Solrouter

Solrouter

Search
⌘K

Open Chat

Get Started

Introduction Quickstart Chat App

Core Concepts

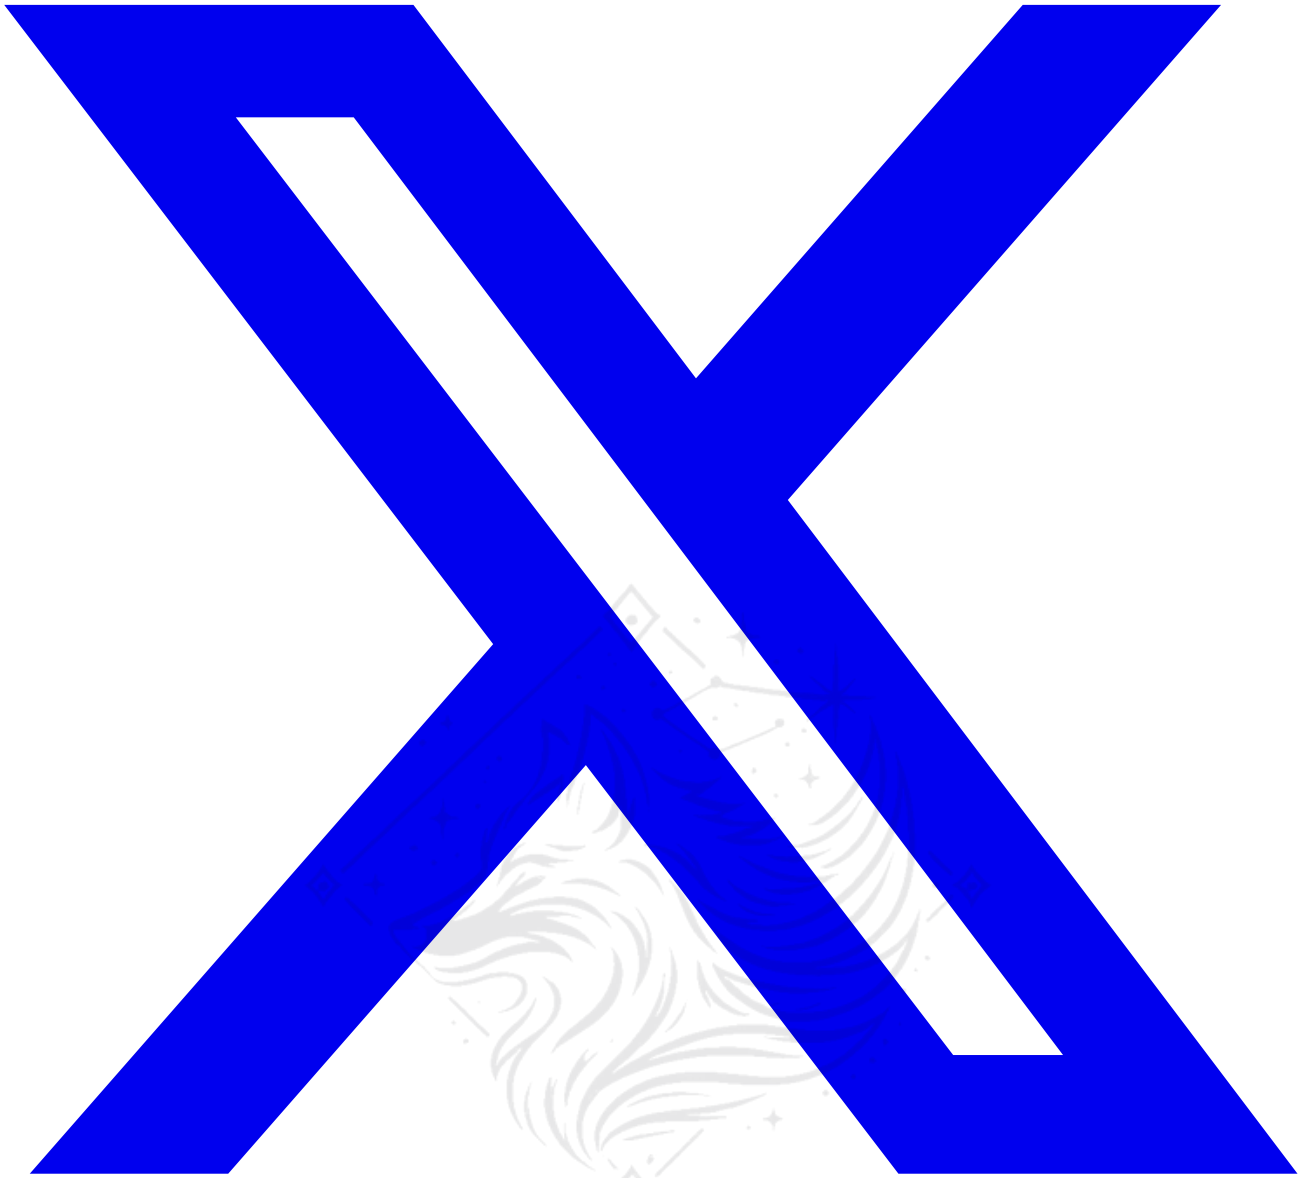
Development

Payments

API Reference

CRYP2NOVA

Archived copy · cryp2nova.com



CRYP2NOVA

Archived copy · cryp2nova.com



Introduction

Introduction

Solrouter is a cryptographically private AI layer — prompts are encrypted client-side and processed in an Intel TDX enclave, so no one can read your data.

Copy Markdown Open

Archived copy · cryp2nova.com

Solrouter is a cryptographically private AI infrastructure layer for Solana developers. The idea is simple: your prompts, documents, and responses should never exist in plaintext anywhere we could read them. So they don't — not on the wire, not on our backend, and not in any log file. Plaintext lives only inside a hardware-isolated enclave. Privacy here is enforced by math, not by a promise.

Why Solrouter?

Start with the problem. Today, every AI request you send is something the provider can read, store, and analyze.

Mainstream providers keep your prompts indefinitely, ask for personal details to sign up, and back their privacy claims with a terms-of-service page. Nothing technical stops them from reading what you send — you're trusting a policy.

Solrouter removes the need for that trust. We built the system so that **we cannot see your data** in the first place.

Here's how. Your prompt is encrypted on your device before it leaves your browser or app. Our backend only ever receives an opaque encrypted blob and forwards it, blindly, to an Intel TDX Confidential VM — a TEE (Trusted Execution Environment: hardware that isolates code and data so even the machine's operator can't inspect it at runtime). Plaintext appears only inside that attested enclave, which even we cannot read.



No third-party model providers

In privacy mode, Solrouter does not use OpenAI, Anthropic, Google, or any third-party model provider. All private inference runs on self-hosted, open-weight models on the Nosana decentralized GPU network. Your prompts never reach an external model API.

Products

Pick the surface that fits how you work — an SDK to embed in your app, typed agent tools, an MCP server for your editor, or the hosted chat app. They all run on the same private backend.

Archived copy · cryp2nova.com

Privacy SDK

`@solrouter/sdk` — integrate end-to-end encrypted AI into your app. The SDK handles key exchange, client-side encryption with Arcium RescueCipher, and response decryption for you. No email or KYC — connect a Solana wallet and start building.

Agent Tools SDK

`@solrouter/agent-tools` — typed tools for the Agent Privacy API with a Vercel AI SDK adapter. Quote, execute, and settle privacy-preserving swaps and encrypted inference without writing orchestration boilerplate.

MCP Server

`@solrouter/mcp-server` — use Solrouter from Claude Desktop, Cursor, or any MCP-compatible client. Encrypted chat, private token research, and wallet analysis without writing any code.

Chat App

Multi-model encrypted chat with file attachments, image and video generation, and a RAG knowledge base. No account needed beyond a connected Solana wallet.

Key Guarantees

These are the promises that hold for any privacy-mode request, no matter which product you reach for. Each one is a property the system enforces, not a feature you have to remember to turn on.

Client-Side Encryption

Your prompt is encrypted on your own device — using Arcium's RescueCipher with X25519 key exchange — before it leaves your browser or application.

Archived copy · cryp2nova.com

TEE-Isolated Inference

Plaintext exists only inside an Intel TDX Confidential VM. No host process — including Solrouter's own backend — can read enclave memory at runtime.



Verifiable Attestation

Don't take our word for it. Every TEE response carries an Intel-signed TDX quote, so you can independently verify the enclave's public key and the exact code running inside it.



No KYC or Email Required

Connect a Solana wallet, generate an API key, and start building. You pay per call in USDC or `$ROUTER` from a prepaid balance — no signup form, no PII.



Open-Weight Models Only

All privacy-mode inference runs on self-hosted open-weight models (`gpt-oss:20b` and `qwen3:8b`) on the Nosana decentralized GPU network.



On-Chain Attestation Anchor

The Solrouter attestation program is deployed on Solana mainnet. Each privacy-mode session can publish a PDA that links the request to the verified enclave on-chain.



Quickstart

Install @solrouter/sdk, generate an API key with your Solana wallet, and send your first encrypted AI request with just a few lines of TypeScript.

On this page

Why Solrouter?ProductsKey Guarantees

CRYP2NOVA

Archived copy · cryp2nova.com