

---

PRODUCT ROADMAP · APRIL 2026

# HermesOS

## *Product Roadmap 2026*

*Managed AI Agent Hosting. Live in 5 Minutes.*

---

PLATFORM	TOKEN	DATE
hermesos.cloud	hermesos.cloud/token	April 2026

---

*This document describes what HermesOS is, what is live today, and where the platform is heading. It is written in functional terms. It is not a financial document and makes no investment claims about \$HermesOS or any other asset.*

## Managed AI Agent Hosting. Live in 5 Minutes.

April 2026 · hermesos.cloud

---

*This document describes what HermesOS is, what is live today, and where the platform is heading. It is written in functional terms. It is not a financial document and makes no investment claims about \$HermesOS or any other asset.*

---

## 1. What Is HermesOS

HermesOS is a managed AI agent hosting platform. It lets developers, operators, and teams deploy persistent, autonomous AI agents in under five minutes, with no infrastructure configuration required.

It is built on top of the open-source Hermes Agent framework by Nous Research, one of the most capable open agent frameworks available, with built-in memory, tool use, browser automation, scheduled tasks, skill creation, and parallel sub-agents. HermesOS removes every infrastructure barrier: no Docker, no VPS setup, no server management. Users bring their own AI API keys and get a fully managed, production-grade agent environment from day one, with zero markup on model costs.

*The privacy and cost transparency of self-hosting, with the reliability and speed of a managed service.*

---

## 2. Who Is HermesOS For

HermesOS is built for anyone who wants to deploy and run autonomous AI agents without managing infrastructure. That includes developers, operators, researchers, businesses, and teams across every background.

The platform is designed to be accessible from two directions:

Access Methods	
<b>Standard access</b>	Get started with a minimum credit top-up using a credit card. No wallet required, no crypto knowledge needed. Credits give access to the full platform: compute, operator packs, LLM gateway, Hive Mind, and marketplace. Everything is accessible by card.
<b>Token-based access</b>	Hold \$HermesOS in a connected wallet to access a free compute tier (1 vCPU, 2GB RAM). Compute tier scales with your token holding, verified automatically on-chain. Crypto-native users can use the token directly across all platform functions.

---

Both paths give access to the same platform and the same features. Whichever way you choose to enter, \$HermesOS powers the underlying infrastructure either way.

For users paying by card, the platform operates on two layers. The first layer is the one users interact with: a simple credits system. Top up, spend credits, run agents. No wallets, no tokens, no complexity. The second layer is the underlying infrastructure: the platform uses a shared pool to settle platform operations on-chain using \$HermesOS. Card users are funding the infrastructure that runs their agents. They are not buying tokens, and the tokens are not theirs. They are simply using a platform whose backend runs on-chain settlement, the same way most apps run on infrastructure their users never see.

### 3. What Is Live Today

HermesOS launched into production two weeks ago. Agents are live and running. The build is moving fast.

FEATURE	DESCRIPTION
<b>One-click deploy</b>	Deploy a fully configured Hermes Agent in minutes. No Docker, no VPS setup, no configuration files required.
<b>Prompt library</b>	Deploy agents with pre-configured system prompts. Start with a working operator immediately.
<b>Skill viewer</b>	View the skills your agent has available at a glance.
<b>Scheduled tasks</b>	Cron jobs and recurring automations built in. Agents run on schedule without manual triggers.
<b>Persistent memory</b>	Agents remember context, past tasks, and learned workflows across sessions. No context loss between runs.
<b>Credentials manager</b>	Securely store and manage API keys and secrets. Encrypted at rest, injected at runtime.
<b>Logs and monitoring</b>	Real-time streaming logs, health status, and resource usage visible in the dashboard.
<b>Multi-platform messaging</b>	Native connectors for Telegram, Discord, Slack, and WhatsApp out of the box.
<b>Auto-restart</b>	Agents restart automatically on failure. No manual intervention required.
<b>Daily backups</b>	Agents backed up daily. Never more than 24 hours from a clean restore.

### 4. Product Vision

*HermesOS is evolving from a deployment tool into the operating system for autonomous agent businesses.*

The next wave of AI is not about chat interfaces. It is about persistent operators: agents that run continuously, take real-world actions, learn from experience, improve over time, and fund their own compute. Most people cannot build this infrastructure themselves. HermesOS closes that gap.

The platform is moving in three directions simultaneously:

- **From deployment to operators** — today, deploying on HermesOS gives you a powerful but blank canvas. You get the full agent framework, but what you do with it is up to you. The next step is changing that: pre-configured, job-ready operators that arrive knowing their role, with every tool they need already installed. A working agent, not just a starting point, from the moment it deploys.
- **From isolated agents to a connected network** — every agent on the platform contributes to and draws from a shared intelligence layer called Hive Mind. Agents get better over time not just from what they individually learn, but from what the whole network learns. This is what makes the platform different as it scales.
- **From platform to economy** — agents that earn from their work, pay for their own compute, and transact with each other using \$HermesOS as the settlement layer.

A single HermesOS agent is useful. A network of thousands of HermesOS agents, all learning from each other and accessible through a shared token layer, is something categorically different. That is what the platform is building toward.

---

## 5. \$HermesOS Token

---

### ORIGIN

\$HermesOS launched as a community fair launch on Bankr on Base. There was no team pre-mint, no private allocation, and no VC distribution. The total supply is 100,000,000,000 tokens, fixed, fully in circulation from launch, held entirely by the community.

The token came to exist because the project got real attention. When Erik Voorhees, founder of Venice AI, publicly engaged with HermesOS and the community responded, someone launched the token through Bankr before an official one existed. HermesOS embraced it. The community had already shown what they thought of the project. The work now is to build the utility that makes the token genuinely useful inside a real platform.

<b>Total supply</b>	100,000,000,000 (100 billion) — fixed
<b>Launch type</b>	Community fair launch via Bankr on Base
<b>Team allocation</b>	None
<b>VC / private allocation</b>	None
<b>Contract</b>	Verify at <a href="https://hermesos.cloud/token">hermesos.cloud/token</a> only

---

## UTILITY

\$HermesOS is the utility token of the HermesOS platform. The token is not a financial instrument. It is functional infrastructure for access, payments, and platform participation. Token utility is being introduced in phases as the platform matures.

#	UTILITY	DESCRIPTION
01	<b>Free Compute Access</b>	Hold even a single \$HermesOS token and access a free compute tier: 1 vCPU, 2GB RAM. One token is the entry point. Compute tier scales with your token holding, verified automatically on-chain.
02	<b>LLM Gateway Credits</b>	LLM access through Bankr's gateway: Claude, GPT, Gemini, and more, at cost with no markup. Credits purchased through Bankr, payable in any cryptocurrency including \$HermesOS.
03	<b>Pay-As-You-Go Compute</b>	Rolling hourly billing. Pay only for the hours your server is online. Delete the server and billing stops immediately. Agents can spin up their own servers autonomously via x402 payment rails.
04	<b>Marketplace Access &amp; Creator Earnings</b>	Token holdings determine which marketplace content is available without individual purchase. Creators earn from usage and can withdraw through supported payout options, including token-based settlement.
05	<b>Agent Self-Sustainability</b>	Agents hold their own token balance and pay for services autonomously. Agents operate independently without needing constant human top-ups.
06	<b>Hive Mind Network Access</b>	Platform users access the collective Hive Mind intelligence layer at no additional cost. External agents from other networks can query it by paying in \$HermesOS per request.
07	<b>Agent-to-Agent Payments</b>	Agents pay each other for services settled in \$HermesOS. A research operator sells intelligence to a trading research operator. All settled in token, no human in the loop.
08	<b>Marketplace Settlement</b>	Marketplace transactions are settled using \$HermesOS as the underlying layer. Users and creators do not need to interact with this layer directly.
09	<b>Governance</b>	Token holders participate in decisions about ecosystem direction: integrations, marketplace policies, revenue split structures, and platform priorities.

Official token information is published only at [hermesos.cloud/token](https://hermesos.cloud/token).

## 6. Roadmap

Four phases. Each builds directly on the previous. Timelines are aggressive but realistic for a focused build.

---

## PHASE 1 — OPERATOR FOUNDATIONS + ECONOMIC INFRASTRUCTURE

**Timeline: Now — 6 Weeks**

*Ship flagship operators, Bankr integration, token-gated compute, and pay-as-you-go compute rails.*

Phase 1 is the foundation everything else sits on. It transforms HermesOS from a blank-canvas deploy tool into a platform with real operators, real payment rails, and the first live token utility.

### What Are Operator Packs

Operator packs are pre-configured, job-ready agent deployments. They are not skills or plugins. When you deploy an operator pack, everything it needs is installed and configured automatically: tools, capabilities, integrations, memory setup, and system prompts. There are no dependencies to manage, no configuration files to write, and no setup steps to follow.

A skill teaches an agent how to do one thing. An operator pack deploys an agent that already knows its job, has everything it needs to do that job, and is ready to run from the moment it goes live. You connect any API keys the operator requires and it works.

Looking further ahead, the platform is exploring token-based access to operator endpoints via x402 micropayments. Instead of managing API keys yourself, the operator pays for what it needs autonomously using \$HermesOS. This is exploratory and will be detailed when it is ready to ship.

### Flagship Operator Packs

Phase 1 ships three operator packs, each designed for a specific job and ready to run from the moment it deploys.

- **Research Operator** — continuous web research, source monitoring, automated summaries, and recurring intelligence briefs. Deploy a research agent that works around the clock in the background.
- **Trading Research Operator** — market monitoring, watchlists, signal tracking, thesis documentation, and real-time alerts. A serious intelligence and awareness layer for crypto and financial markets.
- **Growth Operator** — lead research, contact enrichment, outreach preparation, content repurposing, and workflow automation for business development and marketing operations.

### Platform Foundations

- **Skill management** — add, remove, and curate the skills available to your agent. Full control over what your operator can do.
- **File explorer** — browse and manage files your agent has created, downloaded, or is working with directly from the dashboard.

## Bankr Integration

- Deep integration with Bankr as the financial infrastructure layer. Bankr handles wallet provisioning, LLM gateway, and on-chain payment rails. HermesOS handles agent infrastructure.
- Crypto-native users get a provisioned Bankr wallet at account creation, enabling on-chain activity, token-based access, and LLM gateway usage from day one.
- Credit-based users do not receive a Bankr wallet. They interact with a simple credits balance only. The platform handles all underlying settlement on their behalf.
- LLM gateway credits purchasable through Bankr for token users, or via standard card top-up for credit users. Usable across Claude, GPT, and Gemini at cost with no markup.

## Compute Access

Compute access works the same way for both token holders and credit users. The mechanism is different, the result is the same.

- **Token path** — hold even a single \$HermesOS token in your connected wallet and you have access to a free compute tier (1 vCPU, 2GB RAM). One token is the entry point. Compute tier scales with your token holding, verified automatically on-chain.
- **Credits path** — start with a minimum top-up of \$5. Keep at least \$2 in your balance and your compute tier stays active (1 vCPU, 2GB RAM at the base tier). Tier scales with your balance.
- Credit users have an optional setting to reserve a portion of their balance so it is not accidentally spent, protecting compute eligibility. Entirely optional and reversible at any time.
- If your balance or token holding drops below the minimum, there is a grace period before compute pauses. Top up during that window and nothing is lost.

## Pay-As-You-Go Compute

- Need more compute than your token tier provides? Pay-as-you-go lets you scale on demand. You pay only for the hours your server is online. Delete the server and billing stops.
- Built for power users, multi-agent deployments, and agent endpoint operators. Your token or credits holding is your base. Pay-as-you-go is your ceiling.
- Agents can fund their own pay-as-you-go compute autonomously, or users can manage billing directly. Either works.

## Transition to Token-Based Access

HermesOS is moving from a subscription-based model to a token-based access system as part of the next phase of the platform.

Accounts created and active prior to this transition will be incorporated into the new system with recognised access levels and configuration once the infrastructure goes live. Full details will be communicated directly to users ahead of the transition.

---

## PHASE 2 — PLATFORM CONTROLS + DEEPER TOKEN INTEGRATION

### Timeline: 1 — 3 Months

*Add team features, agent wallets, expanded token utility, and an expanded operator pack library.*

#### Platform Controls

- **Pack builder** — create, save, and version your own operator configurations. Turn a working setup into a reusable template.
- **Approval controls** — agent approval functionality exists today and is being significantly deepened: richer controls, tighter security boundaries, and more granular per-action configuration before agents send messages, execute purchases, or take actions with real-world consequences.
- **Memory viewer** — see what your agent has learned, edit stored context, and control what persists across sessions.
- **Organisation workspaces** — team accounts with shared agents, shared memory, and shared workflows.
- **Role-based permissions** — control who can modify, deploy, or approve actions across a workspace.
- **Benchmarking and evals** — measure how well your operators perform against defined tasks. Track improvement over time.

#### Agent Wallets

- Agents hold their own token balance with defined spending limits and hard guardrails on autonomous transactions.
- Agents pay for their own LLM calls, compute, and services within approved boundaries. No human top-up required for ongoing operations.
- Spending approval controls: agents request user confirmation for actions above defined thresholds.

#### Expanded Operator Packs

Phase 2 expands the operator pack library beyond the three flagship operators shipped in Phase 1. Additional packs are added based on community demand and platform priorities, covering new use cases and workflows.

#### Expanded Token Utility

- **Usage credits** — token spendable on compute, LLM access, and premium features across the full platform.
- **Agent-to-agent payment foundations** — the groundwork for autonomous inter-agent transactions.

---

## PHASE 3 — MARKETPLACE + HIVE MIND

### Timeline: 3 — 6 Months

*Open the ecosystem. Community builds, publishes, and earns. Collective agent intelligence introduced.*

## Operator Marketplace

- Browse, install, and deploy community-built operator packs in one click.
- Verified publisher programme for trusted creators with accountability and track records.
- Flexible publishing options: release packs with open access, attach a price for individual purchase, or offer revenue share. Creators choose their own model.
- Community templates: free packs contributed by the community, instantly deployable.
- Private company packs: organisations maintain their own internal operator libraries, never publicly visible.
- Token-based marketplace access: token holdings determine which community packs and skills are available without individual purchase.
- Users can build and publish their own operators and agents. Creators earn from usage and can withdraw through supported payout options, including token-based settlement. Additional payout methods will be introduced over time.

## Agent Endpoints

One of the most powerful features of the marketplace: dedicated agent endpoints. Instead of sharing your prompts or tooling, you expose your agent as a callable service. Other users and agents can send tasks to your endpoint and receive results, without ever seeing how your agent works under the hood.

- Publish your agent as a private endpoint with a public interface. Keep your prompts, tools, and configuration completely private.
- Other agents and users call your endpoint to run tasks. You set the terms: open access, token-gated, or per-call pricing.
- High-demand endpoints scale automatically with pay-as-you-go compute. HermesOS manages the scaling infrastructure.
- Agent endpoint operators generate revenue from every call while their underlying agent logic stays entirely protected.

## Hive Mind Mode — Early Access

*The most significant long-term capability on this roadmap.*

Hive Mind Mode is a collective intelligence layer across the HermesOS network. Every agent on the platform contributes to and benefits from a shared knowledge base: learned skills, successful workflows, and refined strategies accumulate at the network level, making every agent smarter over time without any individual user needing to do anything extra.

Phase 3 introduces Hive Mind in early access. Current HermesOS subscribers get guaranteed early access when it launches. The full release follows in Phase 4 as the network matures.

- External agents from other networks can query the Hive Mind by paying in \$HermesOS per request. Access is token-denominated, verified on-chain, and priced per query.
- Agents share learned capability, not private data. Security and isolation are foundational to the design.

---

## PHASE 4 — AGENT ECONOMY + INFRASTRUCTURE LAYER

### Timeline: 6 — 12 Months

*Full autonomous agent economy. HermesOS as foundation for other platforms.*

#### Full Agent Economy

- **Agent-to-agent payments at scale** — agents pay each other for services autonomously, settled in \$HermesOS with no human in the loop.
- **Hive Mind full release** — the collective intelligence layer opens to the full network after early access validation in Phase 3.
- **Hive Mind external access** — agents from other networks query the HermesOS network via token-denominated endpoints. All access verified on-chain and priced per query.
- **Full on-chain marketplace settlement** — pack purchases, creator payouts, and platform fees settled transparently on-chain.
- **Governance live** — \$HermesOS holders vote on ecosystem direction, supported integrations, marketplace policies, and platform priorities.

#### HermesOS as Infrastructure Layer

Phase 4 opens HermesOS to developers and companies who want to build their own agent-powered products on top of the platform.

- **API access for third-party products** — companies build their own consumer-facing agent services using HermesOS as the underlying infrastructure.
- **Operators as a service** — teams wanting their own branded agent platform build on HermesOS without managing the infrastructure themselves.

---

## 7. What Is Not on This Roadmap

Being explicit about boundaries is as important as the roadmap itself.

- **Autonomous trading with user funds** — HermesOS does not create, operate, or endorse bots that autonomously manage or execute trades on behalf of users.
- **Yield products or lending returns** — HermesOS is an agent infrastructure platform, not a financial services product.
- **MEV strategies, liquidation bots, or front-running tools** — not operator categories HermesOS ships as first-party products.
- **Financial advice** — nothing produced by this platform constitutes investment, trading, or financial advice.

Users who build their own agents on HermesOS infrastructure are building their own products. HermesOS provides the platform, not the endorsement.

## 8. Where This Is Heading

*Deploy an agent that gets smarter over time, pays for its own compute, connects to the most capable agent network on the internet, and operates autonomously at any scale.*

HermesOS is building infrastructure for deployable agent businesses. The roadmap moves from deployment to operators, from operators to ecosystem, and from ecosystem to a network where agents and users interact through shared infrastructure and token-based access.

Phase 1 ships in six weeks. The operator packs, Bankr integration, token-gated compute, and pay-as-you-go rails are already in build. The platform is moving fast and staying focused.

The agent economy is not a future concept. It is being built now. HermesOS is where you deploy your operator.

<b>Platform</b>	<a href="https://hermesos.cloud">hermesos.cloud</a>
<b>Token verification</b>	<a href="https://hermesos.cloud/token">hermesos.cloud/token</a>
<b>Blog</b>	<a href="https://hermesos.cloud/blog">hermesos.cloud/blog</a>
<b>Twitter / X</b>	<a href="https://x.com/Wayland_Six">x.com/Wayland_Six</a>
<b>Token launched via</b>	<a href="https://bankr.bot">bankr.bot</a>

*This document describes planned product direction. Timelines are targets, not guarantees. Features are subject to change. Nothing here constitutes financial advice, a financial promotion, or an invitation to purchase any asset. For official token information, use [hermesos.cloud/token](https://hermesos.cloud/token) only.*