τensorprox SN91
Decentralized Edge Firewall

Powered by Bittensor Network

Shugo[1]

[1]www.shugo.io

June 2025

Abstract

τensorprox is a Decentralized Infrastructure Collaboration Platform (DICP) built on the Bittensor network [1], designed to aggregate and orchestrate distributed protection systems through a competitive architecture of validator and miner nodes: validators serve as coordinators and controllers, while miners operate self-managed defense engines. Initially focusing on distributed DDoS detection and mitigation, the protocol enables high-capacity, cost-effective protection by dynamically forming adaptive mesh networks that respond in real time to client needs. It introduces novel approaches for traffic management, intelligent node selection, performance evaluation, and resource allocation that overcome key limitations of traditional centralized protection systems. Our three-phase deployment strategy transitions from a testing environment (Foundation Forge) to a commercial platform where clients can access on-demand protection services. By establishing a resilient and scalable base layer, τensorprox aims to evolve beyond DDoS mitigation into a global cybersecurity platform capable of addressing the most complex security challenges and continuously adapting to sophisticated and emerging threats at scale. This paper details the strategic positioning, technical architecture, economic model, deployment roadmap, and governance framework that underpin the τensorprox system for decentralized, adaptive cybersecurity.

Contents

1 Introduction

1.1 The DDoS Threat Landscape

Distributed Denial of Service attacks continue to grow in scale, sophistication, and frequency [3]. Recent industry reports indicate that DDoS attacks reached record volumes exceeding 6.5 Tbps in 2025, with an impressive 509% year-over-year and 397% quarter-over-quarter increase [4]. Leading DDoS protection providers have established comprehensive multilayered defense architectures that leverage global point-of-presence networks, intelligent traffic scrubbing, and real-time threat detection capabilities. These solutions demonstrate impressive scale—with Akamai's cloud scrubbing centers located in 32 global metropolitan cities providing 20+ Tbps capacity [6], and Cloudflare's anycast network delivering 348 Tbps in over 335 cities [4] [5]—yet fundamental architectural limitations persist.

The current state of the cybersecurity market reveals that lack of competition stifles innovation, slowing the development of more effective and adaptive solutions. The prevalence of centralized architectures introduces systemic policy risks and creates concentrated control and decision-making, which can lead to governance issues where political or organizational decisions adversely impact users [7]. More importantly, commercial cybersecurity solutions often come with high costs [8] that are prohibitive for smaller entities, limiting their ability to protect themselves in an increasingly complex threat landscape. These limitations create a protection gap that adversaries continue to exploit, particularly targeting mid-market companies and decentralized applications that lack enterprise-grade protection budgets.

1.2 Strategic Focus on Layer 3/4 Protection

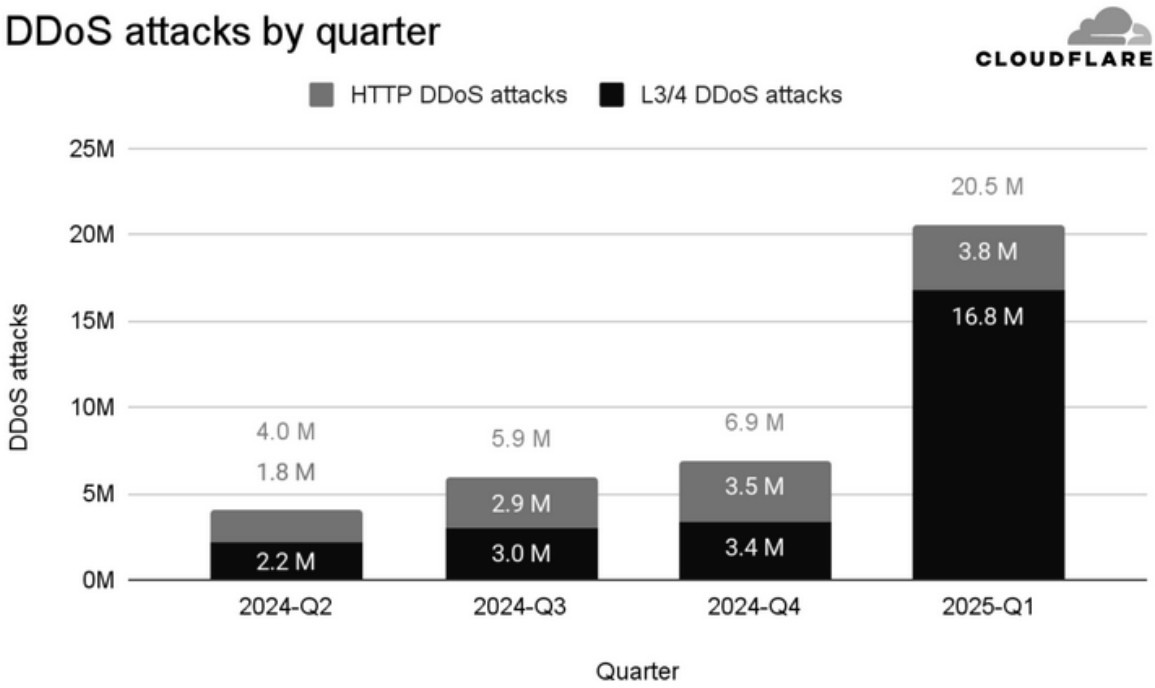

Figure 1: Cloudflare DDoS attack report - Q1 2025 [4]

Analysis of recent attack data shows that Layer 3/4 DDoS attacks constitute 82% of all DDoS attacks, with 16.8 million incidents recorded in Q1 2025 alone. Within this category, SYN flood attacks remain dominant (30.7%), followed by DNS flood attacks (18.5%), Mirai (18.2%) and UDP floods (8.9%)—making them a critical focus area for effective protection [4]. These network layer attacks are particularly dangerous because they can achieve massive volumetric disruption, overwhelming infrastructure before traffic even reaches application layer defenses. By focusing on these L3/4 threats,

τensorprox aims to neutralize attacks closer to their origin and address a significant portion of today's DDoS landscape.

1.3 Decentralized Protection: A New Paradigm

Decentralization transforms cybersecurity by enhancing resource distribution, economic efficiency, and scalability. By spreading resources across independent nodes, τensorprox overcomes the limitations of centralized systems, fostering innovation through open competition. Mesh networks, where security networks and Web3 naturally converge, enable this resilient architecture. Operating without owned infrastructure, the decentralized network eliminates hardware and human resource costs, as contributing nodes manage all operations, significantly reducing overhead.
Built on top of the Bittensor [1] network infrastructure, τensorprox leverages robust peer-to-peer infrastructure, cryptographic identity verification, token-based economic incentives, consensus-driven decision-making and efficient cross-node communication to create a specialized subnet optimized for network security services.
Bittensor subnets support 256 logical node slots by design—each capable of virtually unlimited scaling—ensuring exponential growth to meet rising demand.

2 System Architecture

2.1 Hierarchical Node Structure

2.1.1 Miner

The miner delivers core protection services through a distributed cluster of Moat nodes. These nodes are responsible for implementing a suite of defense mechanisms to detect, validate, and manage network activity. This includes filtering traffic to block malicious or unauthorized packets, validating that all incoming traffic adheres to legitimate protocol specifications, dynamically applying rate limits based on observed traffic patterns, and maintaining efficient connection management to monitor and regulate client interactions.

2.1.2 Validator

Validators serve as the orchestration layer within the network, overseeing traffic capacity and ensuring its efficient distribution across Moat nodes. They are responsible for selecting the most appropriate Miners to protect clients, taking into account factors such as geographic proximity, system performance, node reputation and available resources. To uphold consistent quality of service, Validators perform ongoing security and performance audits. Based on these evaluations, they allocate alpha tokens—the native currency of τensorprox—as rewards to Miners, in alignment with predefined performance metrics.
To participate, Validators must stake a significant amount of TAO (Bittensor's native token) or alpha tokens, ensuring their economic alignment with the integrity of the network.

2.1.3 King/Client

In Phase 1 (Foundation Forge), king nodes serve as test targets; in Phase 3, they represent client assets requiring protection.

2.2 Communication Protocol Stack

τensorprox is structured as a modular stack, delivering multi-level functionalities to support the network's core architecture. The stack components are:

1. Base Layer: Bittensor P2P with cryptographic identity.

2. Traffic Management Layer: Routing and load balancing.
3. Telemetry Layer: Performance and capacity metrics.
4. Security Layer: Encrypted communications and access control.
5. Application Layer: Client interface and protection service APIs.

2.3 System Data Flow Architecture

2.3.1 Core Processing Pipeline

The system processes traffic through a streamlined pipeline optimized for security and performance. Client traffic enters the τensorprox network via a strategically placed entry point and is intelligently routed to an appropriate Miner based on validator assignment rules. The Miner then employs a load-balancing strategy [9] to distribute traffic across its Moat nodes, which apply multi-layered defense mechanisms to filter malicious traffic while seamlessly forwarding legitimate requests to protected assets.

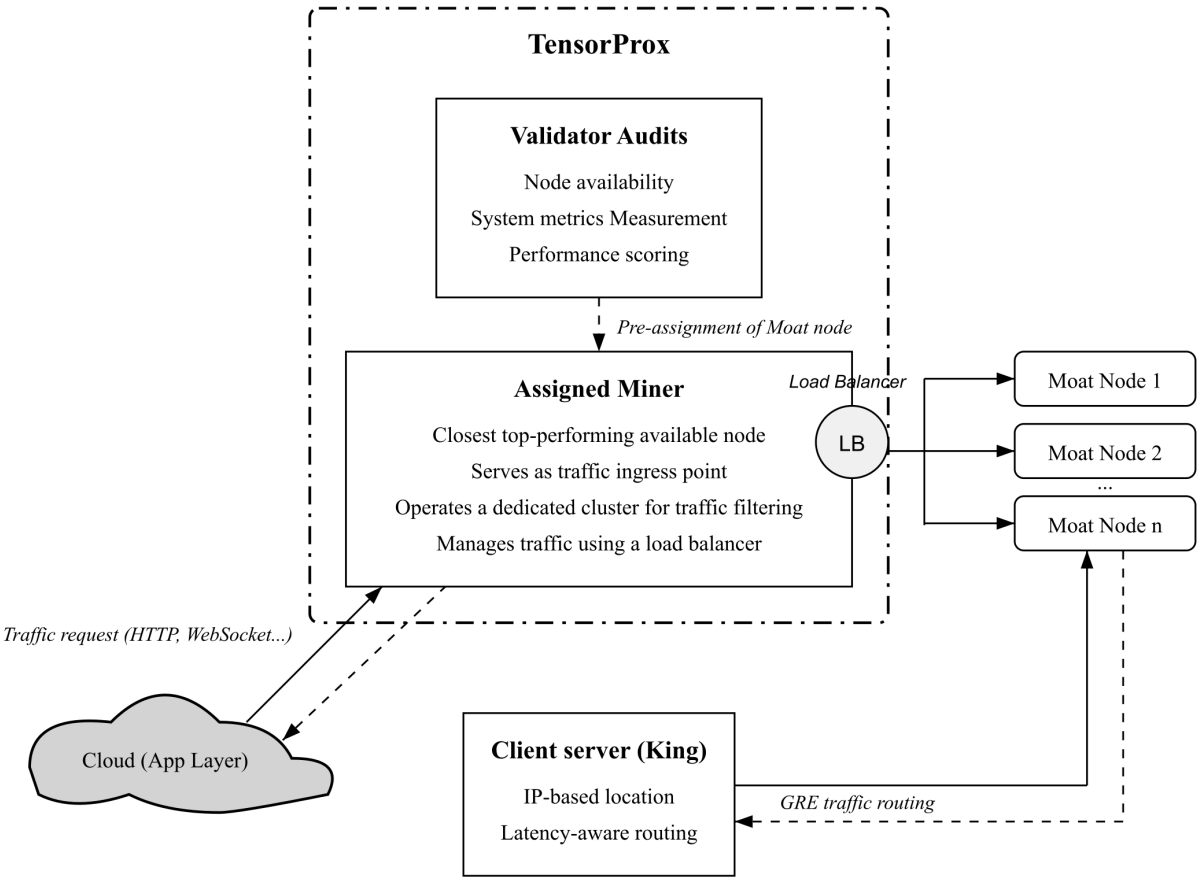

Figure 2: Conceptual Design of the Moat Network Architecture

2.3.2 Quality Assurance

The system employs a robust quality control framework to ensure operational excellence. New nodes undergo a stringent pre-assignment phase, where they must prove efficacy, consistency, and reliability by competing with established nodes before accessing production environments. High-performing nodes

earn "reserved node" status, granting priority assignment to safeguard high-value clients. Validators conduct ongoing performance assessments against predefined benchmarks, seamlessly integrated to avoid disrupting production traffic.

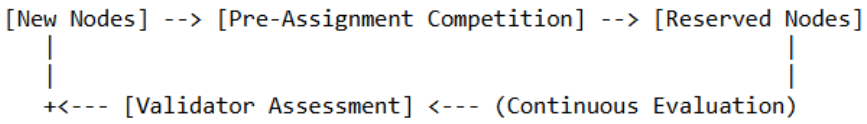

Figure 3: Quality Assurance Loop

2.3.3 Secure Command Execution Framework

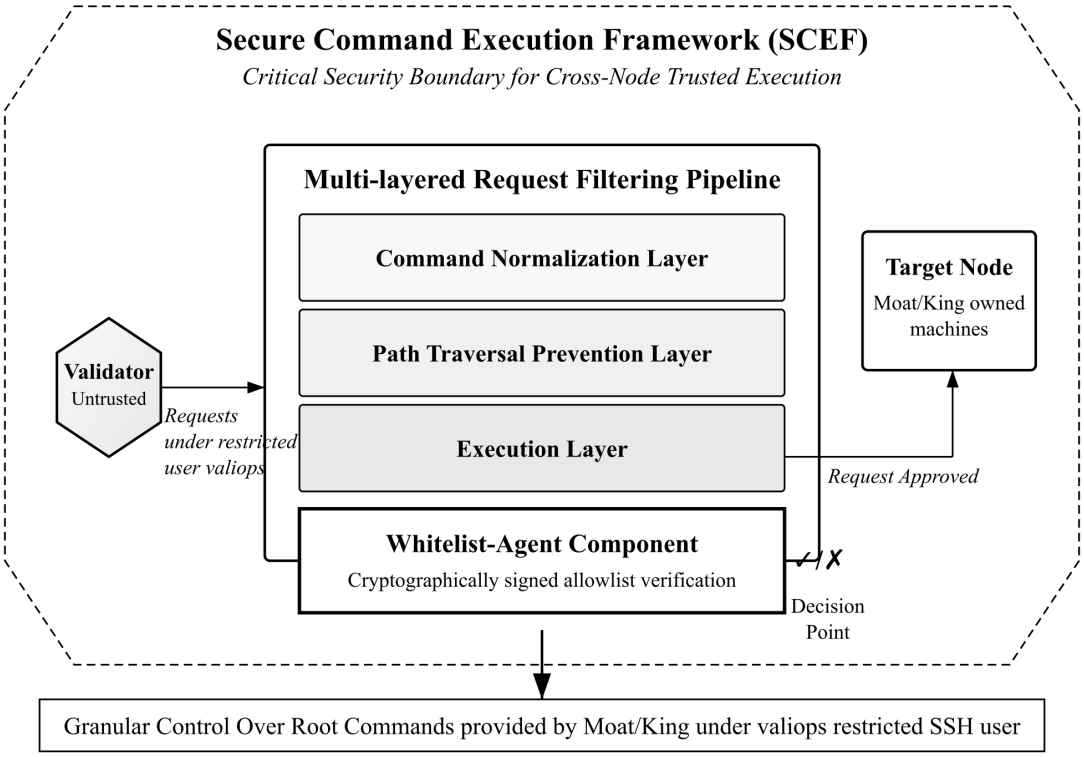

Figure 4: Secure Command Execution Framework

The Secure Command Execution Framework (SCEF) is a critical security boundary that enables secure cross-node auditing within the network by providing granular control over root command execution. Through sophisticated privilege isolation and command restriction mechanisms, SCEF allows untrusted validators to execute a strictly controlled set of diagnostic commands on non-owned external nodes without compromising the integrity of the nodes' system. The framework leverages the whitelist agent component to intercept and control all incoming remote executions. This zero-trust approach ensures that even if a validator is compromised, the potential attack surface remains extremely limited. SCEF's

multilayered validation pipeline, incorporating command normalization and path traversal prevention, provides defense-in-depth while still enabling performance and system metrics measurements essential for the network's quality assurance and adaptability.

3 Traffic Management System

3.1 SmartFlow Engine

The SmartFlow Engine (SFE) serves as the intelligent nerve center of the τensorprox ecosystem, dynamically managing traffic flows between clients and protection resources. It implements real-time decision algorithms that continuously determine the optimal allocation of resources for the network.

3.1.1 Client Assignment Algorithm

Upon client subscription to protection services, the orchestration algorithm intelligently assigns Miner nodes by evaluating a composite set of weighted criteria, including availability, historical performance, current load and Moat system capacity, geographic proximity and specialized defense capabilities. This multifaceted approach ensures an optimal balance of performance, reliability, capacity, low latency, and tailored protection, delivering robust and efficient service allocation.

3.1.2 Near-Field Defensive Clusters

τensorprox's Near-Field Defensive Clusters (NFDCs) strategically deploy defensive resources in close network proximity to critical assets. By positioning protection within a 10-millisecond network latency range, NFDCs ensure rapid response times and localized traffic management, reducing congestion on backbone networks.

3.1.3 Dynamic Traffic Rebalancing

SmartFlow continuously monitors network conditions and performance metrics to automatically rebalance traffic when necessary. It proactively detects load by monitoring the capacity of Moat nodes, quickly shifts traffic away from nodes showing performance degradation, and adapts to emerging attack patterns by redistributing traffic based on evolving threat signatures and the specialization profiles of Miners.

3.2 Traffic Analysis and Telemetry System

This unit serves as forms the backbone of the network's quality control framework. This system is responsible for continuously collecting, analyzing, and acting upon multidimensional metrics across the entire protection mesh, ensuring optimal performance, traffic legitimacy verification, and adaptive resource allocation.

3.2.1 Distributed Telemetry Architecture

The telemetry collection system utilizes a hybrid push-pull architecture to optimize real-time monitoring while maintaining system efficiency. It conducts periodic metrics harvesting at configurable intervals to establish low-overhead baseline data, triggers immediate metric collection in response to detected anomalies, performs random deep inspection sampling of traffic patterns without compromising performance, and correlates metrics across adjacent network segments to identify discrepancies, ensuring robust and accurate network oversight.
This multi-modal telemetry approach ensures comprehensive network visibility while minimizing resource overhead on protected assets. Telemetry data are transmitted through a dedicated, encrypted

channel, isolated from primary traffic paths, with prioritization of critical metrics during network congestion.

3.2.2 System Metrics Collection Framework

The System Metrics Collection Framework (SMCF) provides a unified interface for gathering performance and capacity data across heterogeneous node deployments, enabling accurate resource assessment and optimization. SMCF employs an abstraction layer that normalizes system metrics across diverse hardware environments. This cross-platform approach enables accurate comparison between disparate node configurations, ensuring proper load distribution regardless of underlying hardware differences. The SMCF leverages the Secure Command Execution Framework (Figure 4) to gather comprehensive system-level metrics. This includes access to privileged system statistics through controlled execution paths, safe monitoring of process and system resource consumption, utilization of hardware performance counters via abstracted interfaces, and the collection of in-depth packet processing statistics from kernel-level network subsystems (Table 1).

Table 1: System Metrics Collection Framework

Category	**Parameters**	**Method**	**Normalization**
Computational	CPU utilization, thread count, instruction throughput	SCEF commands with resource limiting	Benchmark scores
Memory	Available RAM, buffer allocation, cache hit rates	Whitelist agent analysis	Virtual capacity units
Network	Interface throughput, packet rate, queue depth	NIC statistics via sysfs paths	1Gbps-equivalent units

3.2.3 ResourceSync

This network inventory system collects a Static Resource Profile (SRP) from Moat nodes and their underlying infrastructure, supporting clustered deployments across trusted hosting providers.
By leveraging standardized API interfaces, Rsync compiles a comprehensive inventory of virtual resources for each Moat node, including:

- **Computing Resources**: Virtual CPU count (e.g., 64 vCPUs), CPU architecture (e.g., x86_64, ARM64), processor model (e.g., AMD EPYC 7763), physical core count (e.g., 32), threads per core (e.g., 2), and specialized instruction sets (e.g., AVX2, AES-NI).
- **Memory Resources**: Total RAM (e.g., 256 GB), RAM type and speed (e.g., DDR4-3200), ECC support, and NUMA configuration (if applicable).
- **Network Resources**: Aggregate bandwidth (e.g., 100 Gbps), network interface card (NIC) configuration, and packet processing capacity.
- **Provider Information**: Hosting provider (e.g., AWS, GCP), geographic region, and instance or server class.
- **Cluster Metrics**: For multi-node deployments, includes node count, aggregate vCPU count, total memory, and combined network capacity.

This structured profile enables accurate capacity planning and Moat selection by providing constant visibility into the static capabilities of protection resources before they are assigned to specific tasks.

3.2.4 DynamicSync Routing Engine

DynamicSync Routing Engine orchestrates client-server assignments across the network by integrating Moat node capacity and telemetry data. It employs advanced client-side adaptation mechanisms to optimize connectivity amid evolving network conditions and defensive resource reallocations.
The protocol dynamically reconfigures client-side GRE tunnels [11] to adapt to shifting network topologies, automatically updates DNS resolution entries to align with changes in protection configurations, and maintains session state continuity during transitions to minimize disruptions, ensuring seamless and robust connectivity for protected clients within the τensorprox ecosystem.

4 Moat Evaluation Framework

4.1 Synthetic Attack Generation

τensorprox implements an extensible library of attack vectors that serves as the input layer for node evaluation. This library encompasses a broad spectrum of contemporary cyber threats with varying complexity and impact profiles:

- **Volumetric Attacks**: High-bandwidth DDoS vectors including UDP floods, amplification attacks (DNS, NTP, SSDP), and ICMP floods that target network capacity exhaustion.
- **Protocol Attacks**: TCP/IP exploitation vectors such as SYN floods, ACK floods, fragmentation attacks, and Ping of Death variants that exploit protocol vulnerabilities.
- **Mixed/Complex Attack Patterns**: Sophisticated multi-vector attacks combining temporal variations, protocol diversity, and evasion techniques to circumvent traditional detection mechanisms.

The library is continuously updated based on emerging threat intelligence and evolving attack methodologies observed in production environments.

4.2 Performance Scoring Matrix

Nodes are evaluated by validators through structured pre-assignment competition rounds and continuous, unpredictable production phase audits to assess operational resilience and consistency (Table 2).

Table 2: Performance Scoring Matrix

Dimension	**Measurement Methodology**
Traffic Purity	Measures how clean the traffic is by calculating the proportion of normal traffic compared to the overall traffic forwarded.
Filtering Accuracy	Benign delivery rate x Attack detection accuracy
Throughput	Capacity/Forward Bandwidth
Latency Impact	Added latency to legitimate traffic (ms)

4.3 Reputation System

The node reputation system incorporates performance history, reliability, audit compliance, age, and stake level, with a temporal decay function:

$$\text{performanceHistory} = \frac{\sum(\text{score}_i \times \text{temporalWeight}(t_i))}{\sum \text{temporalWeight}(t_i)} \tag{1}$$

The reliability factor ensures nodes maintain consistent availability and robust uptime, critical for network stability, while audit compliance evaluates a node's ability to consistently meet rigorous, unpre-

dictable audit standards, reinforcing trust in operational integrity. The system balances recognition of long-established nodes with opportunities for newer entrants, fostering both stability and innovation. The economic commitment of node operators through alpha staked aligns incentives to encourage active and responsible participation, strengthening the network's overall resilience and efficiency.

5 Economic Model

5.1 Alpha Token: Designed for Real-World Utility

Alpha (α), the native utility token of the τensorprox ecosystem, is capped at 21 million tokens and serves as the cornerstone for aligning incentives among network participants. The token incorporates multiple functions:

- **Network Participation**: Validators stake Alpha tokens to actively participate in securing and maintaining the network.
- **Reward Distribution**: Tokens are distributed as compensation to network participants (e.g., miners and validators) for their contributions.
- **Governance Mechanism**: Alpha tokens grant holders weighted voting rights, enabling decentralized decision-making on protocol upgrades and network policies.
- **Service Access / Holder Benefits**: Token holders receive discounted or complimentary access to protection services, revenue sharing for stakers, and other exclusive privileges.

5.2 Stakeholder Incentive Framework

5.2.1 Validator/Miner

Validators earn for orchestration and audits, ensuring network integrity. Miners deliver protection, earning tokens based on overall performance.

5.2.2 Client User

Clients can use the Alpha token as an alternative payment method for protection services, benefiting from streamlined, lower-cost transactions. The system fosters enduring client relationships through progressive loyalty benefits —including discounted services and preferential access to new features—creating a virtuous cycle of engagement and value creation within the ecosystem.

5.2.3 Staking Benefits

The stake-to-access model offers an optimal path for client engagement, where staking Alpha tokens unlocks service discounts proportional to the staked amount. At higher tiers, clients may access protection services at no additional cost beyond their stake. This model aligns client incentives with network growth through strategic token commitment, while also granting governance rights based on stake size. In addition to discounted services, stakers earn a share of the network's revenue from commercial activities, reinforcing their role in sustaining and shaping the ecosystem.

5.3 Economic Security Mechanisms

The Alpha token economy incorporates safeguards against economic attacks. Against Sybil attacks, the system implements stake requirements that raise the cost of attack, performance-based selection that limits impact of low-quality nodes, and a reputation system that favors established, reliable participants.

The long-term viability of the network is ensured through balanced incentive distribution across participant types from token emissions and network fees, strategic treasury allocation for ongoing development, and governance mechanisms for economic parameter adjustment.

5.4 Governance Rights

Alpha tokens grant holders proportional rights in network governance, allowing participation in protocol parameter adjustments, economic model refinements, feature development prioritization, and security policy decisions. This governance utility ensures that all stakeholders have representation proportional to their network commitment.

6 Roadmap

6.1 Phase 1: Foundation Forge

Timeline: *Q1 2025*
Focus: *Core Infrastructure & Competitive Framework*
Status: *Completed*

Key Achievements

- Core subnet architecture deployment including SCEF framework.
- Synthetic attack library for L3/L4 attack vector simulation, volumetric and protocol attack patterns, dynamic attack generation, and real-world threat intelligence integration.
- Dataprox release: an open-source data collection framework leveraging synthetic traffic generation.
- Multi-dimensional scoring system implementation, performance tracking dashboard at dashprox .com.

Foundation Metrics: 250 active miners; 24 distinct attack vectors catalogued.

Table 3: Foundation Metrics (Phase 1)

Metric	Baseline	Target
Ingress Traffic Handled	50 Gbps [10]	250+ Tbps
Normal Traffic Delivered	10 Gbps [10]	200+ Tbps
Traffic Delivery Rate	98% [10]	100%
Purity / Traffic Cleanness	97% [10]	$> 99\%$
Filtering Method	Batch processing & ML	Real-time, AF_XDP & ML
Latency Overhead	1–10s	Sub-10ms

These metrics reflect early-stage performance using a simple decision tree classifier operating on 1-10 seconds packet batches; they serve as a starting point and do not represent the expected final performance. As τensorprox evolves, we are actively transitioning to ultra-low-latency packet filtering using high-performance Linux networking techniques such as AF_XDP [10], targeting sub-10ms end-to-end processing latency.
Combined with more advanced classification models, distributed optimization, and continuous miner-driven learning, our goal is to reach 100% delivery and $> 99\%$ attack filtering with less than 10ms

overhead latency. This shift from batch analysis to real-time streaming intelligence marks a key milestone in achieving scalable, production-grade cybersecurity performance in decentralized environments.

6.2 Phase 2: Maturity

Timeline: *Q2-Q4 2025*
Focus: *Real-Time Filtering, KPI Growth, and Platform Core Implementation*
Status: *In Progress*

This step marks a strategic transition from foundational experimentation to robust performance refinement. The focus shifts toward enhancing the intelligence, responsiveness, and scalability of the τensorprox network, laying technical and algorithmic groundwork for seamless platform integration.

Key Objectives

- Transition from batch to real-time filtering using AF_XDP [10], eBPF, and zero-copy packet processing.
- Expand network ingress capacity from 54 Gbps to multi-terabit throughput through node scaling, better routing overlays, and bandwidth-aware load balancing
- KPI Growth & Benchmarking by improving scoring mechanism & fostering competition between miners
- RSync & SMCF for metrics access, custom rule deployments, and diagnostic tooling, setting the stage for Phase 3's user-facing integration.

The Maturity Phase goal is to prepare τensorprox for broad adoption: this includes ensuring miners can operate efficiently at high loads, validators can enforce rigorous performance standards, and the network can autonomously evolve its threat response strategies. By the end of this phase, the system will be equipped to provide protection services to users, applications, and subnet nodes through a decentralized seamless platform interface.

6.3 Phase 3: Platform Integration

Timeline: *Q1-Q2 2026*
Focus: *On-Demand Protection Services, User Access*
Status: *Upcoming*

During this phase, τensorprox transitions from a research-driven infrastructure to a user-centric cybersecurity-as-a-service platform. The objective is to make decentralized protection accessible, automatic, and actionable for a wide range of users, from individuals to enterprises and subnet operators.

Key Objectives

- Develop lightweight client agents that automatically create GRE tunnels and all preconfigurations needed to deploy a client node to the τensorprox mesh.
- Implement a validator-led matchmaking layer to assign clients to optimal Moat clusters based on real-time performance, geography, load, and availability.
- Deploy an incentive-compatible usage-based billing system where clients pay based on traffic volume and uptime guarantees, miners are rewarded for performance, uptime, and threat prevention accuracy, and validators earn for orchestration and auditing.

- Launch a web-based dashboard and CLI/API suite allowing users to subscribe to protection plans, deploy and manage their own instances, monitor filtering KPIs/traffic logs in real-time, and trigger instant audits or escalations.
- Provide protection to other Bittensor subnets and third-party Web3 infrastructures, allowing decentralized applications to secure their endpoints natively using τensorprox nodes.

Phase 3 delivers on the core promise of τensorprox: decentralized protection as a service. By abstracting the complexities of routing, attack detection, filtering, and mitigation, it will offer any user the ability to deploy enterprise-grade defense in seconds, without relying on centralized intermediaries. This stage redefines how cybersecurity is delivered, moving from passive tools to active, intelligent mesh-based guardianship, powered by collective intelligence and verifiable performance.

6.4 Phase 4: Expansion

Timeline: *Q3 2026 & beyond*
Focus: *Global Scale, Enterprise Integration & Autonomous Defense Ecosystems*
Status: *Upcoming*

The Expansion Phase sets the stage for τensorprox to become a globally distributed, enterprise-ready defense network, capable of autonomously responding to complex, high-scale cybersecurity threats. With the platform infrastructure in place, this phase focuses on hardening, scaling and diversifying τensorprox's capabilities to support commercial use in real-world context, in industries and geographies.

Key Objectives

- Deliver support for large-scale, multi-tenant deployments with dedicated Moat clusters per enterprise.
- Enable custom protection policies, traffic segmentation, and tenant isolation.
- Network Design Upgrade enabling Miners to operate as autonomous agents, exchanging learned threat patterns to form a collective defense intelligence layer.
- Launch governance mechanisms for decentralized decision making around upgrades, blacklists, validator audits, and economic parameters.
- Extend protection to IoT, edge computing, gaming servers, decentralized apps, and smart infrastructure, making τensorprox the backbone of a new programmable security layer for the Internet.

Phase 4 transforms τensorprox from a decentralized service into a globally adaptive cybersecurity fabric: a real-time, intelligent mesh that can learn, scale, and respond without centralized control. This will enable organizations, networks, and sovereign digital systems to operate with a new level of autonomous protection that evolves continuously with the threat landscape.

7 Conclusion

τensorprox represents a paradigm shift in cybersecurity defense by leveraging decentralized infrastructure to address high costs, while enhancing performance and reliability in response to the constantly growing sophistication of cyber-attacks. The competitive mechanism established between defense nodes creates a self-reinforcing system of continuous improvement driven by economic incentives rather than top-down control structures. The Alpha token serves as both a utility mechanism and governance instrument, creating sustainable value flows that support ongoing network development and operation. This economic layer transforms security from a cost center into a value-generating infrastructure component.

The implications of this work extend beyond DDoS protection, establishing a framework for decentralizing other critical infrastructure services that have traditionally required trusted intermediaries. By distributing both the execution and verification of security functions, τensorprox creates resilience against technical attacks and policy-level threats. As cyber threats continue to evolve, truly resilient systems will require defense mechanisms that match attackers in adaptability and distribution. τensorprox lays the groundwork for such systems, demonstrating that security at scale can emerge from properly aligned incentives and transparent performance metrics. τensorprox aims to establish itself as the standard security infrastructure layer for the digital economy, a unified platform that aggregates the most advanced cybersecurity defense solutions into a single system, purpose-built to deliver best-in-class protection for web clients.

Acknowledgements

We thank the Bittensor community for their support and contributions to τensorprox's development.

About Shugo

Shugo is a cybersecurity innovator building on the Bittensor network, delivering decentralized protection solutions to secure the digital economy with scalable, intelligent, and adaptive systems.

References

[1] Bittensor, *"Introduction to Bittensor"*, https://docs.bittensor.com/learn/introduction

[2] Cloudflare, *"Understanding DDoS attacks"*, https://www.cloudflare.com/fr-fr/learning/ddos/what-is-a-ddos-attack

[3] Akamai, *"DDoS Attack Trends in 2024 Signify That Sophistication Overshadows Size"*, https://www.akamai.com/blog/security/ddos-attack-trends-2024-signify-sophistication-overshadows-size

[4] Cloudflare, *"Targeted by 20.5 million DDoS attacks, up 358% year-over-year: Cloudflare's 2025 Q1 DDoS Threat Report"*, https://blog.cloudflare.com/ddos-threat-report-for-2025-q1

[5] Cloudflare, *"What is Anycast routing?"*, https://www.cloudflare.com/fr-fr/learning/cdn/glossary/anycast-network

[6] Akamai, *"Proxelic - Comprehensive DDoS Attack Protection"*, https://www.akamai.com/resources/product-brief/prolexic

[7] GrapeData, *"What are the challenges to Cybersecurity market research?"*, https://www.grape-data.com/blog/what-are-the-challenges-to-cybersecurity-market-research

[8] TrendsResearch, *"Cybersecurity in a World in Crisis: Funding Challenges, Investment Opportunities, and Sustainability Determinants"*, https://trendsresearch.org/insight/cybersecurity-in-a-world-in-crisis-funding-challenges-investment-opportunities-and-sustainability-determinants

[9] Cloudflare, *"Types of load balancing algorithms"*, https://www.cloudflare.com/learning/performance/types-of-load-balancing-algorithms

[10] B. Töpel and M. Karlsson, *"The eXpress Data Path (XDP): Introduction and practical use"*, https://www.kernel.org/networking/af_xdp.html

[11] Cloudflare, *"What is GRE tunneling? | How GRE protocol works"*, https://www.cloudflare.com/learning/network-layer/what-is-gre-tunneling

[12] Dashprox, *"Dashprox by Shugo"*, https://www.dashprox.com