its total size, cost, and estimated processing time.

PPSF and PPSR

Note: The following concept is a hypothesis and requires further research, design, and testing.

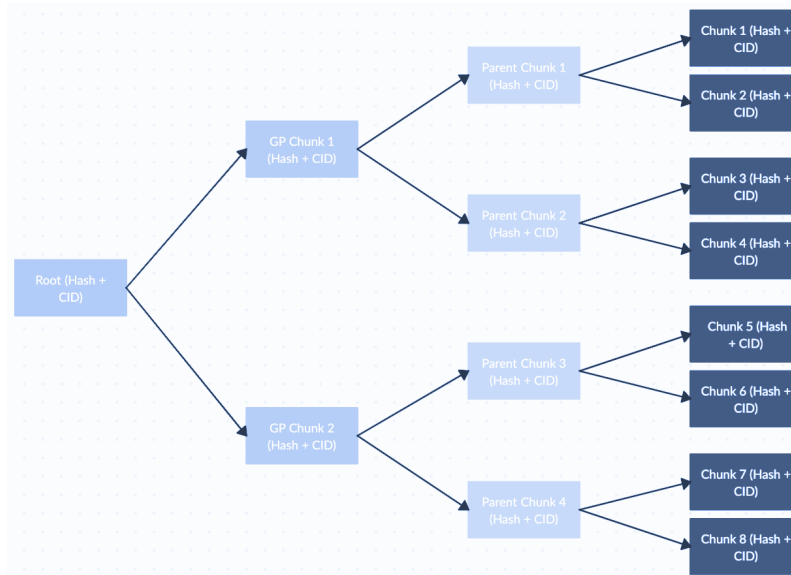
Once the data has been uploaded to the network, it will follow a similar process as standard data. It will be split into several 4KB chunks to establish a Merkle tree and corresponding Merkle root to be sent to the Pensamento Permanent Storage Facilitator (PPSF) for processing to the blockchain.

The PPSF serves as a real-time database construct by adding the Merkle Tree and facilitating the allocation of associated chunks to new blocks on the Pensamento Blockchain—starting from the bottom of the tree and working its way up to the root. As chunks are added to new blocks, the corresponding block information is then tied to the chunks CID within the Merkle tree. This process continues until the system reaches the Merkle root, at which time the Merkle root and tree, along with its updated data, are added as a block on the blockchain.



Once all data chunks have been uploaded to the network and their information is updated within the Merkle tree, the system will add the entire tree and subsequent Merkle root to the final

block. This way, when a user wishes to access the data, the network will only need to call the root block containing the Merkle Root. The system can use each chunk's corresponding hash, CID, and block information to find them in their respective blocks across the blockchain. The Pensamento Permanent Storage Receiver (PPSR) system will handle the retrieval process by using the information found in the Merkle Root to locate each chunk from their respective blocks to recreate the original file.



Permanent storage offers users several benefits, including immutability, decentralization, data integrity, and a one-time fee. Permanent storage opens the door to several use cases, such as financial records, supply chain data, health data, intellectual property, land deeds and titles, notarization and timestamping, etc. However, permanent storage is a more expensive option for users and presents possible disadvantages, including scalability and limited storage capacity. For example, current plans include the Pensamento Blockchain block size being capped at roughly 3,500,000 bytes, or 0.0035 GB per block. This limitation may result in higher costs for the publication and slower data retrieval times depending on the total size of the batch being uploaded.

Staking

Pensamento Cloud storage nodes are required to stake a certain amount of the native \$CLOUD cryptocurrency prior to participating in the network. Similar to the Pensamento Blockchain, Staking serves as both an incentive and a deterrent in the network. As mentioned, storage nodes must stake the minimum amount of \$CLOUD cryptocurrency into the network to be eligible for participation in the lottery and other incentive programs. When a raffle is initiated, selected nodes must put up a stake to participate, and the total reward is split among winning nodes in proportion to their stake in the network. However, If nodes are found to be acting maliciously, behaving dishonestly, or failing to prove ownership of allocated data, a portion or the entirety of their stake can be frozen or slashed from the network and barred from future participation. Users or nodes within the network can also initiate a challenge through a smart contract, triggering the verification process and possible seizure of stake if a host node is found to be inactive.

Staking also offers other benefits to the network. For example, by requiring nodes to authenticate themselves via their Pensamento ID, reputations, and staking requirements, we can better ensure the network remains safe and secure from bad actors. Staking also allows storage providers to increase their position in the network to earn more rewards. Since rewards are distributed to raffle winners in proportion to their stake in the network, nodes can increase their stake to earn higher rewards while increasing their total available storage to increase their chances of being chosen for a raffle.

Network Storage

Lastly, Pensamento Cloud storage nodes are required to allocate 25GB of storage space to their Network subsystem prior to participating in the network. This storage space is to be used to fulfill network needs such as temporary blockchain storage, chat data, Swap data, Pensamento ID data, etc. Network storage is vital to the Pensamento Ecosystem, and as the computer of the ecosystem, we rely on Pensamento Cloud to realize our vision of Web4.

High-Level API Access

Pensamento Cloud's High-Level API Access layer will act as the bridge between the network infrastructure and user applications. Developers building for the Pensamento Cloud network can utilize these sets of functions, protocols, and tools provided by the Pensamento team to create incredible dApps with smooth integration and interaction with the Pensamento Ecosystem.

The API layer will offer developers a collection of native APIs (Application Programming Interfaces) and SDKs (Software Development Kits) to give developers the tools and documentation they need to build applications on top of the Pensamento Cloud network successfully. This layer bridges the underlying infrastructure and the end-user by allowing API access to Pensamento Cloud's mechanisms, like content-based addressing, where developers can locate and retrieve specific data chunks or files using their unique identifiers or hashes.

As a top priority for the Pensamento Ecosystem, the API stack will also incorporate authentication and authorization mechanisms to control network access and resources to ensure security. For example, features like user authentication, role-based access control, and secure communication protocols can be utilized to ensure that only authorized users can interact with the network and perform specific operations.

Lastly, The API stack will also include mechanisms for error handling, logging, and monitoring for developers. By providing developers with the tools necessary to track and handle errors, capture logs for debugging purposes, and monitor the performance and health of the network, the API can help ensure the reliability and stability of applications built on top of the network. By providing a standardized interface and making it easier to build applications, access data, integrate with smart contracts, and manage resources on the network, the API stack truly enhances Pensamento Cloud and developer productivity, promotes interoperability, and enables the creation of a rich ecosystem of applications and services around the network.

Application

Where the API layer allows developers to interact with the core infrastructure of the Pensamento Cloud network, the top Application layer is used by developers to create the end-user experience. The application layer provides a user-facing interface and facilitates the development and deployment of dApps built on top of the network by focusing on allowing end-users to easily interact with the network, access services, and utilize the decentralized storage and computing capabilities.

First and foremost, the application layer is responsible for designing and implementing the user interfaces (UI) and user experiences (UX) that allow individuals to interact with the network and its decentralized applications. We plan to place an extra focus on this layer to develop intuitive interfaces for tasks like uploading and retrieving data, managing storage preferences, configuring privacy settings, and interacting with smart contracts that will provide users with the best possible, streamlined user experience. Our goal has always been to provide a seamless, secure, and user-friendly experience that abstracts the complexities of blockchain—the application layer is where we can accomplish this goal.

Second, the application layer supports developing and deploying dApps on the Pensamento network. We will provide a robust set of developer tools, frameworks, and libraries to enable them to build the best applications in the world by leveraging Pensamento Cloud's decentralized storage, content addressing, and smart contract capabilities.


Lastly, Pensamento Cloud's application layer will facilitate integration and interoperability between the Pensamento Ecosystem. It defines and promotes standard protocols, APIs, and data formats that enable seamless communication and interaction between various network components, as well as different applications, systems, integration, and services worldwide. This promotes the creation of a vibrant ecosystem that extends far beyond the core Pensamento Ecosystem, where

different dApps, services, and users can interact, share data, and collaborate in a decentralized manner, truly providing the foundation for our Web4 future.

Storage Tiers

Pensamento Cloud will provide more options to users looking to adopt the network.


- Standard Storage (No time commitment, no monthly payment, variable pricing)
- Pro Storage (Included with Pensamento Pro subscriptions)
- Permanent Storage (One-time payment, saved forever)
- Corporate Storage (Negotiate storage capacity with recurring monthly cost)



pCloud

Pay-as-you-go


- No Commitment
- Self-Sustaining Economic Model
- Ticketing System
- 1 Ticket = 1 Day
- Prices determined by Oracle
- Re-up your ticket count anytime
- Available to all users



pCloud Plus

Pay-per-month

- Set amount of storage for monthly fee
- Cheaper than traditional cloud storage options
- Great for users looking for traditional cloud experience
- Available to all users



pCloud Pro

Permanent Storage

- One-time fee
- Saved directly to pChain or pChain Plus
- Fully Immutable
- Accessible anytime, anywhere, forever
- Available to all users



Pensamento Swap

Pensamento Swap, or 'pSwap,' is a fully decentralized digital asset exchange built directly into the Pensamento Ecosystem. Our goal with Pensamento Swap is to provide a 'Centralized' experience on a fully decentralized platform, bridging the gap between Web2 and Web3 trading platforms to make it as safe and simple as possible. To achieve this, Pensamento Swap was broken down into three core products: **Swap**, **Swap Plus**, and **Swap Pro**.

Regulation

First and foremost, Pensamento Swap is designed with compliance in mind. We plan to implement strict Anti-Money Laundering (AML) and Know Your Customer (KYC procedures and programs), along with Risk-Based Procedures and others. We aim to ensure the Pensamento platform is not only safe for users but the world as a whole, and we are prepared to offer a trading platform compliant with economies and governments all around the world to restore trust in decentralized systems. With that said, let's take a look at the Pensamento Swap platform.

pSwap

Pensamento Swap is a safe, secure, and fully decentralized digital asset exchange that allows users to buy, sell, and trade cryptocurrencies, NFTs, and more directly to one another. Built as a native solution to the Pensamento Ecosystem, Pensamento Swap functions through a combination of Smart Contracts, the Pensamento Blockchain, Pensamento ID, and other dApps to create a seamless user experience from anywhere within the Ecosystem.

The core focus of Pensamento Swap is to create a 'Centralized' experience on a fully decentralized platform. To achieve this, we plan to offer increased functionality and safety to all

users. This includes things like Biometric and 2-factor authentication, Two-way transaction approval, the ability to set transactions as public or private natively, the ability to set preset trading approval windows for quick trades and immersive experiences on dApps, and even account delegation to set restrictions on who, when, and where a user can trade with a Pensamento ID.

Pensamento Swap is built with security in mind and will ensure that users can safely buy and sell assets from the exchange without fear of scams or rug pulls. To accomplish this, all projects within Pensamento Swap are vetted by Pensamento Team members and must meet or exceed our high-security standards to be added. Once assets are approved and added to the exchange, project teams will be responsible for maintaining project security of their assets, liquidity pools, and much more to ensure the safety and security of all Pensamento users.

pSwap Plus

pSwap Plus takes the Swap platform to a new level, offering users a truly decentralized experience and opening the door to any DeFi asset available without a vetting process. Intended for our professional users, pSwap Plus offers our users the freedom to take their security into their own hands and lifts any restrictions with access to thousands of assets across multiple blockchains while maintaining native support throughout the Pensamento Ecosystem.

pSwap Pro

Lastly, pSwap Pro is a live trading platform built for our most professional users. Built on top of pSwap Plus, pSwap Pro offers pro tools designed for day trading directly on top of the pSwap Plus platform. Users can take advantage of pSwap Pro for a monthly fee or with the purchase of specific Pensamento Pro Plans. All trades are fully supported by the Pensamento Ecosystem, and Pensamento Swap will be a game changer for users looking for a safe, secure, and fully decentralized trading platform without sacrificing a seamless user experience.