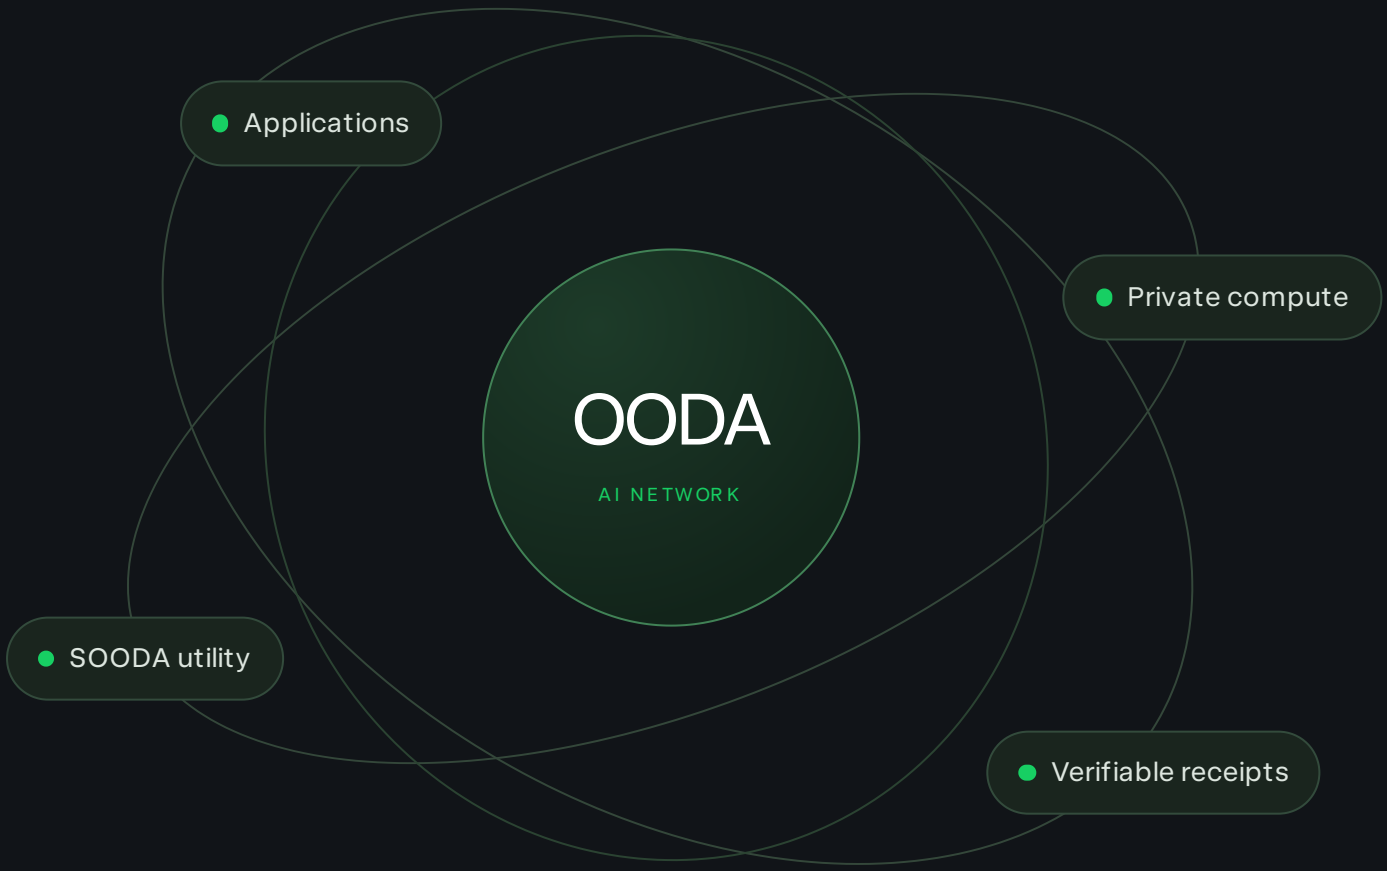
PRIVATE AI · VERIFIABLE COMPUTE

Intelligence. With proof.

OODA AI & the SOODA ecosystem



Litepaper

Version 3.8 · 8 November 2026

ooda.ai ↗

THE DOCUMENT

A clearer view of the network.

The product, the desktop apps, the OODA AI Network and the role of SOODA in everyday AI compute.

01 Executive summary 3

02 Market context 4

03 The problem 5

04 Approaches and trade-offs 6

05 The OODA AI solution 7

06 How OODA AI is used 9

07 Architecture and the OODA AI Network 10

08 Who OODA AI is for 16

09 SOODA utility and loyalty 17

10 Roadmap 25

About this litepaper

This document describes the OODA AI platform, its desktop apps, the OODA AI Network and the SOODA network design. It includes provider earnings and the compute loyalty programme, and distinguishes current capabilities from future deployment. Private subscription terms and payment instructions are provided separately.

Version 3.8 updates the OODA AI Network to the settings each organisation decides: its own computers only by default, App Attest and TPM requirements, Europe only, PII cloaking as a separate choice for text, and the limits of each protection. It supersedes version 3.7 as the project overview; it does not amend any signed agreement.

01 / EXECUTIVE SUMMARY

Make AI useful.
Make trust visible.

OODA AI brings AI applications, organisational knowledge and workflows into one platform, on the web and in desktop apps. Its network adds private compute, contributed computers and verifiable records.

Organisations need more than a capable model. They need control over data, clarity about how a workload is handled and evidence they can inspect. OODA connects those requirements to the tools people use every day.

Use AI

Chat, experts, knowledge, workflows and applications on the web and in desktop apps.

Protect execution

Cloud TEE routes for suitable workloads; sealed, attested and signed requests on the network.

Inspect evidence

Connect an authorised receipt to its verification status and recorded evidence.

THE ROLE OF SOODA

Compute credit. Participation. Network utility.

SOODA connects platform usage and contributed compute with service credits and rewards for qualifying activity. The planned maximum supply is 100 million tokens.

- The OODA platform and desktop apps are operational, and the OODA AI Network serves models and credits providers in EUR today. SOODA delivery, redemption, wallet-connected SOODA earnings, loyalty distribution and the broader production network are rollout milestones. A testnet demonstration is distinct from production token availability.

02 / MARKET CONTEXT

AI is part of everyday workflows.

As AI moves into business processes, governance and visibility become part of the product requirement.

From tool to service

Applications call AI during document handling, service interactions and internal workflows. Compute becomes a recurring part of delivering the service.

Control matters at scale

Data access, execution location and a usable audit trail matter when AI handles sensitive information.

Questions every deployment must answer

Requirement	What the organisation needs to understand
Data access	Which people and systems can access the inputs, outputs and retained records?
Execution boundaries	Where does processing take place, and which controls apply to that workload?
Operational evidence	Which records can an authorised team inspect to verify how a request was handled?

These requirements connect the application experience to the underlying infrastructure. A useful platform makes both the controls and their limits understandable.

What “inference” means

Inference is the processing of an input, such as a prompt or document, to produce an AI output. This is where private data enters a workload and where execution controls matter.

03 / THE PROBLEM

Sensitive work needs clearer boundaries.

A useful AI result is only one part of the requirement. Organisations also need to understand the path their data takes.

Data exposure

Prompts, documents and business context can contain confidential information. Controls must address who can access data during processing, how it is retained and how connected systems handle it.

Limited execution visibility

A log entry alone may not establish which software or hardware processed a request. Independent evidence needs a defined scope and a link to the underlying workload.

Practical adoption

Teams need a usable interface and a path from policy to daily work. Complex infrastructure can slow adoption even when the underlying technology is capable.

Usability, privacy and evidence need to work together.

OODA's objective is to make these controls available through ordinary applications, with enough visibility for technical teams to inspect the underlying records.

04 / APPROACHES & TRADE-OFFS

Different approaches. Different controls.

Enterprise AI can be deployed in several ways. The right assessment depends on the workload and the controls actually provided.

Approach	Strength	What to evaluate
Centralised AI services	Familiar interfaces and straightforward integration.	Data handling, retention, execution visibility and contractual controls.
Private cloud / on premises	Direct control over deployment and internal policy.	Operating burden, hardware isolation and independent evidence.
Decentralised compute networks	Multiple operators and shared coordination.	Privacy, provider quality, latency and enterprise usability.
OODA AI network design	Applications on cloud TEE routes and attested network computers, with inspectable evidence.	Each execution path's availability, protections and verification scope.

Familiar experience | **Defined privacy controls** | **Inspectable evidence**

OODA's approach combines application access, orchestration, private execution and verification. Each layer has a separate role, making it possible to inspect the boundaries between them.

Assess implementations, not labels

Cloud, private and decentralised systems vary. A deployment must be assessed on its actual access controls, data flows, hardware and evidence, rather than its category alone.

05 / THE OODA AI SOLUTION

A familiar platform.
A verifiable foundation.

Users work in applications. The underlying network coordinates execution, evidence and accounting.

01 Usage

Enterprises, developers, internal teams and end users access AI through the web, desktop apps, dashboards and APIs.

↓

02 Applications

Ready-made tools, configurable applications and branded deployments connect to a common platform.

↓

03 Network

Orchestration routes each workload to cloud TEE models, other cloud models or computers of the OODA AI Network, with verification records and the settlement design.

For users

AI, knowledge and workflows in one place.

For organisations

Policy linked to infrastructure and evidence.

For computer owners

Serve models and earn per request.

Privacy and verification properties depend on each application's execution path and configuration.

05 / DESKTOP APPS

The whole platform.
On your own computer.

Native OODA AI apps for macOS, Windows and Linux bring the full platform to the desktop, together with models that run on the computer itself.

macOS
Apple silicon Macs run local models with MLX.

Windows
Local models run with llama.cpp.

Linux
Local models run with llama.cpp.

The full platform
Chat, experts, documents, apps, workflows and the rest of the platform, in a native app.

Local models, also offline
Models downloaded to the computer answer on the computer itself, also without an internet connection.

Local models with platform tools
Connected to OODA, a local model can use platform tools such as web search.

Provide compute
Turn on Provide Compute to serve chosen models to the OODA AI Network and earn for every verified request.

The owner decides
Running local models and providing compute are separate choices. A computer joins the OODA AI Network only after its owner picks models and turns Provide Compute on.

The models a computer can run depend on its memory and the app's engines, so availability can differ between macOS, Windows and Linux.

06 / HOW OODA AI IS USED

From a request to an inspectable record.

Every request follows four steps. The result returns to the application while evidence and accounting are handled alongside it.

01

Submit a request

A user or application submits a prompt, document or task through the platform, a desktop app or an API. If the organisation has turned on PII cloaking, personal details in text are replaced with tokens first.

02

Execute on the chosen route

The request runs on its model's route: a cloud TEE model, another cloud model or an OODA AI Network computer the organisation allows, sealed to that computer's own key. Network models run only when someone picks one.

03

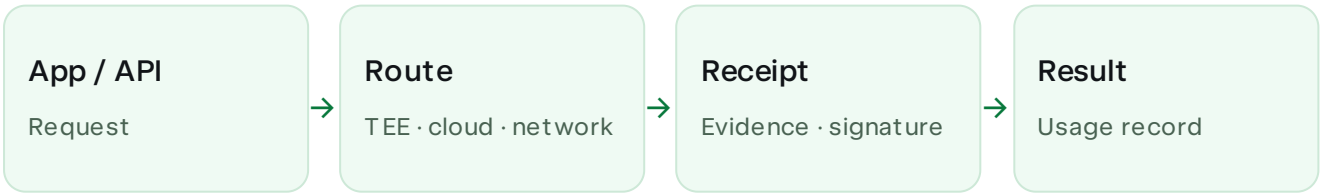
Record verification evidence

A receipt records the evidence available for that workload, and network answers are signed and verified. Authorised users can inspect status and any onchain reference.

04

Return the result and account for usage

The application receives its output. The server meters the consumed compute and supports settlement, including the serving computer's earnings.



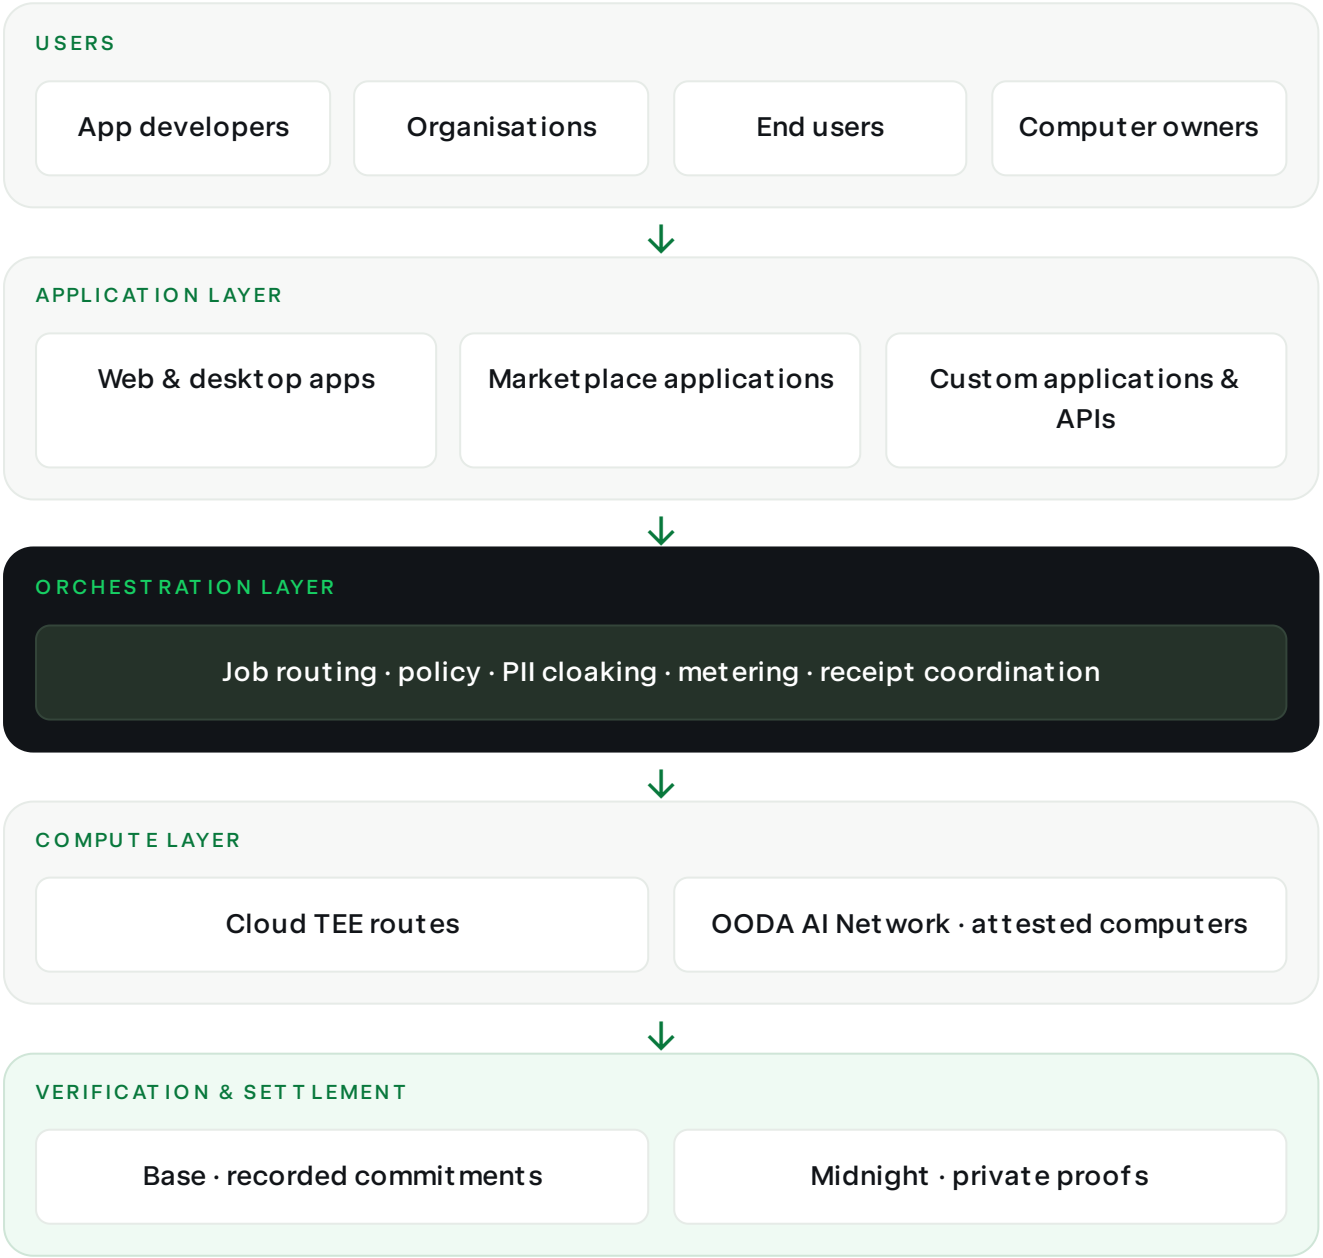
What the proof establishes

Evidence relates to the statement and execution boundary it covers. A receipt or signed answer does not, by itself, prove that an AI answer is correct. Only cloud TEE routes run inside a TEE; network computers are not TEEs.

07 / ARCHITECTURE OVERVIEW

Four layers. Clear responsibilities.

The network separates application use, orchestration, compute, and verification and settlement.



Network design. The OODA AI Network serves models and credits EUR earnings today; SOODA settlement and broader provider participation remain roadmap work. Verification integrations expose the evidence actually available for a receipt.

07 / CORE COMPONENTS

How the layers work together.

Compute routes

Cloud TEE routes execute supported workloads in a trusted execution environment. Computers of the OODA AI Network are not TEEs. Sealed requests, attestation, posture rules, signed answers and organisation rules protect them, as the following pages describe.

Orchestrators

Orchestrators connect requests to compute and manage receipt flows. Multiple orchestrators are part of the resilience design. Batching commitments in Merkle trees can reduce the number of onchain writes while retaining a path to individual records.

Verification and contracts

Receipt commitments and proof status support auditability. Smart contracts are part of the planned coordination, collateral and settlement system. Production addresses and supported mechanisms are disclosed when deployed.

Compute accounting

A metered service-credit balance makes usage understandable in currency terms. The server counts every request itself, for charges and for provider earnings. A balance is distinct from a transferable token balance, a purchase agreement or a provider's settlement record.

Privacy has a defined boundary

Hardware attestation, key handling, software integrity, access policy and metadata exposure all affect the result. Cryptographic evidence supports these controls; it does not replace operational security or organisational review.

07 / THE OODA AI NETWORK

Your computer.
Part of the network.

In the desktop app, a person picks models to serve and turns on Provide Compute. The computer becomes a node of the OODA AI Network.

Pick models
Downloaded to the computer

→

Provide Compute
The computer joins as a node

→

Serve requests
API and platform chat

→

Earn
For every verified request

Customers use these models like any other model on the platform: through the OpenAI-compatible API and the platform chat. The gateway balances requests across the computers serving each model.

LIVE SESSION IN THE APP

Requests	Tokens	Speed	Earned
Served this session	Prompt and answer	Tokens per second	This session

Your own organisation

By default an organisation's requests go only to its own computers, which must pass App Attest or the TPM check. That use only lowers its own costs.

Other organisations

Their computers answer only if an organisation allows them, within its rules. An owner can limit a computer to their own organisations.

Network computers are not TEEs

Apple hardware offers no trusted execution environment to third parties, and Windows and Linux computers are not used as one either. No network computer is described as a TEE; the platform's cloud TEE routes remain TEE. Network computers rely on the protections on the next two pages.

07 / NETWORK SECURITY

Sealed, attested and signed.

Each protection covers a different part of the path between the platform and a computer on the OODA AI Network.

Sealed requests

The platform handles each request, then seals it with HPKE to a key in the computer's Secure Enclave (Mac) or TPM (Windows and Linux). Relays and logs only see ciphertext.

Attested every five minutes

App Attest (macOS 27 and later) proves the genuine OODA app on a real Mac, not the Mac's security settings. TPM attestation (Windows and Linux) keeps the keys in the TPM, bound to Secure Boot (PCR 7) so they stop working if it is turned off or changed, and measures the boot from the TPM's log; it does not prove which program uses the keys. Checks repeat every 5 minutes; a failing computer stops serving organisations that require them at once.

Signed answers, public keys

Every answer is signed by the computer's key and verified by the gateway. The public endpoint /v1/network/attestation lists live computers' keys and posture, never people, organisations or hardware.

Posture rules

Computer	Required for attestation
macOS	Release build · no debugger · SIP on · Hardened Runtime · PT_DENY_ATTACH
Windows	Release build · no debugger · Secure Boot · no test signing or debugging · code integrity
Linux	Release build · no debugger · Secure Boot · ptrace restrictions

A missing posture value never satisfies a rule; a failing computer produces no attestation. From desktop 1.0.40, image, speech and music run in separate signed helper programs that receive each request over a pipe.

07 / PRIVACY ON THE NETWORK

Who may serve. What they receive.

Each organisation's admins decide which computers may serve its requests. The settings apply together.

Setting · default	What it decides
Own computers only · on	Only its own computers answer; off, others may too.
App Attest required · on	Macs need App Attest, its own Macs included.
TPM required · on	Windows and Linux need TPM attestation, its own included.
Europe only · on	Others' computers must connect from Europe (by IP address).

PII cloaking

Detect
Locally, in seven languages

→

Replace
With opaque tokens before leaving the platform

→

Restore
Only for the authorised recipient

Text only, with a stated limit

An organisation can turn on PII cloaking for supported text requests in English, German, Spanish, Portuguese, Swedish, Italian and French; image, speech and music prompts are not cloaked. Cloaking is pseudonymisation: detection can miss values.

RESIDUAL RISK

Why the organisation decides.

A Mac's owner who turns its security off could read prompts in its memory; a Windows or Linux computer's administrator or root can at any time. That is why each organisation decides which computers may answer: by default, only its own.

07 / TECHNOLOGY NETWORKS

Built around Base. Verified with Midnight.

Settlement and privacy-preserving evidence have different jobs in the architecture.

BASE

Settlement & receipts

Base is the planned SOODA delivery and settlement network. The platform can anchor receipt commitments and expose confirmed transaction references for inspection.

Base documentation ↗

MIDNIGHT

Private receipt proofs

Midnight integration supports proofs about authorised receipt commitments without placing prompt or response content in the proof. OODA displays the proof's status and associated evidence.

Midnight documentation ↗

Deployment status

Existing demonstration evidence includes Base Sepolia and Midnight preprod. Production network activation and SOODA delivery are separate milestones. A proof link identifies its actual environment.

These are technology roles. Neither network's name or logo represents an endorsement of a subscription or a guarantee of the project's outcome.

08 / WHO OODA AI IS FOR

One foundation. Several ways to build.

Enterprises & public sector

Teams handling sensitive business information need usable AI tools, configurable access and a record of how supported work is processed.

Document workflows | Internal knowledge | Process automation

Developers & AI builders

Developers connect applications to models, knowledge and workflows through the platform. Configurable and branded applications provide a route from an idea to a deployed service.

Custom applications | APIs | AI assistants

Compute providers & contributors

Anyone with a supported Mac, Windows or Linux computer can serve models to the OODA AI Network from the desktop app. Each verified request earns the computer's organisation EUR today; wallet-connected SOODA earnings are the next step.

Desktop apps | OODA AI Network | Earn per request

Sector examples describe intended use cases. Each organisation remains responsible for evaluating the suitability of its deployment and its own legal and professional obligations.

09 / SOODA UTILITY

Put SOODA to work.

COMPUTE REDEMPTION

1 SOODA = €1

of eligible OODA compute credit

Once redemption launches, each unlocked SOODA covers €1 of eligible OODA compute charges, including applicable VAT.

A service-credit value

The rate defines compute use within OODA. It gives no right to redeem SOODA for cash and does not guarantee a token market price.

No inactivity expiry

SOODA and unused redeemed compute credit do not expire merely because time passes or the account is inactive.

100 million

Planned maximum SOODA supply

Real activity

The basis for programme rewards

Service availability, delivery and applicable restrictions remain relevant. A future contractual allocation cannot be spent as SOODA before delivery and unlocking. SOODA does not represent shares in OODA AI AB or a right to company profits.

Reading this paper, receiving an invitation or making a private subscription does not itself mint tokens or activate compute redemption.

09 / COMPUTE CREDIT & SETTLEMENT

Clear usage accounting. Separate token settlement.

A compute balance expresses available service value. It is consumed as work is performed.

SOODA-to-compute redemption

Unlocked SOODA
Available for lawful redemption

→

OODA compute credit
€1 per SOODA, including VAT

→

Eligible compute
Credit consumed through usage

Redeemed credits are an internal, non-transferable service balance. Compute charges follow the applicable service tariff. Consumed credit is deducted from the balance, with a usage record showing the corresponding work.

Provider settlement

Request served
Tokens counted by the server

→

Answer verified
Finished and signed

→

Organisation credited
EUR today · SOODA next

Today each verified request credits the serving organisation in EUR. Next, wallet-connected providers receive SOODA from designated protocol balances, drawn from allocated tokens within the fixed supply. Any reserve replenishment requires its own operational rules.

> **Different records serve different purposes**
> A compute credit records service value; a token records SOODA ownership; a private agreement records contractual delivery rights. Payment for a future token allocation does not create an immediately spendable compute balance.

This design does not promise automatic market purchases, price support or automatic conversion of every payment asset. Production payment and settlement capabilities are specified in the relevant service terms.

09 / EARNING WITH COMPUTE

Serve models. Earn per request.

Every completed request a computer serves earns its organisation money: a EUR credit today and, as the next step, SOODA in real time for providers who connect a wallet.

TODAY · EUR

Credited per request

- Rates in EUR per million tokens, for input and output, set per model.
- Each verified request is credited to the organisation's balance.
- Use by your own organisation only lowers your own costs.
- Earnings from other organisations are paid out on request, from €100.

NEXT · SOODA

Earned in real time

- Connect a wallet and earn SOODA for every verified request as it completes.
- A better rate than the EUR credit, set in the network terms.
- No payout requests, no minimum and no threshold.
- Transfer SOODA to other wallets. Any wallet holding SOODA can use it on OODA AI.

What makes a request count

Counted by the server
From what it sent and received

→

Finished and signed
Signature verified by the gateway

→

Credited
To the serving organisation

Never on the computer's word

Tokens are counted by the server from what it actually sent and received. Unfinished, empty or unsigned answers earn nothing.

Wallet-connected SOODA earnings are the next step and not yet active; they depend on SOODA delivery and published network terms. SOODA earned this way redeems for compute like other unlocked SOODA. Provider earnings pay for work performed; they are not a return on money invested.

09 / NETWORK PARTICIPATION

Participation linked to useful work.

The network design includes collateral, service access and application participation. These mechanisms are separate from customer compute loyalty.

Provider collateral

Provider self-stake is a proposed addition to today's attestation-based admission. Rewards relate to accepted work; penalties require disclosed rules and evidence.

Buyer service access

The design explores token commitments for service benefits such as routing priority or higher limits. These benefits require a separately activated programme.

Developer participation

Application participation can include collateral, marketplace standards and rewards for activity. Ranking and reward rules must be disclosed before activation.

Delegation

Delegation models described in earlier papers remain a separate technical and legal design topic. This litepaper sets no delegated reward entitlement, rate or launch commitment.

DESIGN BOUNDARY

Service activity drives the reward model.

Holding SOODA alone earns no compute loyalty reward. No annualised or guaranteed investment return is set out in this paper.

The broader network mechanisms on this page remain roadmap features. The weekly customer-loyalty schedule on the following pages is the defined compute programme, effective when that programme launches.

09 / COMPUTE LOYALTY

The more you use, the more you earn.

Earn SOODA on qualifying compute usage. Higher portions of weekly usage earn a higher percentage back in future compute credit.

Calculated weekly | Accumulated monthly | Redeemed for compute

Weekly reward bands

Tier	Eligible weekly compute value	Rate on this band
01	First €100	1%
02	Above €100 to €500	2%
03	Above €500 to €2,000	3%
04	Above €2,000 to €10,000	4%
05	Above €10,000 to €20,000	5%
06	Above €20,000 to €50,000	6%
07	Above €50,000	7%

Progressive bands, not one rate on the whole bill

Each rate applies only to the part of weekly eligible compute value in that band. Entering a higher band does not change the rate on earlier portions.

€500 IN ONE WEEK

€100 × 1%		€400 × 2%		Total
1 SOODA	+	8 SOODA	=	9 SOODA

€9 of future compute credit · 1.8% effective credit back

These are the programme's defined bands for launch. Rewards become available after programme activation and validation of eligible usage. They are service rewards, not a return on money invested.

09 / THE LOYALTY MATRIX

See exactly what usage earns.

The effective credit-back rate rises as more weekly compute falls into the higher bands.

Weekly compute	SOODA earned	Future credit	Effective credit back
€100	1	€1	1%
€500	9	€9	1.8%
€2,000	54	€54	2.7%
€10,000	374	€374	3.74%
€20,000	874	€874	4.37%
€50,000	2,674	€2,674	5.348%
€100,000	6,174	€6,174	6.174%

2,674 SOODA

At €50,000 of weekly eligible compute

6,174 SOODA

At €100,000 of weekly eligible compute

Reading the €100,000 row

The first €50,000 earns 2,674 SOODA. The next €50,000 earns 7%, or 3,500 SOODA. Together, that is 6,174 SOODA: €6,174 of future compute credit and a 6.174% effective credit-back rate.

The table uses eligible compute value after discounts and excluding tax. The credit shown can be used against future eligible charges once distributed and redeemable. It is not an immediate reduction of the invoice that generated the reward.

Values are calculated from the weekly bands, without token-price or foreign-exchange assumptions. Effective percentages are shown to three decimal places where needed. Model-token volumes are not used as a universal reward unit.

09 / WEEKLY USAGE, MONTHLY REWARDS

Every week counts. Rewards add up.

Add each completed week's SOODA to the monthly total.

Completed week	Eligible compute	SOODA earned
Week 1	€100	1
Week 2	€500	9
Week 3	€2,000	54
Week 4	€10,000	374
Month total	€12,600	438

MONTHLY REWARD

438 SOODA

€438 of future compute credit

3.48%

Effective credit back on €12,600 eligible compute

One defined week

Weeks run Monday–Sunday in UTC. Each completed week belongs to the month in which its Sunday falls. Weekly bands reset; earned SOODA carries forward.

Monthly distribution

Monthly rewards are the sum of those weekly awards, distributed after month-end validation. A month can contain four or five qualifying week endings.

What qualifies

Settled compute charges after discounts, excluding tax. Purchased-SOODA payments qualify at the redeemed service value, excluding tax. Free, promotional or reward-funded usage, refunds and duplicate activity do not qualify.

Distributed loyalty rewards are usable for compute without the private-subscription vesting schedule. Programme changes apply prospectively and do not cancel awards already validly earned. No reward is generated by repeatedly spending reward-funded credit.

09 / WAYS TO EARN SOODA

Contribute. Use. Build. Earn.

SOODA rewards recognise qualifying activity across the platform and its community programmes.

Compute loyalty

Earn from settled, eligible compute usage under the weekly bands and monthly accumulation rules.

Marketing campaigns

Complete eligible campaign activities under the published award and validation rules.

Customer referrals

Refer qualifying platform customers through an activated referral or affiliate programme.

Loyalty programmes

Participate in defined customer engagement programmes with stated eligibility and rewards.

Compute provision

Serve models from your computer: EUR per verified request today, SOODA with a connected wallet next.

Surveys & feedback

Contribute validated feedback or research participation when a programme offers SOODA rewards.

Clear rules for each programme

Each programme publishes its eligible actions, award basis and validation process before activation. The compute matrix applies to compute loyalty; it does not set the award rate for every other activity.

Rewards use allocated SOODA within the fixed supply. Customer referral and campaign rewards do not authorise commissions for introducing token investors or raising investment capital.

10 / ROADMAP

From working foundations to a broader network.

Progress is organised around technical and operational readiness. No trading or unconditional token-delivery date is set here.

0

Platform, apps and network core

The platform, the desktop apps and the OODA AI Network operate today, with sealed requests, attestation, signed answers and EUR earnings. Private compute receipts and evidence inspection are validated in defined test environments.

1

Provider readiness

Broaden attested provider coverage, quality measurement and payout operations. Validate security assumptions and settlement rules; complete production testing and operational procedures.

2

SOODA utility activation

Complete lawful token delivery, operational compute redemption and applicable documentation. Activate wallet-connected SOODA earnings for providers, the compute loyalty programme and validated reward distribution.

3

Ecosystem expansion

Expand supported providers, applications, orchestration and contribution programmes. Introduce further mechanisms only with their own published rules.

4

Network maturity

Improve privacy controls, reporting, resilience and hardware coverage as adoption grows. Extend the network to more supported workloads and organisations.

The roadmap describes development direction. Sequence and scope may change with readiness, testing and applicable requirements. Existing contractual rights remain governed by the relevant agreement.

DOCUMENT NOTES & REFERENCES

Scope and technical references.

Project information

This litepaper is a project and technology overview. It is not an executed subscription agreement, a payment instruction or a statement of regulatory approval. Private terms govern individual contractual rights.

Availability & risk

Development, delivery and redemption depend on technical, operational and legal readiness. Tokens and service credits depend on platform continuity; no cash value, resale market or financial return is guaranteed.

Limits of the network's protections

Network computers are not TEEs. App Attest does not prove a Mac's security settings, and TPM attestation does not prove which program uses the keys. PII cloaking is pseudonymisation that can miss values. Provider earnings depend on completed, verified requests.

References

1. Base documentation
   Network architecture, transactions and the settlement infrastructure described in this paper.
2. Midnight documentation
   Privacy-preserving applications and the proof infrastructure described in this paper.

Version 3.8

This edition describes the platform, the desktop apps, the OODA AI Network and its security model, provider earnings, the SOODA compute-credit model and participation programmes. The loyalty matrix defines weekly reward bands and monthly accumulation. Feature availability is identified separately from the development roadmap.

OODA AI

OODA AI AB (publ)
Organisation no. 556737-5489
Grev Turegatan 30 · 114 38 Stockholm · Sweden
ooda.ai