DEcntAI

DECENTRALIZED AI COMPUTE INFRASTRUCTURE

Whitepaper v3.0

May 2026

"The future of AI is not owned by corporations.
It is shared by everyone who contributes to it."

Platform
ai.decntai.com

ICO
ico.decntai.com

Token
DECNT / Solana

Executive Summary

Artificial intelligence is the defining technology of the 21st century. Yet the infrastructure that powers it — the servers, the compute, the model weights — is owned by a handful of corporations whose interests are not aligned with those of the individuals who use it.

DEcentAI is a decentralized AI compute marketplace built on Solana. It enables anyone with a capable computer to earn DECNT tokens by serving AI inference requests, and enables any user to access AI privately, at cost, without subscriptions. The platform has expanded from its initial language-model marketplace into a full multi-modal creative stack — image generation, video generation, text-to-speech, music generation, animated stickers, branded promotional cards, and a composition studio with AI-powered background removal. The platform is live today at ai.decntai.com.

Core Thesis

AI compute is a commodity. Like bandwidth or cloud storage before it, compute should be freely marketable — priced by the unit, not the month. DEcentAI creates the infrastructure for that market.

1B	80%	1%	0.001
DECNT	**Host Share**	**Referral**	**SOL/1000**
Total Supply	Of every prompt fee	Per referred user prompt	Current ICO price
Fixed forever	*Paid instantly*	*Passive forever*	*+20% early bonus*

1. The Centralization Problem

1.1 The Current State of AI Infrastructure

Three companies — OpenAI, Google DeepMind, and Anthropic — control the vast majority of all AI inference globally. This is not merely a market share statistic. It represents a structural dependency that affects every user, every developer, and every organization that relies on AI services.

The consequences of this concentration are measurable and growing:

- Users have no choice over who processes their data or where it is stored
- Every prompt is potentially logged, analyzed, and used to train future models
- Subscription pricing extracts value regardless of actual usage
- A single corporate decision, regulatory intervention, or outage can interrupt service globally
- Innovation is constrained to directions that serve corporate revenue models

> **The Utilization Gap**
> Studies of consumer hardware show average GPU utilization of under 8% globally. The compute needed to power AI at scale already exists in the hands of individuals — it simply has no marketplace to sell into.

1.2 Why Privacy Cannot Be a Policy

Current AI providers make privacy promises in their terms of service. These promises have three fundamental weaknesses that no policy can fix:

- Policies can be changed unilaterally at any time
- Promises cannot be verified by the user — there is no transparency into what actually happens to data after it is sent
- Legal jurisdiction creates complex and opaque data handling rules that most users cannot navigate

DEcentAI takes a different approach: structural privacy. When a prompt is sent through DEcentAI, it travels through an encrypted tunnel directly to the host node. The DEcentAI platform servers are architecturally excluded from the data path. We cannot see your prompts because the system was built so that we never receive them. This is not a policy — it is an architectural fact.

1.3 The Economic Exclusion Problem

AI access has become a subscription economy. ChatGPT charges $20/month. Claude Pro charges $20/month. Gemini Advanced charges $20/month. For a user who needs AI occasionally — a researcher, a small business, a developer in an emerging market — the choice is to overpay for a subscription or go without.

Pay-per-prompt changes the economics fundamentally. A user who sends ten prompts per week should pay for ten prompts. DEcentAI charges fractions of a cent per AI output token, with no subscription, no minimum, and no lock-in.

	ChatGPT Plus	Claude Pro	DEcentAI

Pricing model	$20/month	$20/month	Per prompt
Data privacy	Policy-based	Policy-based	Structural
Data logged	Yes	Yes	Never
Subscription required	Yes	Yes	No
Can earn by hosting	No	No	Yes — 80%
Open source models	No	No	Yes
Offline capable	No	No	Yes

2. Platform Architecture

2.1 High-Level System Design

DEcentAI consists of four principal layers: the user application layer, the marketplace coordination layer, the provider node layer, and the Solana settlement layer. Each layer has clearly defined responsibilities and communicates with adjacent layers through well-specified interfaces.

Layer 4 — User Application
ai.decntai.com web interface. Manages wallet, selects provider, sends prompts, receives responses. All cryptographic operations happen client-side.

Layer 3 — Marketplace Coordination
DEcentAI platform server. Maintains provider registry, routes session setup, handles heartbeat monitoring, manages reputation scoring. Never touches prompt data.

Layer 2 — Provider Node
Host computer running the DEcentAI worker and the runtime stack it has installed (language models, image diffusion, video diffusion, text-to-speech, music generation, sticker pipeline, promotional card renderer, and the Studio compositor). Receives encrypted sessions, runs inference locally, returns encrypted responses. Earns DECNT per task, per output token, or per second of generated media — depending on the capability invoked.

Layer 1 — Solana Settlement
SPL token contract for DECNT. Processes micro-payments per task, per token, or per second of generated media. Payment splits are now configurable through a transparent on-disk specification — defaults remain 80% host, 19% platform, 1% referral. Sub-second finality, fractions of a cent per transaction.

2.2 Session Flow — Step by Step

The following sequence describes a complete prompt transaction from user input to provider payment:

①	User loads ai.decntai.com and authenticates. Their encrypted account wallet key is decrypted server-side for the session.
②	User browses the provider marketplace. Each provider node advertises hardware tier, supported models, price per token, and uptime reputation.

③	User selects a provider. The platform establishes a dedicated SSH tunnel between the user's browser session and the provider's machine.
④	User sends a prompt. The text travels through the encrypted tunnel directly to the provider node — the platform server is not in the data path.
⑤	Provider's local runtime processes the request using the model installed on their hardware. For text, this is a language model loaded into memory. For image, video, audio, or Studio compositions, the request is routed to the appropriate runtime. No network call is made for inference — computation happens entirely on the host machine.
⑥	Provider streams the response back through the tunnel to the user's browser.
⑦	At session end, the platform calculates total output tokens. A Solana transaction is constructed and signed client-side.
⑧	Transaction broadcasts: 80% to provider wallet, 19% to platform, 1% to referrer (if user was referred). Settlement completes in under 1 second.

2.3 Provider Node Architecture

A DEcentAI provider node is any internet-connected computer that meets the minimum hardware requirements. Each node runs a unified worker process that coordinates one or more inference runtimes selected by the operator — language model, image diffusion, video diffusion, audio generation, AI sticker pipeline, promotional card renderer, or the Studio compositor. A single node can host multiple capabilities simultaneously; the worker advertises only the ones the operator has actually installed and enabled.

Tier	Hardware	Models Supported	Est. Tokens/sec
CPU Low	4–8 GB RAM	0.5B–2B models	5–12 t/s
CPU Mid	8–16 GB RAM	3B–7B models	3–8 t/s
CPU High	16–32 GB RAM	7B–14B models	2–5 t/s
GPU Entry	RTX 3060/4060 4–6GB VRAM	7B models	20–40 t/s
GPU Mid	RTX 3070/4070 8–12GB VRAM	7B–13B models	30–60 t/s
GPU High	RTX 3090/4090 16–24GB VRAM	13B–34B models	40–80 t/s
GPU Extreme	24GB+ VRAM	70B models	50–120 t/s
Multi-GPU	80GB+ total VRAM	70B–405B models	100–300 t/s
Apple Silicon	M1–M4 unified memory	All sizes	20–60 t/s

3. Security Architecture

3.1 The Trust Problem in Decentralized Systems

Decentralization introduces a trust paradox: users must rely on unknown third parties to process their data, yet the entire premise of the system is to eliminate dependency on trusted intermediaries. DEcentAI resolves this paradox through three complementary mechanisms: structural isolation, cryptographic verification, and economic alignment.

3.2 Structural Privacy — Why Architecture Beats Policy

The most important security property of DEcentAI is architectural: the platform coordination server is never in the data path of a user's prompt. This is not a policy choice — it is an engineering constraint.

- User prompts travel through SSH tunnels that terminate at the provider node
- The platform server facilitates session setup but does not proxy or log prompt content
- Each session uses ephemeral keys generated for that conversation only
- Even if DEcentAI's servers were compromised, no historical prompt data would be exposed

Security Principle

A security guarantee that depends on a corporation choosing to honour it is not a security guarantee. DEcentAI's privacy is enforced by the network topology, not by our intentions.

3.3 Model Integrity Verification

A provider running a tampered or substituted model is an existential threat to the platform's trustworthiness. DEcentAI mitigates this through SHA-256 model hash verification on every node heartbeat:

- When a provider registers a model, the expected SHA-256 hash is recorded on the platform
- Every 60 seconds, the provider agent computes the current hash of the loaded model weights
- If the hash does not match, the node is immediately suspended from the marketplace
- Users can verify the model hash independently before sending prompts

This approach means users know not just which model they are using, but that the weights they are using are identical to the publicly known weights of that model.

3.4 Hardware Attestation — TPM and TEE

For the highest security tier, DEcentAI supports hardware attestation through Trusted Platform Module (TPM) and Trusted Execution Environment (TEE) verification:

- TPM attestation provides cryptographic proof that the provider's machine has not been tampered with at the firmware level
- TEE environments (Intel SGX, AMD SEV) allow computation to occur in isolated memory regions that cannot be accessed even by the host operating system
- Nodes with verified attestation receive a trust badge in the marketplace and can command higher prompt fees
- The attestation report is published on-chain so any user can independently verify a node's security posture

3.5 Economic Security — Alignment Over Control

The final security layer is economic. A provider who behaves maliciously — serving fake models, logging user data, going offline mid-session — damages their own income stream:

- Reputation scoring reduces future bookings for nodes with poor performance metrics
- Payment is post-completion, not pre-paid, so providers must deliver to receive DECNT
- Node suspension removes the provider from the marketplace immediately upon detected violations
- Referral rewards create incentive for referrers to bring high-quality users, not spam accounts

This alignment means the security of the system improves as the network grows, rather than degrading — a property not shared by centralized systems where a single compromise can affect all users.

4. Incentive Architecture

4.1 The Payment Split

Every payment in DEcentAI is automatically distributed between three parties in a single atomic Solana transaction. The split percentages live in a transparent platform configuration that can be adjusted as the network matures, with all changes published on-chain. The current defaults — and the splits that govern every transaction at the time of this writing — are:

Share	Recipient	Rationale
80%	**Provider / Host**	The majority of value created goes to the party doing the work. This is deliberate — it makes hosting economically attractive without requiring large hardware investment.
19%	**DEcentAI Platform**	Covers infrastructure, development, security monitoring, and marketplace maintenance. The lowest platform fee in the decentralized AI space.
1%	**Referrer (if any)**	Perpetual passive income for bringing users to the network. Every prompt a referred user ever sends earns the referrer 1%, automatically, on-chain.

4.2 Provider Incentives — Why Host?

The provider incentive case is straightforward: idle compute is wasted compute. A mid-range gaming GPU running DEcentAI earns DECNT tokens 24 hours a day with zero marginal effort after initial setup.

- Setup takes one command — python3 setup.py — and approximately 15 minutes
- The node runs passively in the background with no ongoing management required
- Hardware tiers ensure any capable machine can participate — from a gaming laptop to a server rack
- Providers set their own price per token, creating a competitive market that benefits users
- Reputation scoring rewards consistent uptime and fast response times

Network Effect

As more providers join, users experience better availability, lower prices, and more model choice. As more users join, providers earn more. The incentive structure is self-reinforcing — growth benefits all participants.

4.3 User Incentives — Why Use DEcentAI?

For users, DEcentAI offers five distinct advantages over subscription AI services:

- Cost: pay only for what you use — no monthly fee for occasional users
- Privacy: structural guarantee that your prompts never leave the host machine
- Choice: select from multiple providers, models, and price points
- Ownership: earn DECNT by referring others — a perpetual 1% stake in every prompt they send
- Range: a single account and a single wallet access language, image, video, audio, sticker, promotional, and Studio composition tools — without juggling subscriptions to a separate vendor for each modality

4.4 The Referral Flywheel

The referral incentive model creates a powerful growth mechanism. Each user who registers through a referral link permanently links 1% of all their future prompt payments to their referrer. This is not a one-time reward — it is a perpetual stake.

- User A refers User B → earns 1% of User B's prompts forever
- User B refers User C → User B earns 1% of User C's prompts
- User A earns indirectly from the growth of the users below them

This creates genuine incentive to bring quality users — users who actually engage with the platform — because passive income scales with their referred users' activity, not just with signup numbers.

5. Tokenomics

5.1 DECNT Token Overview

DECNT is a Solana SPL token. It is the sole medium of exchange on the DEcentAI marketplace — used for language model prompts, image and video generation, audio synthesis, sticker creation, promotional cards, and Studio composition renders — and the primary store of value for network participants.

Token name	DEcentAI
Symbol	DECNT
Blockchain	Solana
Token program	SPL Token (TokenkegQfeZyiNwAJbNbGKPFXCWuBvf9Ss623VQ5DA)
Contract address	A3GTCErwHYaxjVx2M9N76c85DEjNCVNar6U4bPk7qXvX
Total supply	1,000,000,000 DECNT

Mint authority	Permanently disabled — supply is fixed
Freeze authority	Not set
Decimals	9
Pre-ICO price	1 SOL = 1,000 DECNT
Early bird bonus	+20% — 1 SOL = 1,200 DECNT during Pre-ICO

5.2 Token Distribution

Category	Allocation	Purpose
ICO (Main Sale)	**800M — 80%**	Primary token distribution to the public
Pre-ICO	**100M — 10%**	Early participants at 1,000 DECNT/SOL
Team	**50M — 5%**	Founders and core contributors — 2yr vesting
Ecosystem	**29M — 2.9%**	Partnerships, grants, developer incentives
Reserve	**20M — 2%**	Strategic reserve for future development
Airdrop	**1M — 0.1%**	Community tasks and early adopter rewards

5.3 Why Solana

The choice of Solana as the settlement layer is deliberate and technically motivated:

- Transaction finality in 400 milliseconds — fast enough for real-time prompt payments
- Transaction fees of fractions of a cent — economically viable for micro-payments per AI token
- Proven ecosystem — Phantom wallet, Jupiter DEX, Raydium liquidity pools
- SPL token standard — interoperable with all Solana DeFi infrastructure
- No EVM congestion — predictable fees regardless of network activity

6. Roadmap

Q1 2026 Completed ✓	► DECNT token launched on Solana mainnet — supply fixed, mint disabled ► AI marketplace live at ai.decntai.com — users, providers, payments working ► Pre-ICO platform live at ico.decntai.com with airdrop task system ► Referral system with permanent 1% prompt share ► SHA-256 model hash verification on every heartbeat ► 9 hardware tiers, 54+ supported language models across the catalog ► SSH tunnel-based provider connectivity ► Email verification, AES-256 encrypted wallet storage
Q2 2026 Completed ✓	► DEcentAI worker shipped — unified runtime dispatcher replacing the legacy

	Ollama agent ► TPM and TEE hardware attestation for premium node tier ► Provider reputation scoring system with on-chain records ► Image generation marketplace — DreamShaper, Stable Diffusion, SDXL-Turbo, Z-Image-Turbo, OpenJourney ► Provider withdrawal UI — direct DECNT transfer from provider wallet
Q3 2026 Mostly Shipped ✓	► Audio generation — text-to-speech (piper, NeuTTS) and text-to-music (Stable Audio) on capable nodes ► AI Stickers + Promo Cards — animated Telegram stickers via @decntai_stickers_bot, branded launch graphics ► Studio Compositor — multi-track composition into MP4 / WebM / GIF / animated WebP / TG-sticker ► AI Background Removal — ML-based subject segmentation integrated into the Studio Compositor ► Configurable payment splits — host / platform / referral percentages adjustable via on-disk configuration
Q4 2026	► AI Agent Orchestrator — natural-language workflows that chain multiple capabilities into one request ► DEX listing on Raydium and Jupiter — DECNT tradeable on open markets, plus Wormhole bridge to Ethereum and BSC ► DEcentAI DAO — DECNT holders vote on platform parameters, including payment-split percentages and staking rewards ► Mobile app + Enterprise SLA tier — iOS / Android client and dedicated provider pools for business customers ► CEX listings on major centralised exchanges, supported by traction from the agent orchestrator and multi-chain availability

7. The Long-Term Vision

7.1 AI as Public Infrastructure

The most important technologies of the past century — electricity, the internet, telephony — began as luxuries controlled by private monopolies and ended as public infrastructure accessible to all. AI compute is following the same trajectory. The question is whether the public infrastructure phase happens through government regulation of existing monopolies, or through decentralized alternatives that make monopoly economically unviable.

DEcentAI's thesis is the latter. When millions of GPU-equipped individuals can earn DECNT tokens by providing AI inference, the marginal cost of compute approaches zero. At that point, subscription pricing for AI becomes as absurd as subscription pricing for internet search.

7.2 Personal Compute Sovereignty

The long-term vision extends beyond AI prompts. DEcentAI's architecture — verified nodes, pay-per-use settlement, structural privacy — is a foundation for any computation that requires trust without centralization. Future applications include:

- Fine-tuning and training on distributed hardware

- Privacy-preserving data analysis — run computations on sensitive data without exposing it
- Edge AI inference for IoT and embedded systems
- Scientific computing — molecular simulation, climate modeling, genetic analysis
- Personal AI assistants that run exclusively on hardware controlled by the user

7.3 Why This Matters Now

The window for establishing decentralized AI infrastructure is narrow. As AI systems become more capable and more economically critical, the incentive for incumbents to prevent competition grows. Regulatory capture, API restrictions, and terms-of-service changes are already being used to limit third-party access to AI capabilities.

DEcentAI is being built now, while the technology is young enough that decentralized alternatives can compete on merit. Every provider node added to the network, every user who experiences structural privacy, every referral link shared — these are not just transactions. They are votes for a different kind of AI future.

> **The Bet**
> We are betting that enough individuals value privacy, economic participation, and technological autonomy to build a network that makes centralized AI monopolies economically irrelevant. We think that bet is correct. The Pre-ICO is your opportunity to join before the market prices in that possibility.

8. The DEcntAI Advantage

DEcntAI occupies a position in the AI ecosystem that no existing platform currently serves. Subscription services price for daily power users and exclude the long tail. Developer infrastructure prices by GPU rental and serves engineers, not end users. DEcntAI is neither — it is the marketplace layer that connects compute supply to compute demand at the granularity of a single task. Five properties of the platform are individually notable; in combination they define what the network is for.

8.1 Pay for Outputs, Not for Time

The fundamental economic mismatch in current AI services is that users pay for access while providers pay for capacity. A user who needs ten prompts a month pays the same monthly fee as one who needs ten thousand; a developer hosting a demo pays hourly GPU rental even when no one is using the demo. Neither model reflects the unit of value that actually matters: a completed AI task. DEcntAI prices what users want. A generated image, a finished sticker, a composed video — each carries a small fixed cost in DECNT regardless of how long the underlying compute takes or how often the user returns. Idle providers receive nothing. Active users pay only for what they receive. The subscription tax that excludes the long tail of users — researchers, students, hobbyists, occasional creators — disappears.

8.2 Privacy as an Architectural Property

Every centralized AI service offers privacy as a policy: terms of service describing how data is handled, how long it is retained, who can access it, and under what circumstances. These policies may be sincere, but they are unverifiable, unilaterally changeable, and irrelevant in the face of a single security breach, regulatory order, or corporate acquisition. DEcntAI offers privacy as an engineering constraint. The platform coordination layer is not in the data path of any prompt; SSH tunnels carry requests

directly from user browsers to provider machines. There is no log to subpoena, no database to leak, no policy change that would expose historical data. This is the difference between trusting a corporation and trusting topology — and topology does not change its mind.

8.3 A Two-Sided Marketplace, Not a One-Way Service

Most AI platforms have one side: paying customers. Every dollar of revenue flows to the platform operator, who pays for capacity from a cloud provider. There is no participant on the supply side who benefits from network growth, and no economic mechanism by which the platform's success accrues to anyone outside its corporate structure. DEcntAI inverts this. Anyone with a capable computer can become a provider, and providers receive the largest share of every payment for the work their machine performs. As demand for AI tasks grows, the marginal earnings of providers grow with it. Idle GPU hardware — gaming PCs, lab workstations, decommissioned mining rigs, prosumer machines — becomes income-generating infrastructure. The platform's role becomes coordination, not capacity provision, and the network scales through participation rather than through capital expenditure.

8.4 An Integrated Creator Stack

Single-purpose AI services have proliferated: one generates images, another generates video, a third removes backgrounds, a fourth produces speech. Combining these into a finished piece of content — a transparent animated sticker, a branded promotional graphic, an audio-backed product reveal — requires stitching multiple services together, each with its own account, pricing, and output format. DEcntAI consolidates these capabilities under a single account, a single wallet, and a single composition layer. The Studio Compositor accepts a background from any source, a foreground with optional AI-powered subject extraction, and an audio track from generated or uploaded sources, then produces a finished output in any of five formats — including the native Telegram sticker format and the WebM video format with true alpha transparency. Future development extends this with natural-language agent orchestration, in which the user describes the desired output and the platform chains the relevant capabilities to produce it automatically.

8.5 Resilience Through Distribution

Centralized AI services share a single point of failure: the corporation that owns them. A regulatory order, a security compromise, a strategic pivot, or simply an unilateral terms-of-service change can disrupt access for every user simultaneously. DEcntAI distributes risk across the network. No single provider being offline affects users beyond a brief moment of routing to a different provider. No single regulatory action can compel the network as a whole to behave any particular way. The DEcntAI organization itself can disappear and the protocol — the providers, the encrypted tunnels, the Solana settlement layer, the token — continues to function. For users in jurisdictions with AI restrictions, for use cases that touch sensitive information, for applications that need to outlast any single corporate sponsor, this property is decisive. AI compute becomes infrastructure rather than service.

9. Conclusion

DEcentAI is not a whitepaper about future possibilities. It is documentation of a working system. The marketplace is live. Providers are earning DECNT. Users are generating language model responses, images, videos, audio, animated stickers, promotional cards, and composed multi-track outputs — all without their data leaving the host machine. The token is on Solana mainnet with its supply permanently fixed.

What comes next is scale. More providers means more model choice and lower prices for users. More users means more income for providers. The referral system means growth compounds — each new user potentially brings their own network.

The Pre-ICO at 1 SOL = 1,000 DECNT (plus 20% early bird bonus) is the earliest point at which external participants can join the network. After the Pre-ICO, prices will increase. After the DEX listing on Raydium and the Wormhole bridge launch in Q4 2026, price will be determined by the open market across multiple chains.

Buy DECNT
ico.decntai.com

Start Hosting
ai.decntai.com

Join Community
t.me/decntai1

"Your hardware. Your rules. Your earnings."

DEcentAI — May 2026 — decntai.com