

World Quantum Computer (WQC) v0.4

~ A Decentralized Quantum Simulation Protocol to Shatter Centralized Hegemony ~

v0.4 restates v0.3 with **public protocol terminology** and makes the **SNARK wrap** of π_{Root} explicit for Layer-2 validity settlement. Internal engineering labels (Policy C, E1–E5, MP1/OP1, R1+R2, Phase A–C, and similar) are replaced with descriptive names; slice, proof, and economics semantics are unchanged.

1. Preface: Rebellion Against Quantum Hegemony (Anti-Hegemony)

1.1 The Centralized Quantum Hegemony and Its Threat

Throughout the history of modern computing, the internet and cryptography have served as shields protecting individual sovereignty and privacy. Yet the "dawn of quantum computing" we face today carries an acute crisis of centralization—one that threatens to overturn the decentralized balance of power built over decades.

Developing and operating practical quantum computers (QPUs) requires capital investments on the order of hundreds of billions of yen, dilution refrigerators that maintain millikelvin temperatures, and highly specialized academic teams. These physical and economic constraints have produced an **oligopoly of quantum computing resources** dominated by a handful of nation-states, national-scale technology giants (Big Tech), and military-industrial complexes.

The future this monopoly portends is censorship of scientific freedom and the entrenchment of inequality. Access to the computational power that will decide "the next century of wealth and security"—new drug discovery, breakthrough molecular simulation of novel materials, financial market optimization, and the breaking of public-key cryptography (RSA, elliptic-curve cryptography, and the like) that underpins modern secure society—is now fully controlled by the private clouds of gatekeeping monopolists. They hold the power to vet **who** computes and **for what purpose**, and to block or censor any computation that conflicts with their interests or national strategy at any time.

World Quantum Computer (WQC) is an explicit rebellion against this privileged-class **quantum hegemony**—an open-source protocol that achieves the **democratization of computational resources** through cryptoeconomics.

1.2 Breaking Physical Limits: From One Giant Machine to a Planetary Nervous System

Centralized players are obsessed with manufacturing "a single giant machine" loaded with ever more physical qubits. Yet scaling superconducting circuits, ion traps, and similar hardware hits exponential physical limits—the memory wall and the hardware wall—in noise (decoherence) control and manufacturing cost.

The WQC protocol abandons the dogma of building one physically perfect computer. Instead, it takes the approach of **cryptographically connecting all existing computing devices on Earth**—commercial GPUs, SoCs with unified memory, surplus data-center capacity—and transforming them into **a single distributed nervous system**.

Throughout human history, only a tiny privileged class could own supercomputers or physical quantum machines. WQC makes the very concept of "owning a computer" obsolete. People around the world contribute hardware to the protocol and share one another's compute power; the participants of the network **become the computer itself**. That is WQC's ultimate vision.

"We are the Computer."

(We do not own the computer. We **are** the computer.)

2. Physical and Structural Challenges to Solve

Contemporary computing science and the ecosystem of distributed ledger technology (DLT) face three seemingly insurmountable barriers. WQC defines these not merely as technical hurdles but as **structural defeats** that block the expansion of human intelligence—and seeks to dismantle them.

2.1 Censorship of Abundance: Oligopoly of Access and the Privatization of Science

Physical quantum processing units (QPUs) operated by centralized firms such as IBM, Google, and Rigetti—or petascale simulation clouds owned by Big Tech with thousands of nodes—have become fully **closed gardens (censored spaces)**.

Access requires not only exorbitant usage fees but also passing the **purpose-of-use reviews** imposed by providers. This is serious censorship against the essence of scientific inquiry: **freedom and anonymity**.

For example, researchers affiliated with certain nations or regions, or DAOs (decentralized autonomous organizations) pursuing open-source development of alternative energy or new drugs that conflict with corporate interests, live under the constant risk of having access cut without notice. When the source of wealth and intelligence—computational resources—is controlled by corporate will, the result is structural **censorship of abundance** that slows scientific progress.

2.2 The Memory Wall: Exponential Explosion of the State Vector

The most classical yet destructive approach to simulating quantum computation on general-purpose hardware is **state vector simulation**. To track the full quantum superposition (all probability amplitudes) of n qubits, this method must materialize 2^n complex numbers in memory.

Memory demand grows **exponentially** with qubit count.

Qubits (n)	Required Memory	Hardware Limit
30 Qubits	~16 GiB	Typical consumer PC ceiling
40 Qubits	~16 TiB	High-end data-center server clusters
50 Qubits	~16 PiB	Entire memory of world-class petascale supercomputers (beyond practical limits)

Building a single colossal supercomputer that can fully hold a state vector beyond 50 qubits is **physically impossible** under modern semiconductor physics and cost-effectiveness (ROI). This is **the Memory Wall**. Approaches that confine resources to one chassis or a closed cluster linked by a single high-speed internal bus have already hit theoretical limits.

2.3 The Futility of PoW: A Paradigm Shift from Thermodynamic Waste to Useful Work (PoUW)

Since Bitcoin's birth in 2008, Proof of Work (PoW) has proven the most robust method for eliminating centralized intermediaries and achieving trustless consensus. From a thermodynamic perspective, however, conventional PoW carries a fatal flaw.

Mining farms worldwide consume gigawatts executing **mathematically meaningless brute-force computation**: repeatedly changing a nonce until a hash function output meets a difficulty target. The resulting hash is consumed only to finalize the next block—dumping enormous heat and carbon into the environment while contributing **not a single bit** to human knowledge or scientific progress.

What is needed now is a paradigm shift that converts this thermodynamic waste into **useful work**—binding the security cost of maintaining the network directly to the executable cost of computations that advance humanity's frontier: quantum chemistry, cryptography, machine learning, and beyond (**Proof of Useful Work: PoUW**).

3. Technical Breakthrough: Hyper-Distributed Tensor Contraction

WQC overcomes the physical limits of holding quantum state vectors beyond 50 qubits by transforming quantum circuits into **tensor network (TN) representations** and applying **deterministic dynamic slicing**. This chapter defines the mathematical foundation by which tens of thousands of distributed Worker Nodes cooperatively compute a single massive quantum circuit in a **trustless** environment.

3.1 Tensor Network Slicing and Dynamic Optimization of Bond Dimension

Universal quantum circuit simulation beyond 100 qubits—the target of this protocol—is impossible even with all storage on Earth under conventional state vector methods. WQC makes it feasible through dynamic tensor network slicing.

Large quantum circuits are modeled not as state vectors loaded wholesale into memory, but as **collections of tensors**—gates as matrices, quantum states as indexed legs (dimensions). Simulating the full circuit reduces to **contraction**: collapsing tensor indices one by one.

The WQC Orchestrator analyzes the topology of the submitted circuit graph and dynamically identifies cut lines (slice axes) where **bond dimension** is minimized.

By fixing a specific index i (dimension d) and splitting the circuit, **slicing** decomposes one colossal task into d mutually independent **slice tasks**.

$$T_{\text{global}} = \sum_{k=1}^d T_{\text{slice}}^{(k)}$$

The greatest advantage of this mathematical decomposition is that **Worker Nodes require no real-time communication (synchronization)**. Each slice computes in parallel independently, enabling planetary-scale asynchronous grid computing unaffected by network latency.

3.2 Deterministic Proof of Useful Work (D-PoUW)

WQC eliminates the power waste of nondeterministic hash brute force entirely. **The process of executing quantum simulation (tensor contraction) and cryptographically proving its validity** becomes the network's security anchor (PoUW).

Worker Node n receives tasks not governed by luck but **deterministic** incentive rewards proportional to computational complexity and the cost of generating the zk-STARK proof π that guarantees it. Client payment and worker receipt are related but not identical:

$$\text{Total Fee (Planck)} = \text{Gas}_{\text{quantum}}(C, \pi) \times \text{BaseFee} \times 10^9$$

Where:

- $\text{Gas}_{\text{quantum}}(C, \pi)$: Endogenous **quantum gas** (dimensionless gas units) integrating memory and compute cost of contracting circuit C , plus algebraic cost (NTT, etc.) of generating zk-STARK proof π .
- BaseFee: Gas unit price in **Shannon per gas** (10^{-9} \$WQC / gas), set by network-wide supply–demand. The factor 10^9 converts Shannon→Planck so settlement math stays in Planck integers (§4.2).
- Total Fee: What the client pays for the task (before escrow refunds of unused headroom).

Of each settled Total Fee, **20% is burned** and the remaining **~80%** is the worker reward budget R_{net} (further split across compute / PCS roles in the normative economics spec). Tokenomics detail: §4.2–§4.3.

ASIC resistance in the WQC network—preventing governance oligopoly by giant specialized miners—is permanently enforced not by artificial memory-hard functions but by two **endogenous hardware barriers**:

1. Tensor Contraction Bottleneck (Tensor Access Phase):

Contracting indices of huge tensors exhibits **low arithmetic intensity**—the ratio of random memory buffer access to FLOPs is extremely high. No matter how densely compute units are integrated, the **memory wall** makes memory bandwidth the bottleneck.

2. Proof Generation Bottleneck (STARK Proving Phase):

Encoding contraction validity as AIR (Algebraic Intermediate Representation) constraints, then running NTT (Number Theoretic Transform) and FRI polynomial interpolation via Plonky3 and similar stacks. This demands ultra-wideband memory space for permuting millions of elements at speed.

Mining ASICs optimized only for generic ALU density are completely disadvantaged in ROI versus manufacturing cost for terabyte-scale random access and wideband memory (HBM, etc.). Commercial high-end GPUs (NVIDIA RTX series, etc.) and consumer SoCs with wideband unified memory (Apple M series, etc.) therefore achieve the highest execution efficiency by design—**forcing hardware democratization**.


3.3 Deterministic Verification via zk-STARKs and Recursive Proof (Recursive Aggregation Pipeline)

Having the Orchestrator or Layer-2 smart contracts re-execute and verify each result from tens of thousands of nodes would instantly collapse the system. WQC resolves exploding verification cost through **recursive proof composition**.

Each Worker Node performs contraction on its assigned slice and simultaneously generates cryptographic proof π_{slice} of mathematical validity. Binding the following **public inputs** into the STARK circuit prevents task substitution, wrong slice branches, data tampering, and theft (front-running) of other nodes' results:

$$\text{Public Inputs} = \{\text{circuit_id}, \text{sub_task_id}, \text{node_id}, \text{slice_id}, \text{output_result_hash}\}$$

Rather than verifying every π_{slice} directly, the Orchestrator (or aggregator) performs **recursive composition** in a tree (Proof Tree): generating **one new proof that verifies the fact that proofs π_1 and π_2 are both valid**.



Through this compression pipeline, the validity of thousands or tens of thousands of node computations converges to a single root proof (π_{Root}).

Aggregation is designed in stages: starting from native leaf-level STARK verification, progressing to binary tree combination via Aggregation STARK that binds child proof digests, and ultimately reaching full in-circuit recursion. Off-chain verifiers (Orchestrator FFI, aggregators, clients) check π_{Root} directly. For **on-chain** settlement, raw π_{Root} remains too large for L2 gas at current profiles; Layer-2 therefore verifies a succinct **SNARK wrap** π_{snark} of that root (§3.5), not the STARK blob itself. Exploiting mathematical asymmetry—heavy generation, instant verification—WQC gains overwhelming scalability.

3.4 Execution Model: Practical Quantum Algorithms and Deterministic Sampling

The domain of useful work (PoUW) that WQC targets is not limited to computing a specific amplitude of a single quantum state vector. Applying practical algorithms such as VQE for molecular energy calculation, QAOA for combinatorial optimization, and dynamic circuits for quantum error correction to real-world problems requires an execution model that treats the resulting **measurement histograms** and **expectation values** as deterministic quantities free of probabilistic fluctuation.

Distributed networks, however, cannot satisfy this requirement straightforwardly. Even when different nodes compute the same circuit, a unified verification framework is needed to handle the probabilistic nature of quantum sampling, the destruction of quantum correlations by measurement gates, and conditional branching within circuits. WQC resolves this challenge through the following four components.

3.4.1 Polymorphic Output Modes

What constitutes the result of a quantum circuit computation differs entirely depending on the algorithm's purpose. VQE requires the **expectation value of a Hamiltonian**, QAOA is decided by the **frequency of measured bit strings**, while amplitude-focused workloads need a **single complex probability amplitude** at a specific register.

Rather than forcibly stuffing these into a single data format, WQC provides three output modes optimized for each use case. Each mode produces its own **leaf STARK** (and optional auxiliaries); `statevector_scalar` is the lightest path, not the only leaf type.

- **Statevector Scalar Mode (`statevector_scalar`)**: Pinpoint-contracts the complex probability amplitude corresponding to a specific register (e.g., $|0\rangle^{\otimes n}$). The lightest leaf prove/verify path—no Born / histogram auxiliary.
- **Measurement Sampling Mode (`sample_counts`)**: Performs deterministic pseudorandom sampling for a specified number of trials (`shots`) on measurement gates, generating bit-string occurrence frequencies (histograms). By chaining the seed value, different nodes agree on an identical histogram. Leaf proofs include Born-rule binding (`DistributionAir`).
- **Observable Expectation Mode (`expectation`)**: For a Hamiltonian expressed as a linear combination of Pauli operators, algebraically computes the expectation value directly from the post-contraction state vector. Because this is not shot-based statistical estimation, consistent results are obtained across nodes even in a distributed environment.

3.4.2 Measurement-Preserving and Observable-Preserving Slicing

Tensor network slicing is a powerful technique for parallelizing computation. However, indiscriminately fixing the qubit axes targeted by measurement destroys the quantum-mechanical superposition (coherence) of the measurement results, preventing correct sampling or expectation values. To resolve this trade-off, WQC imposes constraints on slice axis selection per output mode.

- **Measurement-preserving slicing**: In sampling mode, qubit axes subject to terminal measurement (`MEASURE`) are automatically excluded from slice partitioning. Only non-measured qubits are compactly sliced, so contraction executes while preserving the measured quantum state in full, subject to Worker node memory constraints.
- **Observable-preserving slicing**: In expectation mode, fixing qubit axes on which the requested Pauli operators (X, Y, Z) act non-trivially is prohibited. This enables efficient parallel dispatch for circuits exceeding 28 qubits by

extracting only the relevant qubit neighborhood.

3.4.3 Trajectory Separation for Dynamic Quantum Circuits (Mid-Circuit Measurement)


Quantum error correction and quantum walk simulation require **dynamic circuits** that perform mid-circuit measurement and conditionally branch subsequent gates based on the result. Such destructive measurements, however, cannot be handled by existing STARK proof frameworks that assume purely unitary operations without measurement.

WQC solves this by fully separating **unitary traces** from **trajectory segments**. First, only the unitary portion with all destructive measurements removed is subjected to STARK proving. Then, the measurement events and classical register evolution occurring in each shot are independently bound as separate trajectory segments. By cryptographically chaining the state digest immediately before the first measurement (`unitary_link_digest`), consistency between the unitary proof and trajectory segments is guaranteed. This successfully integrates the validity of complex dynamic circuits containing classical conditional branching into the proof tree.

3.4.4 PRNG Binding and Born Rule Zero-Knowledge Proof

The foundation of consensus lies in **deterministic verification**. Yet quantum-mechanical sampling is inherently probabilistic: if each node independently used physical randomness for measurement, even from an identical circuit and state, different nodes would generate different histograms. Quorum cannot be established under such conditions.

WQC resolves this dilemma of **deterministic consensus** and **probabilistic sampling** through a **deterministic resampling pipeline**.



The concrete flow is as follows. The Orchestrator distributes a unique seed value (`sample_seed`) to all Workers. Each Worker runs inverse cumulative distribution sampling using deterministic pseudorandom numbers seeded from this value over the exact probability distribution (Born probabilities) obtained from tensor contraction. All Workers thereby derive an identical histogram.

Furthermore, that the probability distribution was correctly derived from the quantum state vector is attested as a partial zero-knowledge proof using a dedicated STARK circuit (`DistributionAir`). The client can mathematically verify that the presented sampling results contain no probability fluctuation or tampering, but are a deterministic outcome derived purely from the laws of quantum mechanics and the designated seed.

3.5 SNARK Wrap of π_{Root} for On-Chain Validity

Recursive aggregation solves *verification fan-out*: one root STARK attests the whole task. It does not by itself solve *on-chain gas*. A production π_{Root} (RecAgg over Plonky3 / Circle PCS) is still far larger than what an L2 block can verify as a native STARK. WQC therefore inserts a **succinct wrap layer**:


$$\pi_{\text{Root}} \longrightarrow \pi_{\text{snark}}$$

An off-chain wrap prover (a Redis/CAS worker alongside the composer) produces a **Groth16 argument over BN254** whose public inputs bind the settled task commitments—notably $\text{keccak256}(\pi_{\text{Root}})$, the economics receipt hash, and RecAgg header digests. The L2 settlement contract runs only a pairing check on π_{snark} and then pays, burns, and refunds atomically.

Two wrap depths are distinguished:

- **Thin wrap** — settle-aligned public inputs and RecAgg header bindings. Suitable for private rehearsal and KPI locks; it does **not** claim bit-for-bit equivalence with off-chain `verify_root_proof` (host-only PCS payloads are not fully re-checked in-circuit).
- **Thick wrap** — Plonky3 verifier gadgets in-circuit, thickened toward equivalence with `verify_root_proof` , after which mainnet can treat on-chain finalize as full root validity.

Phase 3 may still use **optimistic settlement** as a **rehearsal / fallback** (commit root/receipt hashes, challenge window, dispute via off-chain STARK verify) while the wrap track matures. **Mainnet validity finality** is defined by the SNARK-wrap path—thick wrap toward $\equiv \text{verify_root_proof}$ —not by optimistic commit alone. Normative cryptographic detail lives in the protocol specs (zk-STARK for π_{Root} ; zk-SNARK for π_{snark}).



4. Tokenomics & Governance

WQC's economic design is grounded in cryptoeconomics where token value is determined purely by the cycle of **computational resource provision and consumption**, completely excluding governance oligopoly by capital. To prevent rounding error and precision loss on computers, all currency and fee calculations in the protocol use 18-decimal fixed-point arithmetic (arbitrary-precision integers) without floating point.

Currency Unit System


Currency and economic value in this protocol are strictly defined and managed by three units:

Unit	Symbol	Fraction of 1 \$WQC	Origin and Use
WQC	\$WQC	10^0 (1)	Base currency unit.
Shannon	\$sWQC	10^{-9} (one billionth)	Base unit for quantum gas (Gas Price). (Named for Claude Shannon, father of information theory)
Planck	\$pWQC	10^{-18} (one quintillionth)	Absolute smallest unit in the protocol. All code operates on integers in Planck units. (Named for the Planck constant, physics' minimum length)

4.1 Supply Curve and Fair Launch Incentive Design

Total \$WQC supply is fixed at **210,000,000 \$WQC** (integer form: $210,000,000 \times 10^{18}$ Planck) with **no additional issuance under any circumstances**. The protocol adopts a **Pure Fair Launch** with no pre-sale or VC allocation.

To secure initial contributors and sustained ecosystem development, total supply is strictly allocated and locked by smart contract at the following ratios:



4.1.1 Foundation Allocation Liquidity Restrictions (Vesting Contract Specification)

15% of total supply (31,500,000 \$WQC) is allocated to the **WQC Foundation** for sustained development funding and founder team incentives. To prevent sudden inflationary pressure on markets and align founder and community interests long-term, the following **linear vesting** is enforced by L2 smart contract:

- **Cliff:** For 12 months from network launch (Genesis), not a single Planck of foundation allocation may be withdrawn (fully locked).
- **Linear Release:** After cliff, liquidity unlocks evenly in Planck units over the remaining 48 months (5 years total), second-by-second (timestamp-dependent).

4.1.2 Founder Mining (The Early Adopter Advantage)

In the earliest period after Genesis when participation is lowest, it is fully legitimate on-protocol for the founding team to deploy compute resources (GPU/VRAM) and run D-PoUW—and it represents the greatest opportunity for **real yield**. Like Bitcoin, game-theoretic incentives ensure nodes that take early risk and support the network at peak efficiency earn the most \$WQC (Planck); founders can build wealth openly and fairly.

4.2 Quantum Gas Market Price Mechanism

When clients request quantum circuit simulation, fees consumed are computed as **quantum gas**. Required gas for one task is:

$$\text{Gas}_{\text{quantum}}(C, \pi) = \alpha \cdot \text{VRAM}_{\text{peak}}(C) + \beta \cdot \text{FLOPs}(C) + \gamma \cdot \text{ProvingTime}(\pi) + \delta \cdot \text{SampleCost}(C, \text{shots})$$

Here `SampleCost` is a cost term for measurement post-processing and histogram aggregation when `output_mode` is `sample_counts` or `expectation`. For `statevector_scalar` only tasks, $\delta = 0$.

The transaction fee clients actually pay (Total Fee) multiplies this gas by dynamically varying gas unit price `BaseFee` according to network congestion:

$$\text{Total Fee (Planck)} = \text{Gas}_{\text{quantum}}(C, \pi) \times \text{BaseFee} \times 10^9$$

Units: `BaseFee` is quoted in **Shannon / gas** (1 Shannon = 10^{-9} \$WQC = 10^9 Planck). Multiplying by 10^9 yields Planck integers—the only unit used in ledger and on-chain settlement. Example: Gas = 1, BaseFee = 1 Shannon/gas $\Rightarrow$ Total Fee = 10^9 Planck = 1 Shannon. This is the same Total Fee introduced in §3.2; worker net receipt after burn is §4.3.

When the global pending task queue exceeds a threshold, `BaseFee` rises automatically to encourage node entry (supply increase). When demand falls, `BaseFee` converges to a floor, offering clients low-cost computation. All multiplication and addition in orchestrator internals and on-chain use arbitrary-precision integers (Go `big.Int`, Rust `U256`, etc.)—**not a single Planck of rounding error**.

4.3 20% Automatic Burn (Deflationary Burn) and Economic Sustainability


To provide permanent deflationary pressure (upward price pressure) on the WQC ecosystem, **20% of total fees (Total Fee) paid by clients is immediately sent to a burn address by smart contract at settlement and permanently destroyed (burned)**.

Net worker budget R_{net} (the ~80% of Total Fee after burn; not the same symbol as client Total Fee in §3.2) is enforced by integer arithmetic:

$$R_{\text{net}} = \lfloor \text{Total Fee} \times 80 \div 100 \rfloor$$

$$\text{Burn Amount} = \text{Total Fee} - R_{\text{net}}$$

(Here $\lfloor \dots \rfloor$ denotes integer division truncating remainder. Sub-Planck remainder from truncation is added to Burn Amount automatically.)



4.3.1 Mathematical Consequences of Deflationary Economics

With this 20% automatic burn, the higher WQC usage (compute demand), the faster circulating \$WQC supply decreases.

$$\frac{d(\text{Supply})}{dt} = -0.20 \times \sum_i \text{Total Fee}_i$$

In this absolute deflation model, even if total supply shrinks and per-\$WQC market value soars extremely, the protocol supports payments subdivided to 18 decimal places (Planck units), so **settlement gridlock (liquidity crisis) cannot occur in principle**.

In a future of soaring value, setting BaseFee as low as 1 Planck level lets clients always request quantum computation at fair, tiny costs aligned with real fiat rates, while all token holders (founders, miners, long-term investors) enjoy sustained asset appreciation—a mathematically guaranteed **flywheel**.

5. Development Roadmap: Path to the Sovereign Network

WQC evolves from single-machine verification through public testnet validation to a mainnet release with on-chain settlement, and ultimately to a **fully autonomous quantum compute network (Sovereign Network)** censored by no nation or corporation.

5.1 Phase 1: Foundation — Complete

Establishing single-node simulation engine and minimal end-to-end protocol pipeline proof.

- **Universal Gate Set:** `wqc-core` (Rust) supports major quantum gates (H, X, Y, Z, T, CNOT, CCNOT, etc.), enabling up to ~30-qubit state vector simulation on consumer hardware.
- **End-to-End Vertical Slice:** Connected Orchestrator, Worker Node, and Core; demonstrated pipeline from client submission through node compute to result aggregation.
- **Deterministic PoUW Prototype:** Abolished meaningless hash mining; confirmed "quantum circuit simulation + zk-STARK proof generation" operates as designed in a closed environment.

5.2 Phase 2: Scaling & Swarm Distribution — devnet Complete, Public Testnet Pending

Building the skeleton of a planetary Swarm on **libp2p**, connecting bidding, slice partitioning, verification, and off-chain economics. **devnet e2e for the items below is complete**; Phase 2 closes when the **Public Testnet** (§5.2.2) is live.

5.2.1 devnet PoC Status

The following core modules have passed end-to-end (e2e) verification in the current development environment (devnet):

- **libp2p Swarm Communication**: P2P implementation of task announcement via Gossip protocol, bid stream, dispatch, and result return.
- **Permissionless Bidding and Quorum**: Signed bidding, weighted lottery selection, node capability matching, slice-level majority vote (floating-point epsilon agreement and counts exact match).
- **Compact-register slicing**: BFS slice tree from bid-pool max-qubit capability. Each leaf is a compact-register circuit (width $N - C$ after fixing C tensor legs) executed asynchronously.
- **Leaf-to-root proof aggregation**: From leaf STARK verification to a single root-verification STARK via Aggregation STARK constructing a binary tree.
- **Polymorphic output modes and Born rule proof**: Sampling (`sample_counts`) and expectation (`expectation`) modes, dynamic circuits (mid-circuit measurement), deterministic resampling via seed binding, and `DistributionAir` implementation for probability distribution validity.

5.2.2 Public Testnet

As the culmination of Phase 2, a **Public Testnet**—an open verification environment—will be released prior to mainnet launch. This testnet maintains Swarm logic equivalent to mainnet while prioritizing operational stability and efficient early debugging under the following specifications:

- **Open-Source Repositories and Permissionless Participation**: All relevant source code will be made fully public. Anyone worldwide can download a Worker node client, contribute their local compute resources (GPU/VRAM), and participate in verification on a trial basis.
- **Centralized Orchestrator Operation**: To guarantee task coordination and infrastructure stability, a single Orchestrator managed by the protocol foundation will operate during this phase.
- **Redis-Based Test Economy and Faucet**: The economic zone and settlement log adopt the Redis-based high-speed off-chain architecture carried over from devnet. A **Faucet** web interface will provide anyone with free test \$WQC tokens, enabling stress-testing of fee rules, operator ledger, client escrow, and economics receipts under real traffic.
- **No Physical QPU Proxy**: At this stage, no physical QPU connection will be made; verification focuses purely on classical distributed simulation.

5.3 Phase 3: Mainnet Launch & On-chain Settlement — Next Target

Following sufficient operational track record and security audits on the Public Testnet, the protocol transitions to the production environment (Mainnet), achieving full societal implementation of the tokenomics.

- **Atomic on-chain settlement**: End the provisional Redis-based operation of the testnet; deploy complete on-chain settlement on L2 smart contracts. After an economics receipt, one settlement action pays Workers, burns 20%,

and refunds unused escrow. Genuine \$WQC with real value and the Foundation's 5-year linear lock-up (vesting) contract are activated on-chain.

- **SNARK-wrap validity path (mainnet gate):** A wrap prover reduces π_{Root} to π_{snark} (Groth16 / BN254). L2 `finalizeWithProof` verifies the wrap, then executes the same pay / burn / refund. Thin wrap is an intermediate statement for rehearsal and KPI locks; **thick wrap toward $\equiv$ `verify_root_proof` is the soundness gate for mainnet validity** (§3.5).
- **Optimistic path (rehearsal / fallback only):** Commit root and receipt hashes with a challenge window and off-chain STARK dispute. Useful for private L2 rehearsal and as an operational fallback while the wrap gate matures — **not** a substitute for the SNARK-wrap mainnet soundness claim.
- **Stable Operation Beyond 50 Qubits:** Optimize the Swarm power of globally distributed wideband nodes to stably accept and process universal quantum circuit simulations exceeding 50 qubits—beyond the limits of a single superclass supercomputer—as commercial and scientific tasks.

5.4 Post-Mainnet Evolution: Full Sovereignty and Hardware Integration

After mainnet launch, once 50-qubit-class simulation is established as distributed infrastructure and the network achieves economic and technical stability, the following long-term expansion roadmap will be initiated.

- **DHT-Based Orchestrator Full Decentralization:** Eliminate dependence on a single Orchestrator (a centralized single point of failure and censorship risk). Transition to an architecture where multiple aggregators coordinate autonomously via DHT over P2P for task partitioning, assignment, and aggregation, completing a true **Sovereign Network**.
- **Physical QPU Proxy Integration:** Go beyond a network of classical simulator nodes; enable external independent physical quantum computers (QPUs) to be plugged in as Workers via a standard interface. This elevates WQC from a "planetary-scale distributed classical simulator" to an **ultimate distributed quantum computing fabric fusing real machines with simulators**.

6. Conclusion: We are the Computer.

WQC is a **cryptoeconomic** answer to quantum-compute oligopoly: open the work, not the fridge.

1. **Physical limits** — tensor-network slicing turns Earth's idle GPUs/SoCs into one asynchronous grid.
2. **Useful proofs** — zk-STARK D-PoUW for leaves and RecAgg roots; **SNARK wrap** for L2 gas (§3.5).
3. **Practical outputs** — scalar, sampling, and expectation as first-class modes with deterministic PRNG / Born binding.
4. **Fair launch** — fixed supply, no pre-sale or VC allocation; 20% fee burn; rewards only for attested work.

WQC is not a corporate product or a state utility. It is open-source intelligence at planetary scale with no centralized master. People connect hardware, share compute, and **become** the computer.

"We are the Computer."

(We do not own the computer. We **are** the computer.)