Ninja Whitepaper ver.0426 — Basic Edition (English)

1. Abstract

Decentralized finance (DeFi) has grown rapidly over the past several years, with Total Value Locked (TVL) exceeding $130 billion at its early-2026 peak (at all-time highs in ETH-denominated terms). At the same time, "Agentic Finance" — where AI agents autonomously execute transactions — is emerging at speed, fundamentally reshaping the participant structure of DeFi. Institutional investors, retail users, and AI Agents: a diverse set of actors is flowing in, and a new financial battleground is forming on-chain.

Behind this growth, however, hacking losses continue at a scale of $3.4 billion per year (Chainalysis, 2025). The root cause lies in the paradoxical structure of DeFi itself.

Blockchain was designed as "transparent finance where all transactions are public." In reality, however, contract bytecode is unreadable to humans, the internal allocation structures of Vaults are invisible, and transaction calldata appears as nothing more than strings of cryptographic data. **A financial system that proclaims transparency has become the most opaque financial system** — this is the DeFi paradox.

Agentic Finance deepens this paradox further. AI Agents make autonomous decisions and execute transactions, yet "what they decided and why they chose that transaction" is invisible from the outside. Even when humans delegate asset management to an Agent, there is no way to monitor the delegate's actions. **Agents that proclaim autonomy have become the least accountable actors** — this is the Agentic Finance paradox.

Ninja is the Agentic Security Controlplane that resolves this dual paradox. Ninja Intelligence Core provides contract trustworthiness assessment, transaction intent analysis, and full position visibility through three Intelligence domains (Entity / Action / Position). All Intelligence follows a two-layer architecture: deterministic Layer 1 (establishing facts) and LLM-augmented Layer 2 (translating facts into human language), achieving both reproducibility and explainability. Furthermore, ShoGun (Govern / Block) executes policy control and automated actions based on Intelligence, delivering an end-to-end solution that goes beyond "knowing" to "blocking."

Ninja is not a "DeFi transparency tool." It is the transparency and security infrastructure for DeFi and Agentic Finance.

This whitepaper describes Ninja's technical architecture, the large-scale EVM contract analysis database "not-so-smart-contracts" (covering over 90 million contract deployments and

comprehensively analyzing 15+ million unique deduplicated bytecodes), product roadmap, business model, and competitive landscape.

2. Problem — Structural Opacity in DeFi and Agentic Finance

DeFi's growth is remarkable. ETH-denominated TVL has reached all-time highs, and emerging protocols such as Morpho have surpassed 1.4 million users, achieving rapid growth. Simultaneously, the DeFAI (DeFi + AI) sector is expanding rapidly, ushering in the age of Agentic Finance where AI agents autonomously transact on-chain.

Yet beneath this growth, four structural problems remain unsolved.

1. Contract trustworthiness is unknown When users access a DeFi protocol, the means to determine whether its contracts are safe are extremely limited. Whether a contract is verified, whether it uses a proxy structure, whether it appears on the OFAC list, whether its bytecode characteristics resemble known exploit patterns — it is unrealistic to expect individuals to verify each of these factors.

2. Transaction intent is opaque Users cannot accurately understand what a transaction they are about to sign actually does. Calldata analysis, chains of contract calls, the reasonableness of approval amounts — these require specialized knowledge to evaluate, making them breeding grounds for phishing and approval scams.

3. No holistic view of positions Tools that display wallet balances exist. However, no product provides an integrated view of Vault internal allocation structures, the health of dependent protocols, and the overall risk profile of assets. Users operate their assets without knowing the full picture.

4. AI Agent behavior is unmonitorable The era where AI Agents autonomously interact with DeFi protocols and execute transactions has begun. Yet there is no way for the delegating human to know which contracts an Agent accesses, what decision criteria it uses to execute transactions, or what risks it takes. In the Resolv incident of March 2026, a compromised key was used to mint approximately 80 million USR without collateral, resulting in losses of approximately $25 million. Secondary damage was amplified as automated scripts in adjacent protocols continued to operate even after the depeg. Cases where Agent autonomy becomes a security vulnerability are now a reality.

The combination of these problems sustains $3.4 billion in annual hacking losses, countless phishing scams, and

persistent barriers to DeFi entry for institutional investors and AI Agent operators.

3. Vision

Ninja's vision is to build the security and transparency Controlplane for DeFi and Agentic Finance.

"Controlplane" does not mean a tool that provides individual security functions. It means a platform that centrally governs security decisions, policies, and responses. Both users and AI Agents can know the counterparty (Entity), read the action (Action), track assets (Position), and on that basis define policies and execute automated actions through ShoGun.

Ninja scales from the everyday safety of individual DeFi users to monitoring AI Agent behavior to serving as the security infrastructure for institutional investors and protocols. It delivers not only "visibility" into DeFi but also "implementation of accountability" in Agentic Finance.

Continuity from DeFi to Agentic Finance DeFi has been shunned by regulators due to "ambiguity of responsible parties." Protocol operators, frontend providers, individual users — who bears responsibility for what has been technically difficult to determine. Paradoxically, however, the arrival of AI Agents is driving this problem toward resolution. Deterministic action logs generated by AI Agents, combined with the explainability afforded by the Layer 1/Layer 2 two-layer architecture, create a technical foundation to prove "who made what decision and why." DeFi is part of the same continuum evolving toward Agentic Finance, returning to the essence of finance — verifiability, accountability, and transparency. Ninja is the infrastructure that supports this transition.

---

4. Architecture — Overall Structure of Ninja

Ninja's architecture consists of three layers.

```
Ninja (Agentic Security Controlplane)
│
├── Ninja Intelligence Core (Know)
│   ├── Entity Intelligence  — Know the counterparty
│   ├── Action Intelligence  — Read the transaction
│   ├── Position Intelligence — Track assets
│   └── Learning             — Evolve with data
│
├── Ninja Delivery Layer (Deliver)
│   ├── NinjaScan (Telegram Bot)
│   ├── Dashboard
│   ├── API
│   └── MCP
│
└── ShoGun (Govern / Block)
    ├── Policy Engine
    ├── Whitelist Management
    ├── Approval Flows
    └── Automated Actions
```

Ninja Intelligence Core is the intelligence engine at the heart of the Ninja platform. Through three Intelligence domains (Entity / Action / Position) and Learning, it generates the Intelligence that forms the basis of all decisions. This separation ensures that Ninja's analysis results are consistently accessible from all channels — NinjaScan, Dashboard, API, and MCP — and also function as triggers for ShoGun's policy engine and automated actions. It provides the same Intelligence to AI Agent calls as it does to human users.

Ninja Delivery Layer (Deliver) is the interface that delivers Ninja Intelligence Core's Intelligence to users. On Day 1, it provides on-demand analysis through NinjaScan (Telegram Bot), with Dashboard and API additions from Day 2 onward. MCP (Model Context Protocol) is the standard interface through which AI agents reference Ninja as a Single Source of Truth (SSOT), offered as part of the Delivery Layer.

ShoGun (Govern / Block) handles the definition and execution of security policies based on Intelligence. It includes whitelist management, approval flows, and automated actions (alert issuance, position freeze proposals, etc.). In the future, it will expand to cover AI Agent behavior policy definitions as well.

5. Ninja Intelligence Core — 3 Intelligence Domains + Layer 1 / Layer 2

5.1 Design Principle: Two-Layer Architecture

All Intelligence in Ninja strictly follows the two-layer architecture below.

Layer 1 (Deterministic Layer)

- Establishes facts. 100% reproducible. Does not use LLMs
- Always returns identical output for identical input
- Based on on-chain data, bytecode analysis, known pattern matching, etc.
- The output of this layer forms the foundation of all decisions

Layer 2 (LLM-Augmented Layer)

- Transforms Layer 1 output into a form easily understood by humans
- Supplements and explains facts but does not override Layer 1 determinations
- What the LLM generates is "explanation," not "judgment"

This design secures reproducibility and explainability — the most critical properties for a security product. LLM outputs are probabilistic and can return different results for identical inputs. Grounding security judgments on LLMs not only causes misclassifications but also makes it difficult to reproduce the basis of a determination during audits or incident response. Layer 1 establishes facts; Layer 2 merely translates those facts into a human-readable form — this principle is the foundation of Ninja's trustworthiness.

This two-layer architecture is also critically important for monitoring AI Agent behavior. If Agent transactions are merely "interpreted" after the fact by an LLM, the interpretation itself becomes probabilistic and non-reproducible. Layer 1 deterministically establishes the factual structure of transactions, and Layer 2 translates those facts into an auditable form for humans — this architecture is the technical foundation for implementing Agent accountability.

5.2 Entity Intelligence — Know the Counterparty

Entity Intelligence provides multi-dimensional trustworthiness assessment of contracts and addresses.

Layer 1 evaluation criteria:

- Bytecode ML analysis: Extracts features from contract bytecode and detects malicious patterns using machine learning models
- Risk Score computation: Scoring that integrates multiple evaluation axes with weighted aggregation

- Blacklist/OFAC matching: Cross-referencing against sanctions lists and known malicious addresses
- Contract attribute verification: Verified status, proxy structure, time elapsed since deployment, etc.

Layer 2 augmentation:

- Natural language summary generation of evaluation results
- Prioritized explanation of risk factors
- Comparative context with similar contracts

Operating modes:

- **On-demand:** Instant evaluation initiated by users via the `/check` command
- **Continuous:** Automatic re-evaluation when changes (upgrades, ownership transfers, etc.) are detected in dependent contracts

5.3 Action Intelligence — Read the Transaction

Action Intelligence analyzes transaction calldata to determine "what the transaction is trying to do." It processes not only transactions that humans sign but also transactions that AI Agents autonomously generate and execute through the same analysis pipeline.

Layer 1 analysis pipeline:

```
ParseCheck
  → TxIntentMapper (intent classification)
    → ContractArtifactResolver (contract information resolution)
      → Parallel Checks (multiple verifications executed concurrently)
        → Layer 1 Output (structured facts)
          → LLM Augmentation (Layer 2)
```

1. **ParseCheck:** Syntactic parsing and basic validation of transaction data
2. **TxIntentMapper:** Mapping from calldata function selectors to intent categories (13 types)
3. **ContractArtifactResolver:** Resolution of target contract ABI, source code, and metadata
4. **Parallel Checks:** Concurrent execution of multiple verification logics (approval amount validation, known exploit signature matching, reentrancy checks, etc.)
5. **Layer 1 Output:** Structured output of categories (13 types), alert codes (17 types), and traffic light signals (SAFE / WARNING / DANGER)
6. **LLM Augmentation:** Transforms Layer 1 structured output into a human-readable analysis report

12 category examples: Transfer, Swap, Approve, Deposit, Withdraw, Claim, Stake, Bridge, Governance, Contract Deployment, Multi-call, Unknown

Traffic light system:

- **SAFE:** Matches known patterns with no anomalies
- **WARNING:** Contains elements requiring attention (high-value approvals, unverified contracts, etc.)
- **DANGER:** High-risk elements detected (known exploit signatures, association with sanctioned addresses, etc.)

Operating modes:

- **On-demand:** Instant analysis initiated by users via the `/scan` command
- **Continuous:** Automatic detection and analysis of transactions from monitored wallets (including AI Agent wallets)

5.4 Position Intelligence — Track Assets

Position Intelligence detects all positions in a wallet and provides a holistic view including Vault internal allocation structures.

Key features:

- Full position detection: Comprehensive detection of tokens, LP positions, Vault shares, staking positions, etc. held by the wallet
- Vault allocation mapping: Structural analysis of how user assets within a Vault are actually allocated across protocols and pools
- Anomaly detection: Detection of sudden position fluctuations, Vault strategy changes, dependent protocol anomalies, etc.

Operating modes:

- **On-demand:** Full position scan at wallet registration
- **Continuous:** Periodic sweep for change detection and alerting

5.5 Learning — Evolve with Data

Ninja evolves continuously through three mechanisms.

1. **User reports → Ground Truth → Model retraining:** Feedback from users (false positive reports, suspicious transaction reports, etc.) is accumulated as Ground Truth and used to improve detection model accuracy

2. **Alert code accumulation:** Each time a new incident occurs, its pattern is analyzed and added as a new detection rule (alert code). The scope of detection expands continuously over time
3. **External knowledge absorption:** Continuous ingestion of information from external data sources such as Forta Network, GoPlus Security, and DeFi Llama to update Ninja's knowledge base

---

6. not-so-smart-contracts — Large-Scale EVM Contract Analysis Database

Ninja maintains a proprietary large-scale EVM contract analysis database called "not-so-smart-contracts." Specifically, it **covers over 90 million contract deployments and comprehensively analyzes 15+ million unique bytecodes after deduplication** (i.e., the actual analysis target count after consolidating re-deployments of identical bytecode). The name "not-so-smart-contracts" is a tribute to the eponymous smart contract vulnerability pattern collection published by Trail of Bits (Crytic) at https://github.com/crytic/not-so-smart-contracts. While Trail of Bits' repository serves as an educational resource for vulnerability pattern study, Ninja's not-so-smart-contracts is a production database that comprehensively analyzes EVM contracts at scale — differing in both purpose and magnitude. This database strengthens Ninja Intelligence Core in two directions.

6.1 Action Intelligence Enhancement: selector_mappings.parquet

Function selectors and method name mappings extracted from the analyzed contract corpus are stored in `selector_mappings.parquet` and integrated into the `SELECTOR_METHOD_MAP` referenced by Ninja Intelligence Core's Action Intelligence.

This enables high-accuracy resolution of method names from function selectors even for transactions involving unverified contracts whose ABIs are not publicly available. Intent classification of transactions that were previously categorized as "Unknown" has improved significantly, expanding Action Intelligence coverage.

6.2 Entity Intelligence Enhancement: contract_embeddings.npy

Features are extracted from each contract's bytecode and stored as vector embeddings in `contract_embeddings.npy`.

This enables similarity comparisons between new contract bytecode and the corpus of known contracts. For example,

determinations such as "87% similar to the corpus of past exploit-targeted contracts" become possible, providing quantitative evidence for Entity Intelligence risk assessments. Even for unknown contracts, risks can be detected early through structural similarity to past exploit patterns.

7. Product & Customer Roadmap — Day 1 to Day 5

Ninja's product is released in five stages from Day 1 to Day 5. Each Day progressively expands target customers, channels, Intelligence domains, operating modes, and pricing models.

Day 1: Foundation — B2C Customer Base Acquisition

- **Target customers:** B2C individuals
- **Channel:** NinjaScan (Telegram Bot)
- **Intelligence:** Entity Intelligence + Action Intelligence
- **Commands:** `/scan` (Action analysis — implemented and publicly launched) / `/check` (Entity evaluation — to be added shortly after the public launch)
- **Operating mode:** On-demand (user-initiated)
- **Pricing:** Free tier only (Free: 3 uses/day). The primary objective of Day 1 is customer base acquisition; monetization begins in earnest from the next Day
- **Parallel initiative:** Secure 1-2 B2B pilot customers from Day 1 to prepare for B2B expansion from Day 2 onward

The aim of Day 1 is to build an individual user base as rapidly as possible through NinjaScan, a low-friction channel. Technical differentiation is driven not by LLMs but by ML accuracy — the benchmark precision of the Ninja Risk Engine — allowing users to experience Ninja's core value for free across the Entity and Action axes.

Day 2: Continuous Monitoring — B2C Passive Monitoring + Monetization Launch

- **Target customers:** B2C individuals (continued)
- **Channel addition:** Dashboard
- **Intelligence addition:** Position Intelligence + early Learning (off-chain report data integration)

- **Operating mode addition:** Continuous monitoring (passive alerts incorporating external and report data)
- **Pricing:** Full launch. Three plans: Free / Lite / Pro. Higher plans expand usage limits and features
- **Prerequisite:** Willingness-to-pay validation from Day 1 is complete, and B2B pilots are proceeding in parallel

Day 2 is the inflection point from On-demand to Continuous and the phase where B2C monetization begins in earnest. The Dashboard addition enables visual representation of the full Position picture, and the early Learning version incorporating external and report data begins operation.

Day 3: Protocol Expansion — B2B API + ShoGun + MCP Support

- **Target customers:** DeFi protocol operators
- **Channel addition:** B2B API (three-tier structure: Read / Monitor / Control APIs)
- **Feature additions:**
  - Protocol-specific monitoring items (flash loan countermeasures, Vault strategy deviation detection, governance proposal analysis)
  - ShoGun (Govern / Block) initial release (policy engine, whitelists, approval flows)
  - **MCP (Model Context Protocol) support** — The standard interface through which AI agents reference Ninja as a SSOT. Provided as part of the Ninja Delivery Layer
- **Intelligence:** Full-scale Learning operation

Day 3 is the phase where Ninja expands into a B2B platform. DeFi protocol operations teams are the primary customers, and it addresses protocol-specific monitoring requirements. MCP support lays the groundwork for the Day 5 AI-to-AI Controlplane.

Day 4: Institutional Control — ShoGun Expansion

- **Target customers:** CEXs, institutional investors, family offices, compliance teams
- **Feature additions:**
  - Custom governance policy DSL
  - Extended automated actions
  - Automated compliance report generation
- **Channels:** Institutional-grade dedicated Dashboard, SLA-backed API line

In Day 4, ShoGun (Govern / Block) is expanded for institutional use. Custom governance rules can be defined via DSL, enabling automated actions and reporting aligned with compliance requirements.

Day 5: AI-to-AI Controlplane

- **Target customers:** AI Agent providers, AI-to-AI ecosystem
- **Features:** Advancing beyond Day 3's MCP support (AI → Ninja reference) to provide **policy mediation, contract verification, and incident isolation between AI Agents**
  - Inter-Agent transaction policy agreement protocol
  - Cryptographic proof of action logs
  - Automated anomalous Agent isolation
- **Positioning:** The completed form of trust infrastructure for the Agentic Finance era

Day 5 is the completed form of Ninja's long-term vision. In an era where AI Agents autonomously transact and contract with each other, Ninja functions as the policy mediator, contract verifier, and incident isolator.

Integrated Roadmap Table

Phase	Primary Customers	Channels	Intelligence	Operating Mode	Pricing
Day 1 Foundation	B2C Individuals	NinjaScan	Entity + Action	On-demand	Free tier only
Day 2 Continuous	B2C Individuals	+ Dashboard	+ Position + Early Learning	+ Continuous (external data integration)	Free / Lite / Pro
Day 3 Protocol	DeFi Protocols	+ B2B API + MCP	+ Full-scale Learning	+ ShoGun initial	B2B API pricing
Day 4 Institutional	CEXs / Institutions	+ Dedicated Dashboard	(same as above)	+ ShoGun expansion (Governance DSL)	Enterprise pricing
Day 5 AI2AI	AI Agent Providers	+ Inter-Agent Protocol	(same as above)	+ AI-to-AI Mediation	Agent-based pricing

--- ## 8. Competitive Landscape Multiple existing players operate in the DeFi security, portfolio visualization, and Agentic Finance security spaces. This chapter classifies them into three archetypes based on the nature of their capabilities and visualizes the differences from Ninja.

Archetype Definitions

- **Security Alert:** Archetypes specializing in threat detection and notification without Control capabilities. Includes Blockaid, GoPlus, Hypernative, Forta, and others
- **Asset Visualization:** Archetypes specializing in position display without providing security assessments. Includes Zapper, Zerion, DeBank, Exponential.fi, and others
- **Control Layer:** Archetypes that extend into policy execution and automated control. Includes DeFi Saver (partial) and **Ninja**

Feature Comparison Matrix

Archetype	Service	Entity	Action	Position	LLM Explanation	Alerts
Security Alert	Blockaid	△	○	x	x	x
Security Alert	GoPlus	○	△	x	x	△
Security Alert	Hypernative	△	○	x	x	○
Security Alert	Forta	△	○	x	x	○
Asset Viz	Zapper	x	x	△	x	x
Asset Viz	Zerion	x	x	△	x	x
Asset Viz	DeBank	x	x	△	x	△
Asset Viz	Exponential.fi	x	x	○	△	x
Control Layer	DeFi Saver	x	x	△	x	○
Control Layer	**Ninja**	○	○	○	○	○

○ = Offered as a core feature △ = Partially supported x = Not supported

Analysis by Archetype and Differentiation

Security Alert archetype: Excels at "alerting to danger" but lacks a Control Layer. Blockaid (cumulative funding $83M) and Hypernative (cumulative $68M) enjoy broad adoption backed by substantial capital, but the absence of a Control Layer and Position Intelligence represents a structural gap versus Ninja.

Asset Visualization archetype: Excels at visualizing "what you own" but does not provide security assessments. Zapper/Zerion/DeBank/Exponential.fi specialize in position visibility and do not address contract trustworthiness evaluation or transaction intent analysis.

Control Layer archetype: Very few major players exist, and **an AI-ready Control Layer remains an emerging segment still in the demand-validation phase**. Ninja positions itself as one of the early players that vertically integrate all three archetypes (Security Alert / Asset Visualization / Control Layer) in a See → Detect → Control flow while explicitly including AI-readiness in its roadmap (Note: these assessments reflect our independent analysis based on publicly available information and sales-call intelligence as of March–April 2026).

Ninja's Differentiation Points

- Vertical integration of three archetypes (Alert / Visualization / Control Layer)
- Reproducibility and explainability through the deterministic Layer 1 + LLM-augmented Layer 2 two-layer architecture
- AI-focused roadmap extending to Day 5 (MCP support → AI-to-AI Controlplane)

Note

The AI-ready Control Layer market is currently in the demand validation phase. Ninja adopts a staged expansion strategy from B2C → B2B Protocols → AI, making investment decisions for each subsequent phase based on validation results from the current phase.

---

9. Business Model

Ninja's business model consists of two pillars: B2C freemium and B2B API.

B2C (Individual Users)

Three plans: Free / Lite / Pro. Higher plans expand usage limits and features.

Plan	Details
Free	`/check` 3 uses/day + `/scan` 3 uses/day
Lite	`/check` expanded daily uses + `/scan` expanded daily uses
Pro	Unlimited + priority response + detailed report output

Day 1 focuses exclusively on customer base acquisition with the free plan only; Lite/Pro are fully launched from Day 2 to begin revenue generation. The Free tier continues to function as a user acquisition channel beyond Day 1.

B2B API

Three tiers: Read API / Monitor API / Control API. Higher tiers expand functional scope and pricing scales accordingly. Price ranges are benchmarked against published and estimated competitor pricing (Hypernative $3K-$20K, Blockaid $5K-$50K, GoPlus Enterprise $500-$3K, Forta $100-$1K) and finalized through agreement with pilot customers.

Plan	Features	Target Categories (Examples)
Read API	Entity / Action Intelligence reference	Curators, wallet providers
Monitor API	+ Continuous monitoring + alert delivery + protocol-specific monitoring items	DeFi protocols, AI Agent providers
Control API	+ ShoGun policy engine + automated actions + Governance DSL	Institutional investors, compliance teams, CEXs

*Enterprise full-suite offerings (MCP integration, SLA guarantees, dedicated support) are handled on a case-by-case basis. --- ## 10. Team ### Core Team - **CEO Kaneshiro:** 13 years at Accenture. Led numerous AI projects for financial institutions. End-to-end experience from strategy formulation to technology delivery - **CTO Murakami:** Technical leader in large-scale system development. Deep expertise in distributed system design and performance engineering - **Megami:** Hands-on experience at a crypto exchange and Fireblocks (institutional crypto infrastructure). Practical knowledge of security operations - **Kai Otsuki:** General Co-Chair, B4TI (IEEE International Workshop on Blockchain for Decentralized Trust and Digital Identity). Contributor to the Ethereum Foundation and participant in IETF SD-JWT (RFC 9901) standardization - **CMO Suganuma:** Trilingual capabilities (Japanese, English, Chinese). Global marketing and community building

Advisors

- ICANN DNSSEC key manager (Trusted Community Representative). One of a select few security experts worldwide who participate in the most critical cryptographic key management ceremonies for internet infrastructure (name withheld)

---

11. Market Environment and Regulatory Trends

Market Data

- **DeFi TVL:** Over $130B (early-2026 peak; at all-time highs in ETH-denominated terms, source: DefiLlama)
- **Web3 Security Market:** $2.9B (2025) → $15.8B (2032 projected, CAGR 26.4%)
- **Annual Hacking Losses:** $3.4B (2025, source: Chainalysis)
- **DeFAI (DeFi + AI) Sector:** A rapidly growing sector. On-chain transactions by AI agents are forming a new market category
- **Morpho:** Over 1.4 million users. Emblematic of the rapid expansion of Vault-model DeFi
- **Resolv Incident:** March 2026; compromised key used to mint approximately 80M USR without collateral, $25M loss. Secondary damage amplified as automated scripts in adjacent protocols continued operating post-depeg. Demonstrated the security risks of Agentic Finance

These figures show that security demand is expanding in parallel with DeFi growth. The rise of Agentic Finance in particular is creating monitoring demand not just for "transactions signed by humans" but also for "transactions autonomously executed by AI Agents."

Regulatory Trends

The regulatory environment has the potential to support demand for DeFi security services over the medium term.

SEC DeFi Frontend No-Action Position (April 13, 2026) On April 13, 2026, the U.S. SEC Division of Trading and Markets issued a statement regarding DeFi frontend and self-custody wallet UI providers. It took the form of a no-action letter exempting DeFi frontend providers meeting certain conditions from broker-dealer registration requirements.

The key conditions (summarizing 3 of 11 total conditions) are:

1. **User autonomy:** The frontend does not intervene in user trading decisions, and users operate autonomously

2. **Neutrality:** The frontend does not recommend or steer users toward specific transactions or counterparties
3. **Transparency:** Disclosure of conflicts of interest, implementation of cybersecurity measures, and publication of MEV (Maximal Extractable Value) strategies

This measure is time-limited for five years, valid through April 13, 2031.

This is a SEC staff no-action position (not a legally binding rule, but functioning as a de facto signal to the market), and the increase of DeFi frontends directly expands the market for security infrastructure like Ninja. The third condition — "Transparency" covering conflict-of-interest disclosure, cybersecurity measures, and MEV strategy publication — aligns directly with the value provided by Ninja Intelligence Core. As infrastructure for DeFi frontend providers to comply with SEC conditions, Ninja's B2B API has a concrete use case.

Other Regulatory Developments

- **MiCA (EU Markets in Crypto-Assets Regulation):** A comprehensive regulatory framework for crypto-asset markets. Fully enforced since December 2024 (transition period final deadline: July 2026). The scope of application to DeFi protocols is not clearly defined at present, and is at the stage of monitoring future detailed rulemaking. The directional trend toward greater transparency and security requirements is viewed as a tailwind
- **PSA Amendment (scheduled for enforcement June 2026, Japan):** An amendment to the Payment Services Act, primarily targeting crypto-asset exchange operators. The scope of application to DeFi protocols and related services is unclear, but it may lead to strengthened security requirements for DeFi access through exchanges. Additionally, in April 2026, the Japanese Cabinet approved a bill to amend the Financial Instruments and Exchange Act (FIEA) that would reclassify crypto-assets as "financial instruments." This represents a fundamental regulatory shift from PSA (payment instruments) to FIEA (financial instruments), with potential medium-term implications for the regulatory environment of DeFi-related services
- **EU AI Act (entered into force August 2024; key provisions — high-risk AI requirements, transparency obligations, etc. — scheduled for full application from August 2026):** A comprehensive AI regulatory framework. In the Agentic Finance context, explainability and auditability of AI Agent behavior may become subject to regulation in the future. Ninja's two-layer architecture (deterministic Layer 1 + LLM-augmented Layer 2) is consistent with this regulatory direction

Overall, these regulatory trends are positioned as medium-term demand drivers for security services in DeFi and Agentic

Finance. The SEC's DeFi frontend no-action position, in particular, is viewed as a concrete market opportunity.

References

1. DeFi Llama — Total Value Locked (TVL) Statistics. https://defillama.com/
2. Chainalysis — 2025 Crypto Crime Report. Hacking losses data ($3.4B).
3. Morpho Labs — Protocol metrics and growth data. https://morpho.org/
4. Hypernative — Company information and product documentation. https://hypernative.io/
5. Blockaid — Company information and product documentation. 6.3B+ transactions scanned. https://blockaid.io/
6. GoPlus Security — API documentation and company metrics ($4.7M revenue, Jan-Oct 2025). https://gopluslabs.io/
7. U.S. SEC Division of Trading and Markets — Staff Statement on DeFi Front-End Providers (April 13, 2026).
8. European Commission — Markets in Crypto-Assets Regulation (MiCA).
9. Financial Services Agency, Japan — Payment Services Act amendments.
10. European Commission — EU Artificial Intelligence Act.
11. Forta Network — Decentralized threat detection. https://forta.org/
12. Resolv Incident Report — March 2026. Compromised key exploit — unauthorized minting of 80M USR ($25M loss).
13. Web3 Security Market — $2.9B (2025) to $15.8B (2032 projected, CAGR 26.4%). Source: Electronics Media / Market Research analysis (March 2026).
14. Model Context Protocol (MCP) — Anthropic (2024). https://modelcontextprotocol.io/

Ninja — Agentic Security Controlplane for DeFi & Agentic Finance ShoGun — See Through DeFi, Control Agentic Finance