TENRAI

Whitepaper

Tenrai Technologies

tenrai-technologies.org

Whitepaper · September 2026

---

1 Introduction

The cryptocurrency industry has evolved considerably, yet several fundamental problems remain unresolved. Scalability, decentralization and security continue to impose competing requirements, and despite numerous claims of having solved the blockchain trilemma, in practice these trade-offs remain. At the same time, the environment in which decentralized networks operate has changed. Mining hardware has become increasingly concentrated, computational resources are available at data-center scale, and smaller Proof-of-Work networks can face substantial security and economic pressure from large miners, rapidly migrating hash power and industrial-scale infrastructure.

This is particularly relevant to CPU- and GPU-oriented Proof-of-Work. Commodity hardware makes mining widely accessible, but it also means that large external pools of computing power can potentially be redirected toward comparatively small networks. For such networks, relying on hash power alone as the final security mechanism can become increasingly difficult. ASIC-based Proof-of-Work reduces some forms of opportunistic hardware migration, but introduces its own concentration dynamics and does not solve the economic problem of miners who must continuously sell part of their rewards to cover hardware, electricity and operational expenses. Large established networks may be able to absorb this pressure; smaller ecosystems often cannot.

There is also a broader question of resource efficiency. Conventional Proof-of-Work expends substantial computation primarily to establish consensus, while competing blocks, stale work and orphaned branches can cause already-performed computation to contribute little or nothing to the final ledger. From our perspective, there are two possible directions: accept these limitations as unavoidable properties of Proof-of-Work, or redesign the architecture so that mining no longer carries responsibilities it does not need to carry.

Tenrai was created around the second approach.

At the core of Tenrai is **TRAIDAG**, a consensus architecture developed specifically to separate open Proof-of-Work block production from ledger ordering and finality. Miners remain permissionless and can produce blocks concurrently, but the security of the canonical ledger does not depend on a longest-chain or work-derived fork-choice race. Timely registered work is not discarded simply because another miner produced a competing block. This allows Proof-of-Work to remain accessible to CPU and GPU participants while substantially reducing wasted work and removing hash power as the sole mechanism responsible for final consensus.

This separation enables a high degree of parallelism. Tenrai's release profile targets **25 work blocks per second and approximately 250 transactions per second**, while the architecture is designed to scale toward **100 BPS and approximately 1,000 TPS** as networking, validation and implementation capacity are expanded. TRAIDAG

explicitly includes a coordinated 100-BPS operating profile in its protocol design. Transactions can additionally propagate through a millisecond-scale Fast Path before their later consensus finalization. The release profile targets deterministic transaction finality in approximately **5–7.5 seconds**, while further development targets a shorter **2.5-second finalization path**. Unlike probabilistic confirmation, finalization represents a committed state transition rather than a confidence estimate that increases with additional blocks.

Proof-of-Work itself is also intended to evolve. Conventional mining remains the initial open production mechanism, but Tenrai is architected with **Proof of Useful Work** as a long-term direction. The objective is to move part of the computational expenditure away from work whose only purpose is satisfying a mining target and toward verifiable computational jobs that provide external utility. CPU, GPU and eventually specialized hardware can therefore remain valuable participants in the ecosystem without requiring the network to treat raw hash power as its primary finality mechanism.

Post-quantum security represents another fundamental design requirement. Retrofitting quantum-resistant cryptography into an existing blockchain is not a simple replacement of one signature algorithm. Post-quantum public keys and signatures are substantially larger, changing transaction sizes, bandwidth requirements, storage assumptions and achievable throughput. Wallet derivation, public-key handling, network authentication and consensus certificates must also be reconsidered. Google has set **2029** as the target for its own full post-quantum cryptography readiness, while Google Quantum AI published updated 2026 research showing significantly reduced resource estimates for future attacks against the elliptic-curve cryptography used by many cryptocurrencies.

Tenrai therefore begins with post-quantum cryptography as a native design requirement rather than treating it as a future migration problem. Transaction authorisation and consensus signatures use ML-DSA-65, consensus commitments use SHA-384, wallet spend locks use SHAKE256-derived public-key commitments, and post-quantum key establishment is incorporated where required by the protocol. This means that throughput, block structure, transaction sizes and network resource budgets are designed around post-quantum cryptography from the beginning rather than being recalculated after deployment.

The economic model follows a similar philosophy. Tenrai deliberately uses a relatively fast declining emission schedule, reducing protocol rewards by **0.3% every 34,560 cuts**, approximately once per day at the release cadence. Mining and staking receive equal portions of subsidy and transaction fees. The objective is not to make long-term network participation permanently dependent on newly issued coins. As issuance declines, staking becomes relatively more important to the economic structure, while transaction fees and future Proof-of-Useful-Work revenue are intended to become increasingly important incentives for miners. This may also reduce persistent emission-driven sell pressure, although the actual market effect can only be evaluated empirically after the network operates under real economic conditions.

Privacy is another area where Tenrai is designed to provide users with a choice rather than forcing one transaction model on the entire network. The network is intended to support **optional Zero-Knowledge confidential and stealth-style transactions**, allowing users to conceal sensitive payment information while ordinary transparent transactions remain available. The privacy system is separately resource-budgeted so that the additional computational cost of Zero-Knowledge proofs does not simply consume the normal work-block budget or force the base network to abandon its high-BPS design.

Tenrai is also intended to evolve into a broader programmable ecosystem. Future development includes **Smart Contracts**, allowing applications such as tokens, NFTs, decentralized exchanges, launch platforms, escrow systems, games and other decentralized services to operate on top of the same deterministic consensus foundation. Further development is also directed toward higher BPS, increased transaction throughput and shorter finalization times.

Tenrai has been developed from the ground up rather than created as a fork of an existing blockchain architecture. The node, consensus implementation and supporting infrastructure are written primarily in modern Rust, while TRAIDAG follows its own architectural approach to parallel Proof-of-Work production, certified ordering and deterministic execution. The objective is not to optimize a single metric in isolation, but to combine high throughput, deterministic finality, accessible mining, post-quantum security, optional privacy and a path toward useful decentralized computation within one coherent system.

Contents

1 Introduction 1

2 TRAIDAG 4

3 TenQ Proof-of-Work 4

4 Quantum Search Resistance 5

5 Post-Quantum Security 6

6 Zero-Knowledge Stealth Transactions 6

7 Deflation and Block Rewards 7

8 Pre-Mining, Fair Launch and Genesis Allocation 8

9 Future Roadmap 8

2 TRAIDAG

TRAIDAG is Tenrai's consensus architecture for highly parallel block production. Its central design principle is the separation of permissionless Proof-of-Work from ledger ordering and finality. Instead of forcing miners to compete for a single next block, multiple work blocks can be produced concurrently and published into a directed acyclic graph. The graph itself does not determine a winning Proof-of-Work chain. Valid work is admitted through stake-certified registration, after which registered blocks receive a stable position in a deterministic execution queue and are processed exactly once. This removes the classical stale/orphan mechanism in which otherwise valid mining work is discarded simply because another competing Proof-of-Work branch wins the ordering race. Network latency can still influence whether a block reaches registration in time, but a timely registered block is not later discarded by a work-derived fork choice.

By separating block production, registration, ordering and execution, TRAIDAG allows the mining layer to operate with a high degree of parallelism without requiring Proof-of-Work itself to establish the canonical history. Miners remain an open and permissionless source of block production, while the stake-certified consensus layer determines which contributions enter the ledger and in which deterministic order they are executed. This architecture is designed for high work-block rates rather than sequential block production; the specification defines a 25 BPS launch profile and also includes a coordinated 100 BPS rate-test profile, with the relevant registration, execution and accounting parameters scaled accordingly. The result is a system in which increasing parallel block production does not inherently create a larger competing-branch problem, because concurrent valid work does not need to defeat other work in a chain-selection race before it can contribute to the ledger.

3 TenQ Proof-of-Work

TenQ is a GPU-first Proof-of-Work function developed for the high-throughput work layer of TRAIDAG. It is designed so that general-purpose CPUs remain capable of mining and verification, while GPUs provide the natural performance advantage for production mining. Unlike a conventional hash-only Proof-of-Work function, TenQ combines cryptographic hashing with a deterministic, stateful and memory-bound execution process. The production profile uses a 32 MiB epoch cache, a 2 GiB read-only dataset, an 8 KiB private scratchpad for each attempt, 32 logical execution lanes and four dependent 256-instruction programs executed across 1,024 rounds. The final state is reduced and committed through SHAKE256 to produce the 256-bit work value.

A conventional hash such as SHA-256 follows essentially the same fixed sequence of operations for every input. TenQ deliberately does not. Each mining attempt creates its own register and scratchpad state, performs data-dependent dataset accesses and executes a broad mixture of integer operations, including addition, subtraction, XOR, rotations, multiplication, multiplication-high, division, remainder, byte swapping, comparisons and shifts. More importantly, the four execution phases are chained: the program used by a later phase depends on the terminal state produced by the phase before it. The complete execution path therefore cannot be determined independently in advance, but develops as the attempt itself is evaluated.

In simplified form, a conventional hash may be viewed as:

Input → Fixed Hash Rounds → Output

whereas TenQ is closer to:

Header → Attempt Seed → Dataset + Private State → Dependent Programs
→ 1,024 Rounds → Final Reduction → SHAKE256

This structure is intentional. TenQ is not designed around the claim that ASICs or FPGAs can be made impossible. Any public deterministic Proof-of-Work function can ultimately be implemented in specialized hardware. Instead, TenQ targets economic specialization resistance. A competitive specialized implementation must reproduce not only arithmetic throughput, but also a large shared memory hierarchy, irregular memory access, attempt-local writable state, mixed integer execution and dependent control flow. The objective is to reduce the cost-per-proof advantage that can be gained from a narrow fixed-function mining device and to make the development of mining-specific ASIC or FPGA hardware substantially more complex and less economically attractive than it is for traditional fixed-pipeline Proof-of-Work algorithms.

The decision not to pursue absolute exclusion of specialized hardware is also deliberate in the wider Tenrai architecture. Tenrai is intended to support Proof-of-Useful-Work in addition to conventional nonce-based mining, with different classes of computational jobs and potentially very different hardware requirements. In such an environment, specialization is not inherently undesirable. GPUs, FPGAs and eventually application-specific accelerators may be highly efficient for particular useful workloads. TenQ therefore seeks to limit the ability of specialized hardware to dominate ordinary mining purely through fixed-function optimization, without designing the network around the assumption that specialized hardware itself must always be excluded. The current TenQ interface deliberately retains nonce, extranonce and target-based mining, while future useful-compute workloads can introduce their own job formats, result-verification rules and settlement mechanisms.

This balance is also important for high-rate operation. Full evaluation keeps the complete dataset resident for efficient repeated work, while Light evaluation reconstructs requested dataset nodes from the much smaller epoch cache. Both paths implement exactly the same consensus function and therefore produce identical work values. This allows TenQ to combine a relatively demanding mining workload with deterministic verification suitable for TRAIDAG's high work-block rates, including the 100 BPS verification model considered in the production profile.

4 Quantum Search Resistance

TenQ does not claim to eliminate the generic advantage provided by Grover's algorithm. Its output is 256 bits, so an idealized search for a specific 256-bit preimage would reduce from roughly 2^{256} classical evaluations to roughly 2^{128} quantum oracle evaluations. Proof-of-Work mining is different, however, because miners do not search for one predetermined output. They search for any result satisfying the current target:

$$V \leq T$$

If p is the probability that one complete TenQ attempt satisfies the target, the classical search cost scales approximately as:

$$C_{\text{classical}} \approx \frac{c_C}{p}$$

while ideal quantum amplitude amplification scales approximately as:

$$C_{\text{quantum}} \approx \frac{c_Q}{\sqrt{p}}$$

The corresponding comparison is therefore:

$$\frac{C_{\text{classical}}}{C_{\text{quantum}}} \approx \frac{c_C}{c_Q\sqrt{p}}$$

where c_C represents the cost of one optimized classical attempt and c_Q the cost of one complete quantum amplification step.

The relevant property of TenQ is therefore not the removal of Grover's square-root advantage, but the cost of implementing c_Q. For a simple hash-only Proof-of-Work algorithm, the quantum oracle can conceptually be reduced to a reversible implementation of the hash calculation and target comparison. A TenQ oracle must instead reproduce the complete stateful computation: dependent program execution, dataset or cache access, scratchpad and register updates, branch state, final reduction, target comparison and the uncomputation required by a reversible implementation. As a result, the practical quantum cost cannot be inferred from the 256-bit output alone; it depends on the resources required to implement the complete TenQ state machine as a fault-tolerant quantum circuit.

TenQ should therefore be characterized as quantum-hardened rather than Grover-proof. It retains the generic quadratic quantum search advantage, while deliberately making each oracle evaluation substantially more complex than a conventional hash-only Proof-of-Work operation.

This additional protection is particularly relevant because Proof-of-Work is not the sole security boundary of Tenrai. TRAIDAG separates permissionless work production from stake-certified ledger ordering and finality. In the intended Tenrai design, that finality layer is protected independently by post-quantum-secure authentication. TenQ's resistance to practical quantum mining should therefore be understood as an additional layer of defense rather than the mechanism on which the network's entire post-quantum security depends.

5 Post-Quantum Security

Tenrai is designed with post-quantum security across the complete ownership and consensus path rather than adding quantum-resistant cryptography only at a single layer. Wallet ownership originates from high-entropy seed material; the deterministic wallet profile uses an effective 256-bit seed and domain-separated HKDF-SHA384 derivation so that independent cryptographic keys can be generated from the same recovery source without reusing key material between different purposes. A 256-bit seed also retains a large brute-force security margin under generic quantum search, where the ideal Grover bound reduces exhaustive search to approximately 128 bits rather than exposing it to the complete break that Shor's algorithm causes for conventional elliptic-curve public-key cryptography. The confidential-payment profile extends this separation further by deriving ML-DSA spending keys and ML-KEM recovery/encryption keys through independent derivation domains.

Ordinary wallet ownership and transaction authorisation use **ML-DSA-65**, standardised in FIPS 204. Tenrai therefore does not depend on ECDSA, Ed25519 or another discrete-logarithm-based signature scheme for control of funds. The corresponding ML-DSA public key is not stored directly as the ordinary output address lock. Instead, the spend lock contains a compact 40-byte commitment derived from the public key using domain-separated SHAKE256. When an output is spent, the transaction witness provides the ML-DSA-65 public key and signature, allowing every validating node to verify that the revealed public key matches the committed lock and that the required spending authority is present.

Transaction signatures are also bound to their exact protocol context rather than signing an ambiguous transaction blob. Tenrai derives a domain-separated SHA-384 spending digest that commits to the chain, protocol version, transaction identifier, input position and the referenced input's amount and spend lock. The resulting digest is signed with ML-DSA-65. This binds the authorisation to the intended transaction and input while separating normal spending signatures from other authorisation roles used by the protocol. The reference transaction witness contains the selected authorisation profile, the 1,952-byte ML-DSA-65 public key and the 3,309-byte ML-DSA-65 signature.

The same post-quantum signature family is used for the staking and consensus layer. Validator consensus keys use **ML-DSA-65**, and consensus votes and quorum certificates are therefore not dependent on classical elliptic-curve signatures. Validator registration includes the consensus public key together with proof of possession, while the protocol provides an explicit key-rotation path for later replacement of that consensus key. Wallet ownership, delegated stake ownership and validator consensus authority remain separate authorisation roles rather than sharing one universal signing key.

Consensus-critical hashing uses domain-separated **SHA-384** with 48-byte outputs. This includes transaction and block identifiers, state roots, queue roots, validator-set commitments, execution identifiers and cut identifiers. SHAKE256 is used where the protocol requires extendable-output hashing, including the public-key spend-lock construction. These hash functions are not vulnerable to Shor's algorithm; quantum attacks against ideal hash functions provide generic search or collision speedups rather than the structural key-recovery attack applicable to classical public-key cryptography.

Post-quantum cryptography is also used outside the ledger itself. Tenrai's reference peer transport combines **X25519 with ML-KEM-768** in a hybrid TLS 1.3 key agreement, while application-level peer identities are authenticated with ML-DSA. For optional confidential-payment functionality, cryptographic responsibilities are separated again: ML-DSA-65 provides spending authority, while ML-KEM-768 is used for recovery and encrypted receiver data, with the corresponding keys derived independently rather than reusing the wallet's signing key. Tenrai's post-quantum model is therefore not based on a single replacement algorithm. **Wallet recovery begins with high-entropy seed material, ML-DSA-65 protects transactions and validator consensus, SHAKE256 binds public keys to spend locks, SHA-384 protects consensus identities and state commitments, ML-KEM-768 provides post-quantum key establishment and encrypted recovery functionality, and independent derivation domains prevent unrelated cryptographic roles from sharing key material.** This layered design is intended to preserve ownership, authorisation and consensus security even in a future environment where cryptographically relevant quantum computers make conventional elliptic-curve signature systems unsafe.

6 Zero-Knowledge Stealth Transactions

Tenrai is designed to support optional zero-knowledge confidential transactions alongside ordinary transparent payments. The privacy layer, defined by TSP-2, allows native TENRAI value to be transferred without publishing the confidential amount, the receiver's long-term payment descriptor or the spending authorisation material in plaintext. Instead, the sender creates a cryptographically bound encrypted output for the recipient, while a zero-knowledge proof

demonstrates that the transaction is authorised, conserves value exactly and creates valid recoverable outputs without revealing the private witness. Receiver-controlled spending uses ML-DSA-65, while encrypted recovery data is built from ML-KEM-768, HKDF-SHA384 and AES-256-GCM. The intended proof family is a transparent, hash-based ZK-STARK, avoiding a trusted setup and keeping the privacy layer aligned with Tenrai's wider post-quantum design. The receiver's long-term descriptor is exchanged outside the public ledger and remains part of the private witness rather than being written directly into each output. Each confidential output contains a randomized commitment together with an encrypted recovery capsule that can be discovered and decrypted by the intended wallet. This provides a stealth-style receiving model while retaining deterministic seed-based recovery: a user returning with the correct wallet seed can recover currently unspent confidential outputs from the authenticated current state without requiring an archive of every historical transaction. Privacy is deliberately separated from transaction-graph hiding. Public input references, output identities, transaction class, counts, fees and timing remain observable, and entry or exit through the transparent value domain can reveal additional information through normal value-conservation relationships. Tenrai therefore does not claim that every confidential transaction is fully anonymous; the objective is to hide sensitive ownership material and confidential values while keeping verification deterministic and the state model bounded.

To prevent privacy from creating permanently growing consensus state, TSP-2 uses a finite set of reusable live-value cells rather than an append-only history of spent-note nullifiers. When a confidential output is consumed, its cell can later be reused with an incremented generation identifier, preventing replay while allowing obsolete state to be reclaimed. Heavy zero-knowledge proofs are also handled through a separately resource-bounded confidential-payment inbox rather than being placed directly into ordinary work-block payloads. This separates Proof-of-Work production rate from proof-verification throughput and allows privacy capacity to be controlled independently of normal block production; the protocol analysis explicitly includes a 100-BPS operating model. Confidential transactions consume more bandwidth, proof-verification capacity and persistent state than ordinary payments and are therefore charged according to their reserved protocol resources. Their minimum fee is derived deterministically from the selected transaction class, reserved bytes, execution mass and newly created state rather than from the hidden transfer amount. Privacy is thus optional and pays for the additional resources it consumes without imposing the same overhead on users who choose standard transparent transactions.

7 Deflation and Block Rewards

Tenrai has a fixed maximum supply of **200,000,000 TENRAI**. The protocol emission budget for mining and staking is **180,000,000 TENRAI** and follows a continuously declining reward schedule rather than traditional periodic halvings.

At launch, the protocol schedules **15.625 TENRAI per cut**. With a target cut interval of **2.5 seconds**, this corresponds to 34,560 cuts per day and an initial daily emission of **540,000 TENRAI**. Each subsidy is divided equally between the two security layers: **50% is allocated to Proof-of-Work and 50% to staking**, resulting initially in 270,000 TENRAI per day for miners and 270,000 TENRAI per day for staking participants. Transaction fees follow the same 50/50 distribution.

The emission rate is reduced by **0.3% every 34,560 cuts**, corresponding to approximately one reduction per day at the target network cadence. Each new reward level therefore equals 99.7% of the previous one. This creates a smooth deflationary emission curve instead of abrupt reward halvings, with an effective subsidy half-life of approximately **230.7 days**.

At the initial target rate of 25 work blocks per second, approximately **62.5 work blocks** are produced during one 2.5-second cut. The initial Proof-of-Work portion of 7.8125 TENRAI per cut is distributed equally among the accepted work blocks, corresponding to approximately **0.125 TENRAI per accepted work block** at target utilisation. The exact reward per work block depends on the number of accepted blocks, while the total scheduled subsidy remains fixed.

This is particularly important for Tenrai's parallel architecture: **higher block production does not increase monetary emission**. A higher work-block rate changes how the available Proof-of-Work reward is divided between accepted blocks, but does not create additional subsidy. Future increases in BPS can therefore increase network throughput without proportionally increasing inflation.

The relatively fast emission decline is intentional. Tenrai is not designed around the assumption that long-term Proof-of-Work participation should remain primarily dependent on newly issued coins. A rapidly declining coinbase reduces the amount of newly created supply entering circulation and may therefore reduce persistent emission-driven sell pressure over time, while making staking and fee-based network participation increasingly relevant. In the long

term, **transaction fees and future Proof-of-Useful-Work rewards are intended to become the primary economic incentives for Proof-of-Work participants**, while the coinbase subsidy mainly serves the early distribution and bootstrapping phase of the network.

This economic model is enabled by TRAIDAG's separation of Proof-of-Work production from ledger ordering and finality. Miners provide open, permissionless work and receive the corresponding work rewards, but Proof-of-Work does not independently determine or finalise the canonical ledger state. Network agreement remains stake-certified, meaning that increasing or concentrating hash power alone does not replace the stake-based agreement threshold or directly provide control over finalised consensus history. This reduces the need to preserve permanently high coinbase rewards solely to maintain the network's primary finality security and allows the Proof-of-Work economy to transition toward transaction activity and useful computation as issuance declines.

The 0.3% reduction schedule is calibrated to the **180,000,000 TENRAI protocol emission budget**. Approximately **119.88 million TENRAI** are emitted during the first year, approximately **159.92 million cumulatively after two years**, and approximately **179.997 million after ten years**. As rewards are represented in integer base units, the remaining terminal emission continues at progressively smaller amounts over a much longer period, with the current schedule extending for approximately **55 years**.

8 Pre-Mining, Fair Launch and Genesis Allocation

Tenrai will launch without pre-mining and with an open mining and staking start available to all participants from the beginning of the network. A reference mining implementation will be released together with the network and will support GPU mining from day one. This is intended to avoid an information or software advantage in which a limited group could privately use an optimized GPU implementation before comparable software is publicly available. The mining software, node software and other official Tenrai software will not include developer fees or protocol-level usage fees. Participants only pay the transaction fees defined by the network, which are distributed according to the protocol's reward rules.

Although there is no pre-mining, 10% of the total supply will be allocated directly through the genesis state. This allocation is separate from mining and is transparently defined before network launch. Six founding developers will each receive 1% of the total supply, resulting in a combined founder allocation of 6%. These allocations compensate the core developers responsible for the research, development and initial implementation of Tenrai and represent their direct long-term allocation rather than a hidden mining advantage.

The remaining 4% of the genesis allocation is reserved for the continued development and operation of the project. It may be used for ongoing protocol and software development, infrastructure and server costs, ecosystem development, marketing, community events and other expenses required to maintain and expand the network. A portion of this reserve will also be used to establish the initial staking base at network launch. Because stake participation contributes directly to TRAIDAG ordering and finality, the network requires sufficient active stake from the beginning rather than relying on stake to accumulate only after launch. The amount deployed for initial staking will therefore be limited to what is considered necessary for network security and will not represent the entire reserve, allowing independent participants to compete meaningfully for staking rewards from the start.

Tenrai therefore distinguishes between pre-mining and transparent genesis allocation. No coins are privately mined before public launch, and no participant receives an undisclosed computational head start. The genesis allocation is instead defined at protocol launch for founders, long-term project funding and initial network security, while the remaining supply enters circulation according to the public mining and staking rules.

9 Future Roadmap

Tenrai is intended to continue evolving beyond its initial network release. One of the major future developments is **Proof of Useful Work**, extending mining beyond conventional computational work toward jobs that can produce externally useful results. The goal is to support different categories of useful computation and different hardware classes, including CPUs, GPUs, FPGAs and specialized accelerators, while retaining an open and competitive compute ecosystem.

Another major development is **Smart Contracts**, expanding Tenrai into a programmable application platform. Future contract functionality is intended to support applications such as tokens, NFTs, decentralized exchanges, launchpads, escrow systems, games and other decentralized services, while maintaining Tenrai's focus on deterministic execution, security and high-throughput operation.

Further improvements to the core network are also part of the long-term development roadmap. This includes research and development toward **higher block rates, increased transaction throughput and faster finalization**, allowing Tenrai to scale beyond its initial operating profile as implementations, networking and hardware performance continue to improve.

The long-term goal is to progressively expand Tenrai from a high-performance decentralized network into a broader computational and application ecosystem, combining scalable transactions, programmable applications and useful decentralized computation.