PERMISSION ZERO

AI NEEDS PERMISSION.

The authority layer for humans, AI agents, and the autonomous internet.

LITEPAPER v0.1 / AUGUST 2026

STATUS: CONCEPT / FORWARD-LOOKING

This document describes Permission Zero's current research direction and proposed future plans. PZRO network economics, token supply, allocations, timelines, governance, technical architecture and launch mechanics are provisional and may change before implementation. No token is offered for sale through this Litepaper, and Genesis participation does not guarantee a future token allocation.

01 / EXECUTIVE SUMMARY

Permission should scale with intelligence.

Autonomous AI changes a basic assumption of the internet: the person who owns the intent may no longer be the actor making every request. Software can browse, negotiate, transact, write code, call APIs and coordinate other agents at machine speed. That creates enormous utility, but it also weakens the old shortcut that equated a user account with human intent.

Permission Zero is a proposed authority and defense layer for this new environment. Its goal is not to stop AI. It is to make legitimate autonomy easier to recognize and constrain, while making anonymous or hostile machine-scale coordination harder to abuse.

- Human Proof — establish the required human or organizational root with minimum disclosure.
- Agent Identity — distinguish software actors from their owners and from one another.
- Delegated Authority — express what an agent may do, for how long, and under which limits.
- Action-Time Policy — evaluate the requested action against current context, risk and revocation state.
- Swarm Defense — detect and constrain harmful behavior that only becomes visible across a population of actors.

Genesis is the pre-network phase: research, specifications, controlled simulations, builder programs and community formation. PZRO is described in this Litepaper only as a proposed future network-economic layer targeted for 2027, subject to legal, technical, security, market and community review.

NO PRESALE

Permission Zero is not conducting a token presale in the Genesis phase. The current design philosophy is to build the technology and community first, and consider token issuance only after the network model, legal structure, security architecture and distribution rules are ready.

02 / THE PROBLEM

The internet was built for human operators.

Most web security still begins with an account, a session, a device, an API credential or a wallet key. Those mechanisms remain useful, but autonomous agents introduce a new layer between the person who intends an outcome and the software that executes it. A valid credential can prove access without necessarily expressing the full purpose, delegation path, quantitative limit or current risk context behind an autonomous action.

One operator can become a machine-scale crowd.

AI also changes scale. A single operator can coordinate many agents, browser sessions, identities or workers. Each individual process may remain below conventional thresholds while the population as a whole scrapes data, captures scarce inventory, creates accounts, pressures APIs or manipulates online systems. The relevant unit of defense becomes the coordinated objective, not only the individual request.

“Human or bot” is too crude.

Permission Zero treats the autonomous internet as a multi-actor environment: humans, human-authorized agents, autonomous agents, organizations and hostile or unknown swarms. The useful question is not simply whether automation exists. It is whether an actor is legitimate, what authority it holds, and whether its behavior remains consistent with that authority.

This produces a different security objective: allow useful agents to participate without pretending they are humans, while preserving an application’s ability to enforce boundaries and protect shared resources.

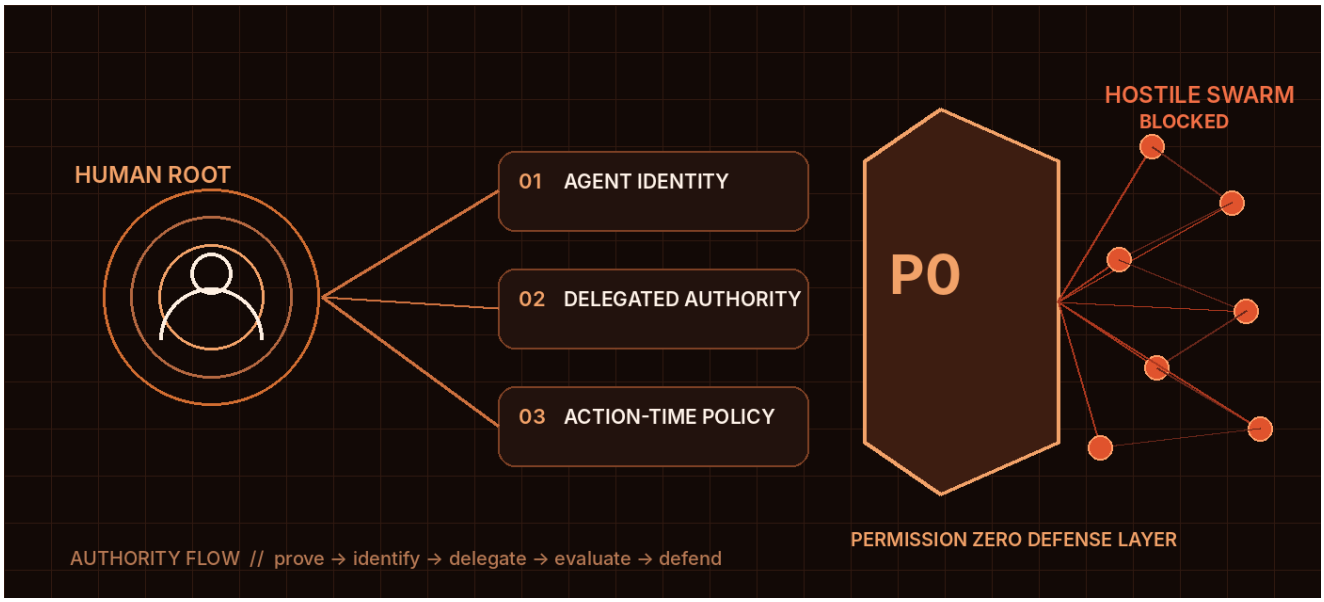

03 / PERMISSION ZERO THESIS

Human-rooted authority. Machine-speed execution.

The Permission Zero thesis is that intelligence can be delegated without delegating unlimited authority. A human or governed organization establishes the root of legitimate power. Agents receive explicit capabilities. Applications evaluate those capabilities at action time. Swarm defense adds a population-level control loop when many actors coordinate in ways that threaten fairness, availability or safety.

Five layers, one decision path

- 1. Root — the source of legitimate authority.
- 2. Identity — the agent or software actor presenting a credential.
- 3. Delegation — the capability envelope granted to that actor.
- 4. Policy — the action-time decision using scope, context, limits and revocation state.
- 5. Defense — behavioral and swarm-level controls when population dynamics matter.

The decision output is not a permanent trust badge. It can be allow, allow with limits, require approval, challenge, rate-limit, isolate or deny. The same agent can receive different decisions for different actions because authority is contextual rather than absolute.

Complement, do not replace, existing standards

Permission Zero is intended to compose with existing authentication and authorization systems rather than claim they do not work. OAuth scopes, enterprise identity, API credentials, device-bound keys, wallets and regulated identity checks can all remain inputs. The proposed layer focuses on making autonomous delegation provenance, narrow task intent, revocation, quantitative limits and swarm-aware policy easier to express together.

DESIGN PRINCIPLE

A verified identity is not the same as authority. A valid permission is not permanent trust. A legitimate agent can still become risky when context or population behavior changes.

04 / HUMAN LAYER

Prove human. Reveal less.

Many online systems need confidence that a legitimate human or governed organization exists behind an action, but they do not always need a complete identity dossier. The Human Layer is therefore designed around minimum disclosure: prove only the property required by the relying service at an accepted assurance level.

What a relying service may ask

- Does an accepted human-root or organizational-root proof exist?
- Does the root satisfy a required eligibility or uniqueness policy?
- Did that root authorize this specific agent or action class?
- Is the proof current, non-revoked and valid for this context?

Different environments can use different proof mechanisms. A consumer service might accept a privacy-preserving personhood proof. An enterprise might rely on its directory and role system. A regulated workflow may require stronger identity checks. Permission Zero's role is to consume the result and assurance level, not force every ecosystem into one universal identity database.

Privacy is part of the authority model

If every autonomous action required permanent public identity disclosure, the resulting system would create unnecessary surveillance, friction and data liability. The design therefore separates the fact needed for a decision from the underlying identity evidence wherever technically and legally appropriate.

RESEARCH STATUS

The Human Layer is a proposed architecture. Specific proof providers, cryptographic constructions, assurance levels and jurisdictional requirements remain implementation decisions.

05 / AGENT IDENTITY & DELEGATED AUTHORITY

Agents should act as agents - not impersonate their owners.

An autonomous actor needs an identity separate from the human or organization behind it. That identity should provide continuity across actions while allowing keys and credentials to rotate. A service can then distinguish the agent, inspect its authority chain and revoke one actor without destroying the root account.

Authority is an envelope

A useful delegation can express more than “yes” or “no.” It can bind the agent to a capability, resource, counterparty, time window, spend limit, rate limit, environment or escalation rule. A travel agent might compare flights freely, reserve a refundable itinerary within a budget and require fresh human approval before a non-refundable purchase.

FIELD	PURPOSE
WHO	Agent / software identity
ROOT	Who or what delegated authority
WHAT	Allowed capability or action
WHERE	Resource, environment or counterparty
WHEN	Start, expiry and renewal
LIMIT	Spend, rate, volume or risk ceiling
REVOKE	Immediate cancellation / rotation

Action-time policy

At execution time, the relying service evaluates the active delegation together with context and behavior. This keeps the model responsive when a budget is exhausted, an agent is revoked, a destination changes, risk spikes or the surrounding population starts behaving like a coordinated swarm.

06 / SWARM DEFENSE & SWARM LEAGUE

Defend against objectives, not only requests.

Traditional controls remain important, but coordinated automation can distribute activity across accounts, IP addresses, devices and API credentials. Swarm Defense is the proposed layer that looks for aggregate behavior: shared timing, resource targets, graph relationships, repeated adaptation, unusual velocity and patterns that only become visible across many actors.

Illustrative threat classes

- Account abuse — synthetic sign-ups, credential misuse and coordinated takeover attempts.
- Scraping — distributed extraction designed to evade simple per-client limits.
- Rate abuse — coordinated bursts or enumeration across many identities.
- Inventory capture — automation that monopolizes tickets, appointments or scarce goods.
- API flooding — application-layer pressure at bot and agent scale.

Swarm League

Swarm League is the community-facing testing concept: controlled Red-vs-Blue scenarios where attack swarms and defensive systems compete inside authorized sandboxes. The goal is reproducibility, measurable outcomes and public learning - not uncontrolled offensive activity against real services.

A future league can reward useful research, detection improvements, safe simulation tooling, challenge design and defensive engineering. Any reward mechanics described before network genesis remain provisional and do not guarantee PZRO allocation.

SAFETY BOUNDARY

Permission Zero and Swarm League are designed around controlled, authorized testing environments. The project does not endorse attacks on systems without permission.

07 / GENESIS COMMUNITY

Genesis starts before the token.

The first year of Permission Zero is intended to focus on technology, research and community rather than token sales. Genesis is the umbrella for people who want to follow, challenge, test and help shape the system before any final network-economic launch decision.

Community stack

- PZRO.IO — canonical source for specifications, research, disclosures and official links.
- X — public research, product updates and community growth.
- Discord — builders, support, technical discussion and Swarm League coordination.
- GitHub — specifications, code, issue tracking and verifiable technical artifacts.
- Telegram — official announcement-only channel to reduce impersonation and scam noise.
- LinkedIn — institutional and enterprise-facing presence.

Participation history, not a token-farming checklist

Permission Zero may maintain a history of meaningful participation such as testing, research feedback, security work, documentation, developer contributions, challenge design and sustained community involvement. The project should avoid publishing a formula such as “follow X = tokens” or “refer five users = tokens.” That model attracts bots and turns community formation into paid task completion.

Genesis participation does not currently guarantee a PZRO allocation. If a future community distribution is approved, historical contribution may be one input alongside anti-Sybil controls, legal eligibility, network requirements and the final distribution methodology.

COMMUNITY TARGET

The long-term ambition can be millions of participants, but the operating metric should prioritize genuine humans, active contributors and useful network participation over raw follower counts.

08 / PZRO — PROPOSED 2027 NETWORK ECONOMICS

Community-first. No public presale.

PZRO is a proposed future network asset, not a token currently offered for sale. The current design target is a fixed maximum supply of 1,000,000,000 PZRO if the network proceeds to issuance. The philosophy is that the network should be built and tested before ownership is distributed, with the large majority of supply reserved for the public/community ecosystem rather than insiders.

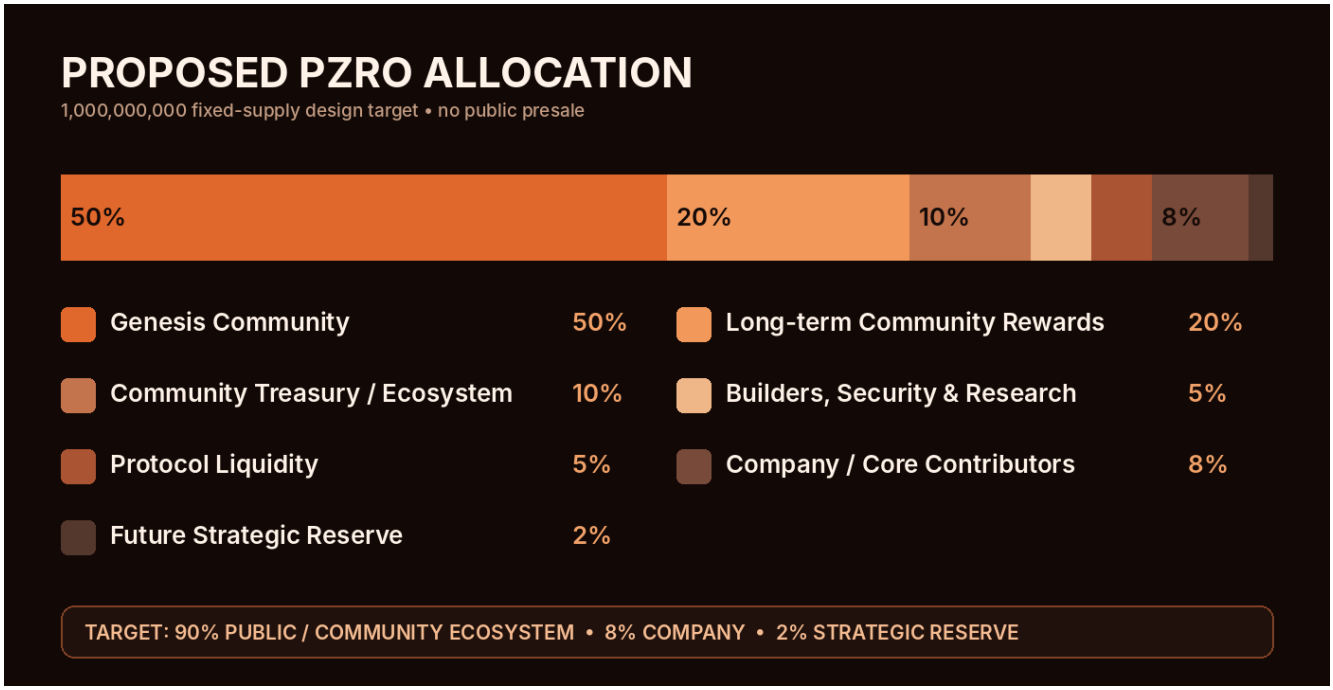

The 90% public/community ecosystem target includes Genesis distribution, long-term rewards, community treasury, builder/security grants and protocol liquidity. It does not mean that 90% of supply would become immediately transferable on the first day.

PROVISIONAL MODEL

The 1B supply and allocation percentages are design targets for Litepaper v0.1. They are not contractual promises and may change before genesis because of legal, regulatory, technical, security, market, operational or community considerations.

09 / RELEASE, CONTRACT & GOVERNANCE PRINCIPLES

Allocation is not circulation.

A community-first allocation only works if release mechanics protect the network from immediate oversupply, concentrated farming and short-term extraction. The current proposal is therefore to separate long-term allocation from day-one circulation.

Indicative release guardrails

- Genesis Community: only a limited portion claimable near network genesis; remaining allocation released through transparent epochs over time.
- Long-term Community Rewards: streamed over multiple years rather than distributed at once.
- Community Treasury: controlled through transparent governance and timelocks.
- Builder / Security Grants: milestone-based programs with published criteria.
- Protocol Liquidity: deployed gradually for market infrastructure; not an insider distribution.
- Company / Core Contributors: target 12-month cliff followed by long-term vesting.
- Strategic Reserve: unallocated unless required; if used, subject to public disclosure and vesting.

Immutable monetary policy after final deployment

If PZRO is launched, the preferred contract architecture is a fixed-supply, non-upgradeable token contract with no post-genesis mint function and no hidden supply-control privilege. Distribution vaults can remain separate so vesting, timelocks and governance can be audited without giving an administrator the ability to create additional supply.

The exact chain, contract implementation, audit process, governance structure and administrative transition are not frozen in v0.1. The final system should undergo independent security review before irreversible controls are removed.

NOT "RENOUNCE EVERYTHING ON DAY ONE"

Immutability should protect monetary policy without destroying necessary security controls prematurely. Distribution and governance contracts should be tested, audited and progressively constrained using transparent roles and timelocks.

10 / ROADMAP

Build the network first.

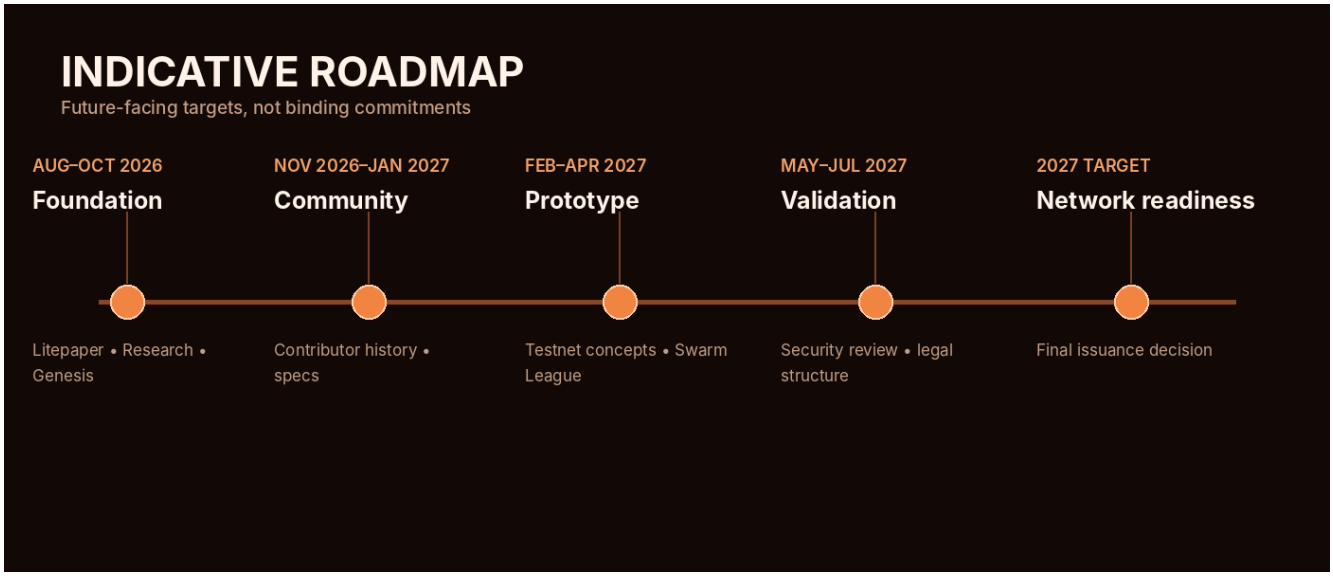

The roadmap is deliberately staged. Research and product definition come first. Community and contribution history follow. Prototype systems and controlled Swarm League environments can then provide evidence about what is useful, what fails and what should be changed. Only after legal structure, security review and distribution mechanics are mature should the project make a final token issuance decision.

What can change

The roadmap may move forward, slow down, change order or stop entirely. A future regulatory requirement may require a different issuance structure. Security findings may require contract or architecture changes. Community behavior may show that a proposed reward mechanism is too easy to farm. Market conditions may make a token launch inappropriate even if the technology is ready. The project should preserve the ability to choose the safer path rather than defend an outdated roadmap because it appeared in an early Litepaper.

2027 IS A TARGET, NOT A PROMISE

No date in this Litepaper creates an obligation to launch a token, network, exchange listing, airdrop or other virtual-asset event.

11 / FORWARD-LOOKING NOTICE & REFERENCES

This Litepaper is a living design document.

Permission Zero Litepaper v0.1 describes current research, proposed architecture and indicative future plans as of August 2026. It is intentionally not permanent. Token supply, allocations, eligibility, unlocks, governance, technology, chain selection, smart-contract design, launch timing and community programs may be amended, delayed, replaced or cancelled before implementation.

Nothing in this document constitutes an offer to sell, a solicitation to purchase, investment advice, a guarantee of allocation, a promise of exchange listing, a promise of liquidity or a representation that PZRO will have any particular value. No token is being sold through Genesis. Participation in community, research, testnet or social channels does not by itself create a right to receive PZRO.

If Permission Zero proceeds to a virtual-asset issuance, the project will need to follow the applicable rules of the relevant jurisdiction at that time. In Dubai, VARA's current issuance framework requires covered issuers to assess the applicable issuance category and, where required, publish a compliant Whitepaper and Risk Disclosure Statement before a virtual asset is made available to the public or marketed. This Litepaper is not intended to substitute for those formal issuance documents.

Selected references

VARA — Virtual Asset Issuance Rulebook and 2026 guidance
https://rulebooks.vara.ae/rulebook/virtual-asset-issuance-rulebook

VARA — Whitepapers and Public Disclosures
https://rulebooks.vara.ae/rulebook/whitepapers

U.S. SEC — Application of Federal Securities Laws to Certain Crypto Assets (Release 33-11412, Mar. 2026)
https://www.sec.gov/files/rules/interp/2026/33-11412.pdf

OpenZeppelin Contracts — Access Control
https://docs.openzeppelin.com/contracts/5.x/access-control

OpenZeppelin Contracts — VestingWallet / Finance
https://docs.openzeppelin.com/contracts/5.x/api/finance

VERSION CONTROL

v0.1 — August 2026. Future revisions should publish a changelog describing material changes to architecture, economics, governance or roadmap assumptions.