AIDP.store — Decentralized AI Compute Network

Litepaper

1 Executive Summary

AI compute demand is exploding globally, driven by LLMs, multimodal models, and inference workloads that scale linearly with users. Traditional cloud services—centralized, expensive, region-limited, and unverifiable—are failing to meet global AI needs.

> *Full cryptographic verification (PoR/PoD) will be introduced during Phase 2, accompanying the decentralization of orchestration and routing. Phase 1 relies on simplified verification mechanisms to support the initial testnet rollout.*

AIDP.store solves this with a decentralized GPU marketplace built on DePIN principles:

- Globally distributed GPU providers
- Cryptographic verification of compute
- Decentralized orchestration (Phase 2)

- Token incentives for contributors
- Treasury-backed buybacks & burns
- Multi-chain settlement (Solana + EVM)

AIDP aims to become a permissionless AI infrastructure layer, enabling scalable, trustless, verifiable AI compute accessible to anyone.

---

2 Introduction

AI inference requires:

High-performance GPUs

- Parallel execution
- Token-streaming response
- Multi-tenant scheduling
- Global low-latency availability

But centralized cloud platforms suffer from:

Problem	Impact
High GPU prices	Limits builders + startups
No verifiable compute	Cannot trust outputs
Regional concentration	High latency for global users
Closed systems	No transparency or auditability
Expensive scaling	Makes global AI deployment impractical

AIDP introduces a global decentralized compute layer powered by thousands of independent GPU providers, incentivized and coordinated through cryptographic proofs, token staking, and decentralized routing.

---

3 System Architecture Overview

AIDP is built using a 5-layer modular architecture for performance and decentralization.

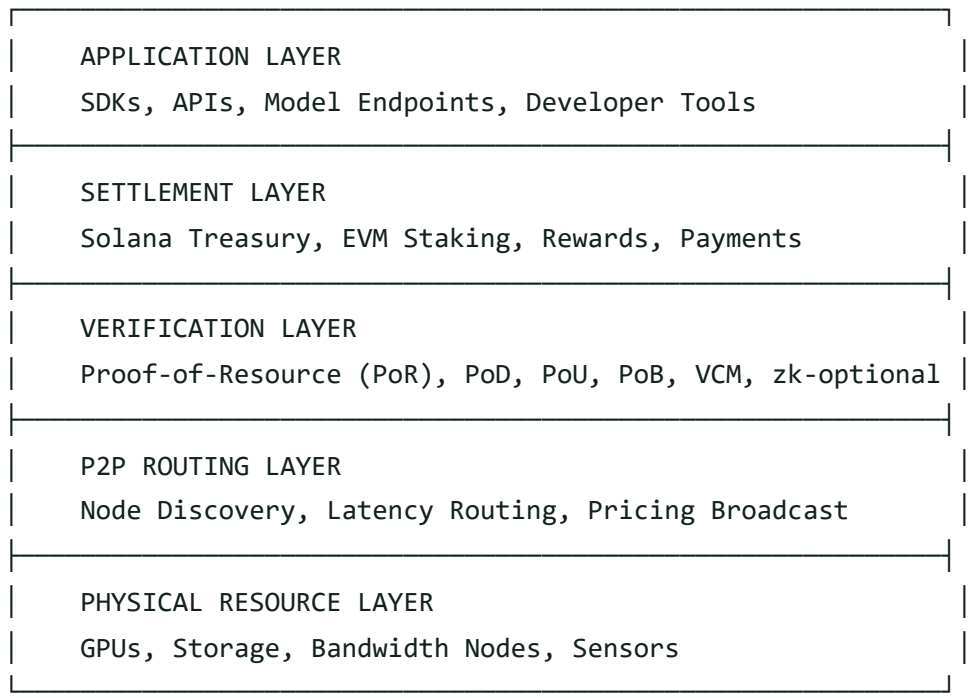

AIDP 5-Layer Architecture

Each layer is modular, upgradeable, and independently scalable.

---

4 Node Types in AIDP

> *Note: The PoR and PoD frameworks are part of Phase 2 of the AIDP roadmap. Their architecture is defined, but they will be implemented and decentralized in later stages as the orchestration layer*

> *evolves.*

AIDP supports four core node categories:

4.1 Compute Nodes (GPU Providers)

Provide GPU cycles for AI inference and general compute.

- Validate performance via PoR + VCM
- Earn rewards based on uptime and quality
- Must stake AIDP to participate

4.2 Storage Nodes

Store:

- datasets
- model checkpoints
- embeddings
- inference caches

Use periodic Proof-of-Storage validation.

4.3 Bandwidth Nodes

Relay inference traffic and provide high-performance networking.

- Validate via Proof-of-Bandwidth (PoB)
- Essential for global inference routing

4.4 Sensor Nodes (Future Expansion)

Edge devices supplying real-time data for specialized inference workloads.

> *PoR and PoD are not live in Phase 1. During the initial testnet, simplified verification and trusted environments are used. Full cryptographic verification will be introduced in Phase 2.*

5 Verification Framework

Trustless compute requires cryptographic verification. AIDP uses:

5.1 Proof-of-Resource (PoR)

Ensures the GPU actually executed the task.

```
PoR = H(Input, Model, Output)
```

5.2 Proof-of-Delivery (PoD)

Confirms output was delivered to the correct user.

5.3 Proof-of-Uptime (PoU)

Random heartbeat challenges ensure node reliability.

5.4 Proof-of-Bandwidth (PoB)

Measures real throughput for bandwidth nodes.

5.5 Verifiable Compute Module (VCM)

Recomputes or samples inference outputs to check correctness.

---

6 Phase 2 — Decentralized Inference & Orchestration Layer

Phase 2 introduces Orchestration Nodes, a decentralized routing layer that replaces centralized schedulers.

6.1 Role of Orchestration Nodes

They:

- Select which GPU provider should run each job

Measure latency, availability, price, VRAM capacity

Validate job completion (PoR + PoD)

Stake AIDP → earn routing rewards

6.2 Routing Algorithm (Simplified)

```
p* = argmin(αL + βC + γ(1−R) + δ/U)
```

Where:

L = latency

C = cost

R = reliability score

U = uptime

This ensures optimal provider selection.

Orchestration State Machine

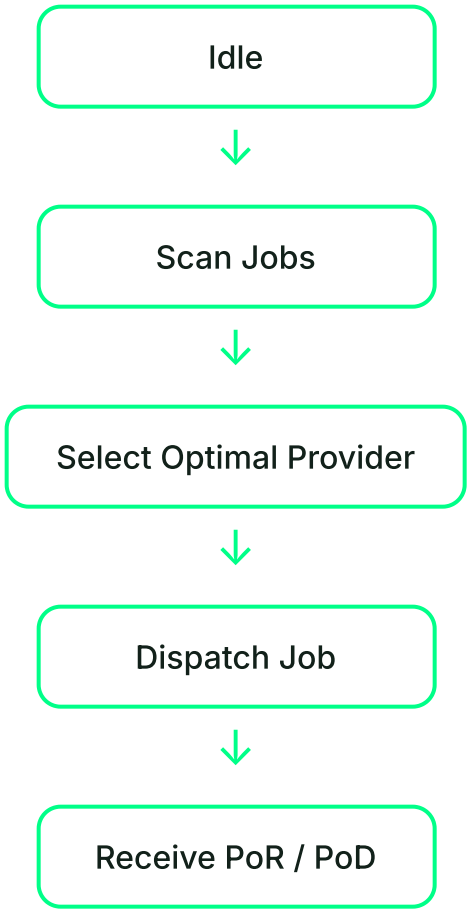

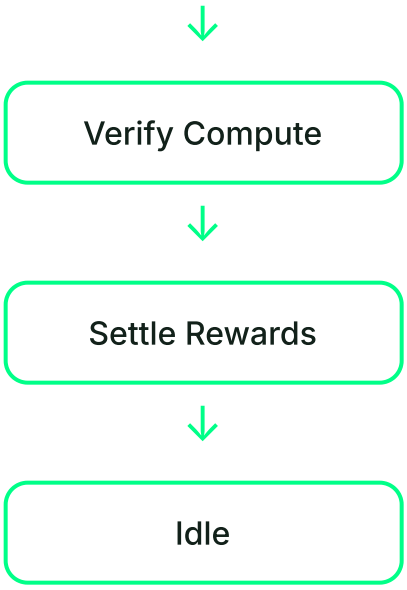

6.3 Benefits of Phase 2

- Fully decentralized scheduling
- Lower inference prices
- Higher reliability
- No single point of failure
- Market-driven compute allocation

---

7 Tokenomics Model

The AIDP token powers the compute marketplace, governance, and decentralized orchestration system. The tokenomics model balances:

- Utility demand (compute payments, staking, routing fees)
- Security (slashing, node bonding)

7.1 Token Supply Overview

Parameter	Value
Total Supply	1,000,000,000 AIDP

Parameter	Value
Circulating Supply	80% (fair launch)
Locked / Vested	20%

The system is intentionally designed with high initial circulation to reduce early manipulation and centralization risks.

7.2 Updated Allocation Table

AIDP Supply Distribution

Category	Allocation	Vesting	Purpose
Team & Contributors	5%	6-month cliff, then monthly	Long-term alignment
Marketing & Growth	5%	Fully unlocked at TGE	Community, KOL outreach
Exchange Listing	3%	Fully unlocked at TGE	Market-making, liquidity
Buyback & Burn Reserve	5%	Unlocks after 3 months	Deflation, price stability
Emergency Treasury	2%	Immediately unlocked	Security events, risk
Total Locked/Vested	20%	—	Long-term sustainability

7.3 Token Utility

AIDP is used for:

1. Compute Payments — Users pay GPU providers in AIDP for inference/compute.
2. Node Staking — Required for Compute, Storage, Bandwidth, and Orchestration Nodes. Slashing applies for misbehavior.
3. Governance Voting — Staked AIDP gives voting power.
4. Routing Priority — Users can stake AIDP for faster routing and higher priority queues.
5. Premium Access — Burn-to-access models/datasets.

7.4 Emission Model

AIDP follows a decreasing emission curve (mild inflation early, deflation later).

```
E(n+1) = E(n) / 2 — every 24 months
```

Ensures:

Strong early growth

Long-term deflation

7.5 Reward Formula (Simplified)

$$R_i = \beta \cdot w_i \cdot U_i \cdot D$$

Where:

β = base reward

w_i = resource weight (GPU=2.2, BW=1.5, ST=1.0)

U_i = uptime score

D = demand multiplier

Compute providers earn more during high network load.

7.6 Platform Fee, Revenue, and Buyback

Each computation request on AIDP pays a platform fee: $f \in [5\%, 10\%]$

Protocol revenue over a period is: $R = f \cdot V$ where V is total computation volume.

Protocol revenue is used for operations and growth.

When business conditions allow, the protocol may use a variable portion of revenue for token buybacks and burns: $B \leq R$

There is no fixed buyback percentage, allowing the protocol to dynamically adjust to market conditions while supporting long-term token value.

Slashing of misbehaving nodes may result in additional token burns.

7.7 Deflationary Condition

Total token supply change: $\Delta S = E - \text{Burn}$ where E is emissions and Burn includes buybacks and slashing.

The token is deflationary when: $\text{Burn} > E$

This condition is expected as:

Network usage increases (V↑)

Protocol revenue grows

Token emissions decrease over time

---

8 Economic Framework

The economic model must incentivize:

Supply (GPU providers)

Demand (AI developers, enterprises)

Reliability (uptime & correctness rewards)

Honest behavior (slashing)

Price stability (treasury buybacks)

8.1 Revenue Sources

AIDP generates revenue through:

1. Compute-as-a-Service (CaaS)
2. Model inference fees
3. Dataset storage fees

4. Bandwidth fees
5. Enterprise API contracts

8.2 Treasury Flow Model

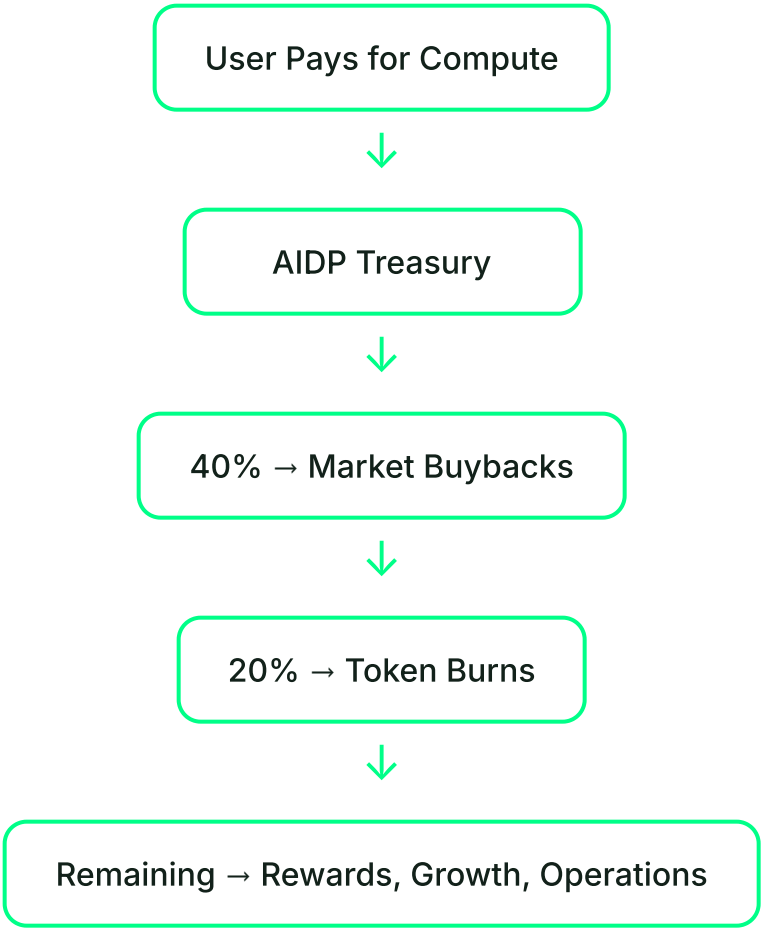

8.3 Treasury Growth Equation

```
T(t+1) = T(t) + Revenue − Expenses − Buyback
```

The treasury strengthens over time, enabling:

- Stability
- Consistent buybacks
- Liquidity support
- Development grants

---

9 Governance Model

AIDP uses a DAO governance framework controlling:

- Emission rates
- Slashing parameters
- Routing algorithm weights
- Treasury allocation
- Ecosystem grants
- Onboarding of new node classes

9.1 Governance Lifecycle

```
Proposal → Review (48h) → Voting (24h) → Timelock (36h) → Execution
```

9.2 Voting Mechanism

```
VotingPower = StakedAIDP × LockDuration
```

This rewards long-term stakers.

9.3 DAO Responsibilities

- Adjust incentive weights
- Maintain economic stability
- Approve treasury expenses

---

10 Security Model

Security is multi-layered:

10.1 Threat Mitigation Table

Threat	Description	Mitigation
Sybil Attacks	Fake nodes	Staking requirements
Fake Compute	Incorrect results	PoR + VCM verification
Provider Downtime	Node unreliable	PoU + slashing
Routing Manipulation	Orchestration bias	Multi-node routing audits
Data Tampering	Model/data modification	Hashing + PoS audit
Bandwidth Fraud	Fake throughput	PoB validation

10.2 Slashing Logic

Nodes lose stake for:

- Failed proofs
- Lost jobs
- Downtime
- Routing manipulation

```
Slash = s × Stake
```

10.3 Security Assumptions

- Honest majority in compute providers
- zk-proofs improve over time
- Orchestration becomes fully decentralized in Phase 2

---

11 Roadmap

AIDP's development roadmap is structured into four phases, aligning technical progress with network decentralization and adoption.

11.1 Phase 1 — Testnet & Core Compute Marketplace (2025)

Key Deliverables

- Public testnet launch
- GPU node onboarding (consumer + datacenter GPUs)
- Core Compute Marketplace (simplified verification; PoR/PoD not yet live)
- EVM staking contracts live
- SDK release (Python + JS)
- Basic routing system (semi-centralized early scheduler)

Outcomes

- First batch of active compute providers
- User testing for inference APIs
- Early network reliability metrics

11.2 Phase 2 — Decentralized Orchestration & Routing Layer (2025–2026)

Core upgrade to eliminate centralized scheduling.

Deliverables

- Introduction of Proof-of-Resource (PoR) and Proof-of-Delivery (PoD) as core verification mechanisms
- Integration of PoR/PoD into the network's decentralized orchestration and consensus layer
- Launch of Orchestration Nodes to remove the centralized scheduler
- Dynamic Pricing Engine V2
- Migration from trusted execution → trustless cryptographic validation

Goal: Fully decentralized, self-verifying compute network.

> *Phase 2 Verification Milestone: AIDP introduces the first production-ready implementation of Proof-of-Resource (PoR) and Proof-of-Delivery (PoD). These proofs ensure the compute node actually performed the requested work, the output was correctly delivered, no spoofing or replay is possible, and verification becomes trustless and decentralized.*

Benefits

- Lower inference cost
- Higher global availability
- Trustless routing
- Better GPU utilization

11.3 Phase 3 — Multi-Chain Expansion & Enterprise Integration (2026–2027)

Deliverables

- Cross-chain compute access (EVM, SVM, modular L2s)
- Enterprise inference SLAs
- Dataset marketplace
- Model marketplace

11.4 Phase 4 — Global Decentralized AI Infrastructure Layer (2027+)

Long-Term Deliverables

- zk-Proof inference pipelines
- Confidential inference (secure enclaves, ZK-ML)
- Multi-org model hosting
- Sovereign compute clusters
- Edge inference (IoT → GPU → Cloud)
- Unified cross-DePIN compute layer

Vision

AIDP evolves into a global infrastructure backbone for verifiable, decentralized AI compute.

Focus

Scaling throughput

Expanding liquidity

Strengthening treasury reserves

Partnering with AI/ML companies

Compliance routing: geo-restricted compute

100,000+ GPU providers onboarding target

Institutional onboarding

---

12 Architecture Diagrams

These diagrams illustrate the core system design.

12.1 System Architecture Diagram

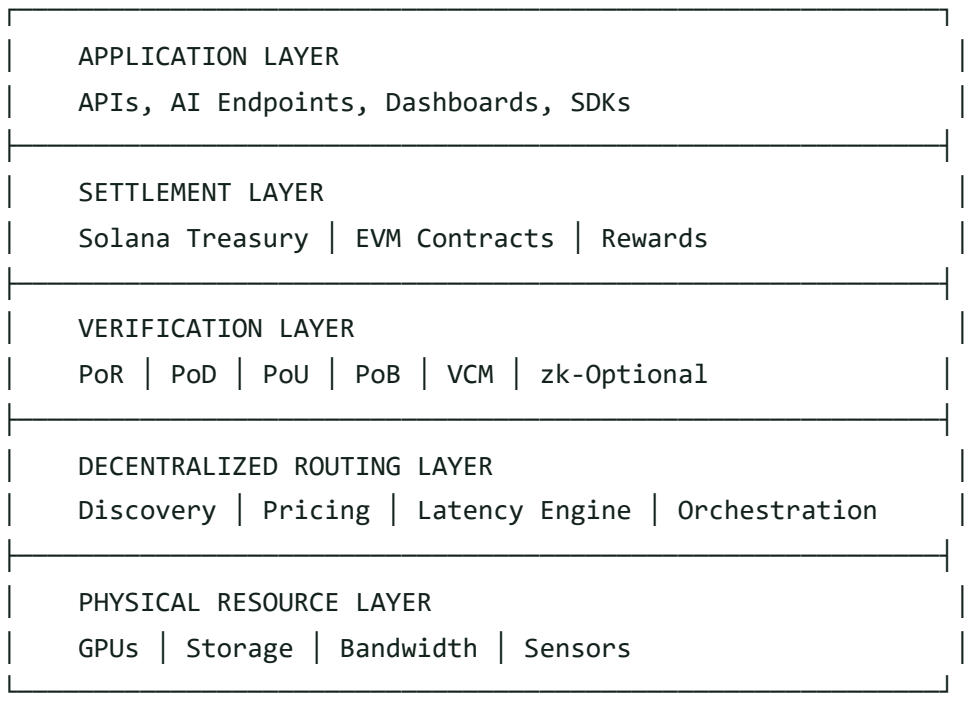

Complete 5-Layer Architecture

12.2 Node Type Diagram

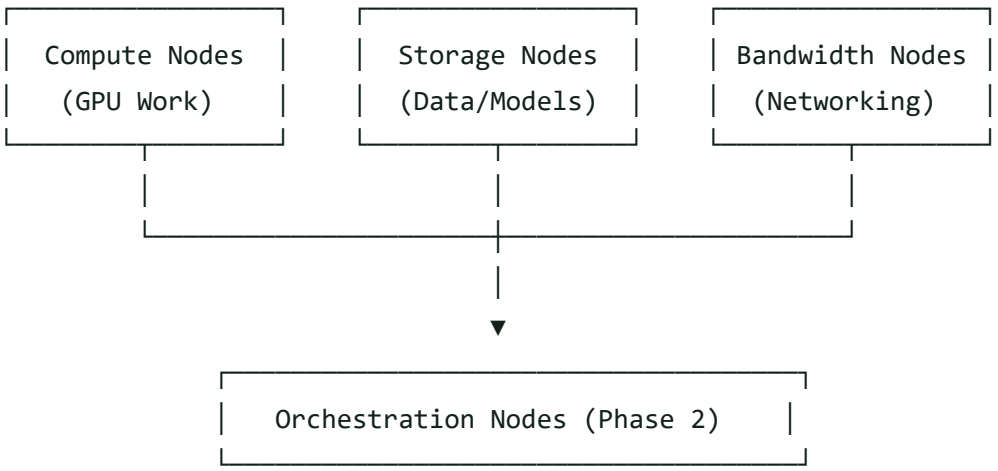

Node Types and Orchestration

12.3 Token Flow Diagram

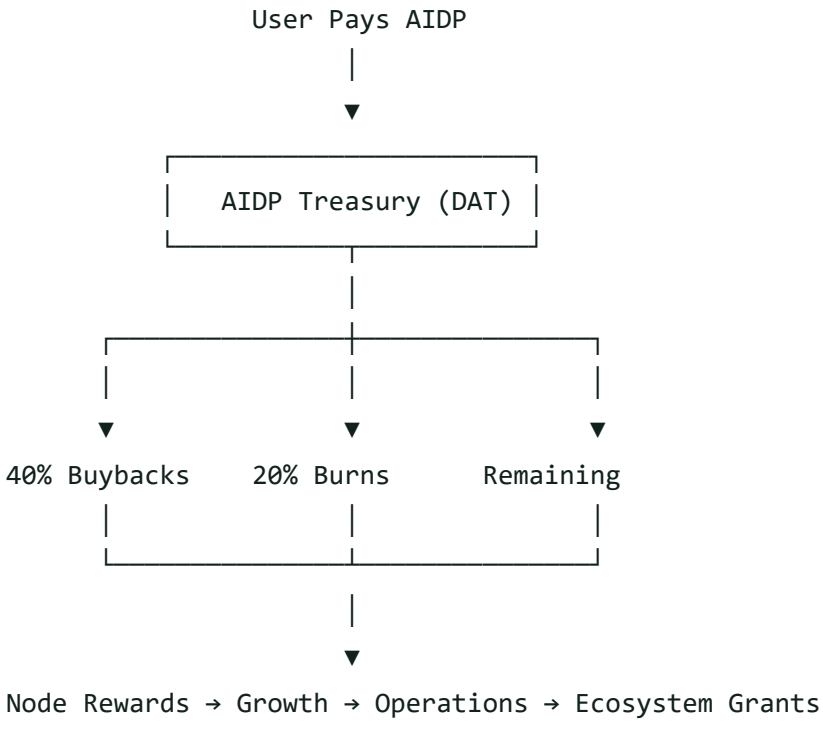

Token Economics Flow

12.4 Orchestration State Machine

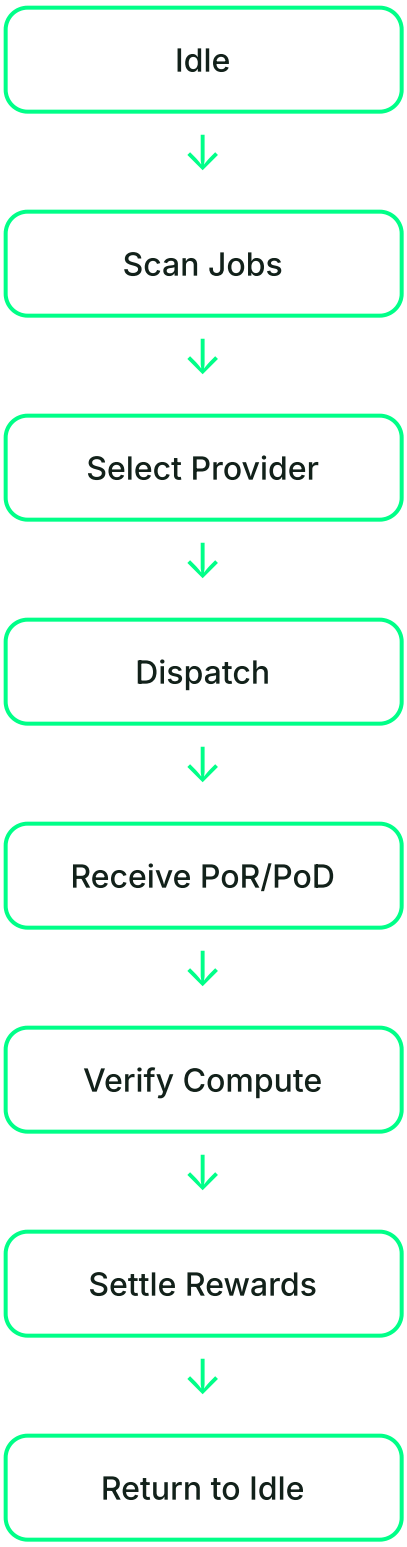

13 Compliance Overview

AIDP follows a hybrid compliance model aligned with UK and EU regulatory expectations.

Core Principles

- Optional KYC for large providers
- GDPR-aligned data framework
- Audit-ready compute logs
- DAO-controlled treasury governance
- On-chain revenue transparency

Enterprise clients get:

- SLAs
- Compliance routing (geo-bound compute)
- Private inference clusters

---

14 Future Extensions

AIDP aims to support:

- zkML inference verification
- Confidential compute
- Multi-chain GPU orchestration
- Incentive markets for datasets and model hosting
- Automated demand-driven pricing

---

15 Conclusion

AIDP.store delivers a next-generation decentralized compute network for AI workloads, powered by:

- A modular 5-layer architecture
- Cryptographic proof systems
- Decentralized Orchestration Nodes (Phase 2)

- Strong token utility
- A deflationary economic engine
- Treasury-backed incentives
- A global network of GPU contributors

AIDP is positioned to become a foundational layer for the decentralized AI economy—scalable, permissionless, verifiable, and economically sustainable.
