AI Consensus Audit Whitepaper

Evidence-Weighted AI Security for Smart Contracts

Version 1.0
July 2026
Website: AIConsensusAudit.com
Token: $AICA

Executive Summary

Smart contracts secure billions of dollars in digital assets, but the security review process remains slow, expensive, fragmented, and often inaccessible to early-stage builders. AI has changed what is possible. Modern frontier models can read smart-contract code, identify vulnerabilities, explain exploit paths, and recommend fixes. But a single AI model is not enough. One model has one reasoning style, one set of blind spots, and one tendency toward false positives or missed edge cases. **AI Consensus Audit** introduces an evidence-weighted multi-agent audit engine for smart-contract security. Instead of relying on one AI auditor, AIConsensusAudit.com runs code through a panel of independent AI auditors, reconciles their findings, weighs the evidence, challenges disagreements, and validates serious issues against the actual code. The result is a faster, cleaner, and more actionable security report designed to help builders identify risk before deployment, before manual audit, and after code changes. **$AICA** is the official payment token of AIConsensusAudit.com. It gives users a crypto-native way to access AI-powered smart-contract scans and unlock the platform's preferred scan rate. **Standard full scan: $100**
Full scan with AICA: 50% DISCOUNT. $AICA is designed around real platform utility: scan payments, API access, continuous re-scans, ecosystem incentives, security research grants, and continuous development of the AI Consensus Audit engine.

Important Notice

This whitepaper is provided for informational purposes only. $AICA is designed as a **utility token** for accessing services within the AI Consensus Audit ecosystem. It is not intended to represent equity,

ownership, revenue rights, profit share, passive income, or any claim on the company, platform, treasury, or future business performance. Nothing in this document should be interpreted as financial, investment, legal, tax, or regulatory advice.

---

1. **The Problem**

Smart-Contract Security Is Still Too Expensive, Too Slow, and Too Reactive

Smart contracts are unforgiving. Once deployed, vulnerable code can lead to:

- **Loss of funds**
- **Frozen assets**
- **Protocol insolvency**
- **Unauthorized access**
- **Oracle manipulation**
- **Broken accounting**
- **Governance failures**
- **Permanent trust damage**

Traditional manual audits remain essential, but they are often expensive and time-consuming. Many teams delay security review until late in the development cycle because the cost and coordination burden are high. Automated scanners help, but most are limited by narrow rule sets, chain-specific assumptions, or high false-positive rates. AI scanners improve accessibility, but many rely on a single model. That creates a new problem: **A single AI auditor is a single point of failure.** One model may catch real issues. But it can also miss subtle vulnerabilities, misunderstand protocol logic, or produce generic warnings that are not actually exploitable. For smart-contract security, that is not enough.

2. Why Single-Model AI Audits Are Limited

One Model Means One Set of Blind Spots

AI models can reason over code, but they do not all reason the same way. One model may be stronger at access-control analysis.
Another may be better at economic attack paths.
Another may catch re-entrancy patterns.
Another may better explain accounting invariants. Relying on only one model creates several risks:

- **False negatives:** real vulnerabilities are missed.
- **False positives:** harmless code is flagged as dangerous.
- **Weak prioritization:** generic warnings crowd out serious issues.
- **Context failure:** the model misunderstands the protocol's design.
- **No second opinion:** one model's conclusion becomes the final answer.

A basic multi-model system improves this by running several models and surfacing what they agree on. But simple agreement is not enough. Multiple models can agree on the same false positive. Multiple models can miss the same subtle exploit. And sometimes the most important vulnerability is found by only one strong reviewer. That is why AI Consensus Audit does not use simple majority voting.

---

3. The AI Consensus Audit Solution

Consensus Is Not Majority Vote. Consensus Is Evidence That Survives Challenge.

AI Consensus Audit is an optimized AI audit system built around **evidence-weighted consensus**. The platform does not simply ask, "How many models found this?" It asks:

- **Is the finding tied to specific code?**
- **Is the vulnerable path real?**
- **Can an attacker trigger it?**
- **What preconditions are required?**
- **Does the issue affect funds, permissions, pricing, accounting, or protocol state?**
- **Did other auditors identify protections that block it?**

- **Can the issue survive adversarial review?**
- **Is the exploit path realistic or only theoretical?**

This approach is designed to produce a ranked report based on the strength of the evidence, not model popularity. **Agreement helps. Evidence decides.**

---

4. How the Audit Engine Works

From Code to Risk Report in Minutes

AI Consensus Audit uses a multi-stage review pipeline.

4.1 Code Intake

Users can paste a contract, upload code, or submit a repository. The platform accepts smart-contract code across multiple ecosystems, including:

- **Solidity**
- **Vyper**
- **Rust**
- **Move**
- **Cairo**
- **CosmWasm**
- **Other smart-contract languages**

The system begins by ingesting the codebase and preparing it for analysis.

4.2 Static Context Mapping

Before AI auditors review the code, the platform builds structured context. This may include:

- **Contracts**
- **Functions**
- **Modifiers**
- **State variables**
- **External calls**
- **Payable functions**
- **Admin-only functions**
- **Token interactions**
- **Oracle references**
- **Signature verification logic**
- **Upgradeability patterns**
- **Critical execution paths**

This context helps AI auditors focus on the most security-relevant areas of the code.

4.3 Independent Specialist AI Auditors

The code is then reviewed by multiple independent AI auditors. Each auditor works separately and does not see the others' findings during the initial review. This reduces groupthink and helps preserve independent reasoning. Auditor perspectives may include:

- **Access-control auditor**
- **Re-entrancy and external-call auditor**
- **Accounting and invariant auditor**
- **Oracle and economic-risk auditor**
- **Signature and authentication auditor**
- **Upgradeability and storage-layout auditor**

- **State-machine and edge-case auditor**
- **General smart-contract security auditor**

Each auditor returns structured findings with:

- **Title**
- **Severity**
- **Confidence**
- **Affected contracts**
- **Affected functions**
- **Code references**
- **Root cause**
- **Attack path**
- **Preconditions**
- **Impact**
- **Recommended fix**
- **False-positive considerations**

4.4 Finding Normalization and Deduplication

The system normalizes the auditors' outputs and groups related findings. This prevents the final report from showing multiple versions of the same issue. Findings are clustered based on:

- **Affected function**
- **Vulnerability class**
- **Root cause**
- **Code path**
- **Attack path**
- **Impacted asset or state**
- **Semantic similarity**
- **Overlapping code references**

The goal is simple: **One real issue should appear as one clear finding.**

4.5 Evidence-Weighted Scoring

Each finding cluster is scored using multiple signals. The system considers:

- **Agreement strength**
- **Evidence quality**
- **Exploitability**
- **Impact**
- **Contradicting evidence**
- **Code specificity**
- **Reviewer confidence**
- **Existing protections**
- **Affected assets and permissions**

A finding is not automatically trusted because several models mention it. A finding is also not automatically dismissed because only one auditor found it. A single well-supported critical issue can outrank a vague warning repeated by multiple models.

4.6 Disagreement Review

When auditors disagree, the system treats the conflict as useful signal. For example:

- One auditor flags a missing access control.
- Another auditor identifies a modifier that appears to block the issue.
- A third auditor suggests the modifier can be bypassed under a specific state condition.

Rather than averaging these opinions, the platform performs an adversarial review. It checks:

- **Modifiers**
- **Require statements**
- **Role checks**
- **Function call paths**
- **State transitions**
- **User-controlled inputs**
- **Potential bypasses**
- **Whether protections are complete or partial**

The finding may be confirmed, downgraded, marked uncertain, partially validated, or refuted.

4.7 Exploitability Validation

High-impact findings are tested against the actual code path. The validator asks:

- **Who is the attacker?**
- **What permissions are required?**
- **What contract state must exist?**
- **What sequence of calls is needed?**
- **Are the inputs attacker-controlled?**
- **Do existing checks block the exploit?**
- **Can the issue affect funds, ownership, accounting, pricing, or protocol state?**

If the issue appears exploitable, it is prioritized.
If the code prevents it, the finding is downgraded or removed. This reduces generic scanner noise and improves the usefulness of the final report.

4.8 Final Report Generation

The final report includes:

- **Consensus Risk Score**
- **Ranked findings**
- **Severity levels**
- **Confidence scores**
- **Exploitability notes**
- **Code-specific explanations**
- **Attack paths**
- **Disagreement summaries**
- **Suggested fixes**
- **Architecture map**
- **Generated invariants**
- **Recommended next steps**

The report is designed to help teams answer: **What should we fix first, why does it matter, and how confident is the system that the issue is real?**

---

5. Consensus Risk Score

A 0–100 Signal Based on Evidence, Not Votes

The Consensus Risk Score gives users a simple read on the contract's security posture after AI review. The score may incorporate:

- **Severity of findings**
- **Confidence level**
- **Exploitability**
- **Evidence strength**

- **Model convergence**
- **Disagreement results**
- **Number of confirmed issues**
- **Potential impact**
- **Presence of critical code paths**

The score is not a guarantee of safety. It is a structured risk signal designed to help teams prioritize review and remediation. A higher-risk score indicates that the audit engine found more serious, more credible, or more exploitable issues. A lower-risk score indicates fewer detected issues or findings with lower severity, lower confidence, or weaker exploitability evidence.

---

6. Platform Deliverables

What Users Receive from a Scan

Every full scan is designed to provide a practical security review package.

Consensus Risk Score

A single 0–100 risk read based on evidence strength, exploitability, severity, disagreement, and model convergence.

Ranked Findings

Issues sorted by priority, so teams know what to fix first.

Plain-English Risk Explanations

Each finding explains what is wrong, why it matters, and how it could affect the protocol.

Exploitability Notes

High-impact issues include analysis of whether the vulnerability appears realistically exploitable.

Disagreement Summary

If auditors disagreed, the report explains how the conflict was resolved.

Suggested Fixes

Findings include remediation guidance to help developers patch the issue.

Architecture Map

The platform generates a map of contracts, functions, relationships, external calls, and critical paths.

Generated Invariants

AI Consensus Audit surfaces properties the contract should preserve, which can support testing, fuzzing, and manual review.

Continuous Re-Scan

Teams can re-scan code as it changes to catch new issues introduced by fixes, upgrades, or commits.

---

7. Target Users

Built for Builders, Protocols, and Auditors

AI Consensus Audit is designed for:

- **Smart-contract developers**
- **Protocol teams**
- **DeFi builders**
- **NFT and gaming projects**
- **DAO developers**
- **Security researchers**
- **Independent auditors**
- **Launchpads**
- **Accelerators**
- **Web3 agencies**
- **Teams preparing for manual audit**
- **Teams reviewing fixes before deployment**

The platform is especially useful before a manual audit, after major code changes, before upgrades, and during active development. AI Consensus Audit does not replace expert human auditors. It helps teams start with a cleaner, more structured view of likely risk.

8. $AICA Token Overview

The Official Payment Token of AIConsensusAudit.com

$AICA sits at the convergence of **crypto, AI, and smart-contract security**. It is the native payment token for AIConsensusAudit.com and is designed to give users a crypto-native way to access AI-powered security scans. The core utility is simple: **Standard full scan: $100**
Full scan with AICA: 50% OFF. Users who pay with $AICA unlock the preferred scan rate while receiving the same evidence-weighted AI audit report. The scan does not change. **Only the payment rate changes.**

---

9. $AICA Utility

Product-First Token Utility

$AICA is designed around direct use inside the AI Consensus Audit ecosystem. Primary and future platform utilities may include:

Scan Payments

Users can pay with $AICA to run smart-contract scans on AIConsensusAudit.com.

Preferred Scan Rate

Paying with AICA unlocks the native-token scan rate: $50 instead of $100.00 for a full scan.

API Access

$AICA may be used for API-based scans, developer integrations, and automated security workflows.

Continuous Re-Scans

Teams may use $AICA for repeated scans as code changes over time.

Advanced Reports

$AICA may be used for premium report features, PDF exports, deeper validation passes, or repository-level scans.

Security Research Grants

$AICA may be used to reward active contributors who improve the ecosystem through verified security research, benchmark contributions, tooling, or audit-engine improvements.

Ecosystem Incentives

$AICA may support builder onboarding, scan credits, integrations, partnerships, and community programs.

Future Auditor Network Participation

If the platform expands into a broader auditor or module marketplace, $AICA may be used for access, quality-control mechanisms, or participation in ecosystem workflows.

---

10. Tokenomics

Fixed Supply Utility Token

Total Supply

100,000,000 $AICA

$AICA has a fixed total supply of **100 million tokens**. The tokenomics are designed to support product usage, ecosystem growth, security research, liquidity, and continuous development of the AI Consensus Audit platform.

11. Token Allocation

Category	Allocation	Tokens	Purpose
Ecosystem Growth & Scan Incentives	**30%**	**30,000,000 $AICA**	User onboarding, scan credits, builder campaigns, integrations, community growth
Treasury, Operations & Continuous Development	**20%**	**20,000,000 $AICA**	AI model costs, infrastructure, security improvements, product development, platform maintenance
Team & Core Contributors	**15%**	**15,000,000 $AICA**	Founder, team, and contributor allocation tied to continued platform work
Security Research & Auditor Grants	**10%**	**10,000,000 $AICA**	Verified research, benchmarks, tooling, auditor contributions, bug research
Liquidity & Payment Rails	**10%**	**10,000,000 $AICA**	DEX liquidity, checkout support, payment routing, crypto-native payment infrastructure
Public / Community Launch	**10%**	**10,000,000 $AICA**	Initial public and community distribution

Category	Allocation	Tokens	Purpose
Advisors & Strategic Partners	**5%**	**5,000,000 $AICA**	Legal, security, AI, crypto, ecosystem, and partnership support
Total: 100% / 100,000,000 $AICA			

12. Allocation Details

Ecosystem Growth & Scan Incentives — 30%

30,000,000 $AICA. This allocation supports ecosystem growth and user adoption. Potential uses include:

- **Scan credits**
- **Builder onboarding**
- **Developer campaigns**
- **Protocol partnerships**
- **Integration incentives**
- **Community programs**
- **Hackathon support**
- **Launch partner programs**

The goal is to help more teams access AI-powered smart-contract security review.

Treasury, Operations & Continuous Development — 20%

20,000,000 $AICA. This allocation supports the long-term operation and improvement of AIConsensusAudit.com. Potential uses include:

- **AI model and inference costs**
- **Cloud infrastructure**
- **Platform maintenance**
- **New auditor agents**
- **Benchmarking**

- **API development**
- **Security upgrades**
- **Repository scanning improvements**
- **Continuous re-scan features**
- **Protocol integrations**
- **Ongoing product development**

This category is designed to ensure the platform can keep improving as AI models, smart-contract ecosystems, and security threats evolve.

Team & Core Contributors — 15%

15,000,000 $AICA. This allocation is reserved for the team and core contributors responsible for building, maintaining, and improving the platform. Team and core contributor tokens are subject to vesting.

Team Vesting Schedule

- **3-month cliff**
- **12 equal monthly unlocks after the cliff**
- **Fully vested after 15 months total**

This structure aligns team allocation with continued platform work, maintenance, product delivery, and ecosystem development.

Security Research & Auditor Grants — 10%

10,000,000 $AICA. This allocation supports active contributions to the AI Consensus Audit security ecosystem. Potential uses include:

- **Verified vulnerability research**
- **Benchmark dataset contributions**
- **Audit methodology improvements**
- **Prompt and model evaluation**
- **Security tooling**
- **Auditor modules**
- **Bug bounty support**
- **Research grants**

Rewards from this allocation are intended for active contributions, not passive holding.

Liquidity & Payment Rails — 10%

10,000,000 $AICA. This allocation supports crypto-native access to the platform. Potential uses include:

- **DEX liquidity**
- **Checkout support**
- **Payment routing**
- **Liquidity infrastructure**
- **Cross-chain payment support**
- **Wallet-based scan payments**

The goal is to make it easy for users to acquire and use $AICA for platform services.

Public / Community Launch — 10%

10,000,000 $AICA. This allocation supports initial community access and public distribution. The purpose is to make $AICA available to users, builders, early adopters, and community members who want to use the token inside the AI Consensus Audit ecosystem.

Advisors & Strategic Partners — 5%

5,000,000 $AICA. This allocation is reserved for advisors and partners who support the platform through:

- **Security expertise**
- **AI expertise**
- **Legal and compliance guidance**
- **Crypto infrastructure**
- **Ecosystem partnerships**
- **Business development**
- **Technical integrations**

13. Team Vesting Table

For the **15,000,000 $AICA** Team & Core Contributors allocation:

Period	Unlock Amount
Months 0–3	**0 $AICA**
Month 4	**1,250,000 $AICA**
Month 5	**1,250,000 $AICA**
Month 6	**1,250,000 $AICA**
Month 7	**1,250,000 $AICA**
Month 8	**1,250,000 $AICA**
Month 9	**1,250,000 $AICA**
Month 10	**1,250,000 $AICA**
Month 11	**1,250,000 $AICA**
Month 12	**1,250,000 $AICA**
Month 13	**1,250,000 $AICA**
Month 14	**1,250,000 $AICA**
Month 15	**1,250,000 $AICA**
Total unlocked after Month 15: 15,000,000 $AICA	

14. Platform Payment Flow

How $AICA Is Used in the Product

The platform payment flow is designed to be simple.

1. User submits a contract or repository.
2. AI Consensus Audit estimates scan scope.
3. User chooses payment method.
4. Standard payment price is **$100**.
5. AICA payment price is $50.00.
6. User connects wallet and pays.
7. The audit engine runs the scan.
8. User receives the final ranked consensus report.

$AICA creates a direct utility loop: **Use token → access scan → receive report → improve code security.**

---

15. Use of Platform-Received $AICA

$AICA received by the platform may be used to support ecosystem operations. Potential uses include:

- **AI model costs**
- **Infrastructure**
- **Continuous development**
- **Security research grants**
- **Ecosystem incentives**
- **Payment infrastructure**
- **Liquidity and checkout support**
- **Platform maintenance**

This keeps the token tied to platform usage and operational utility.

16. Product Roadmap

Planned Development Phases

The roadmap below reflects intended platform development areas and may evolve based on technical progress, user demand, and ecosystem needs.

Phase 1: Core Scan Engine

Focus:

- **Single-file scans**
- **Independent AI auditor panel**
- **Evidence-weighted scoring**
- **Ranked findings**
- **Plain-English reports**
- **$AICA payment support**
- **Free scan access**
- **Standard and $AICA pricing**

Phase 2: Repository-Level Audits

Focus:

- **Full repository scans**
- **Multi-file context**
- **Dependency mapping**
- **Architecture maps**
- **Contract interaction analysis**
- **PDF-style reports**
- **Improved deduplication**
- **Expanded language support**

Phase 3: Continuous Re-Scan

Focus:

- **GitHub/repository integration**
- **Re-scan after commits**
- **Change-aware analysis**
- **Fix verification**
- **Regression detection**
- **CI/CD support**

Phase 4: API Access

Focus:

- **Developer API**
- **Programmatic scan submission**
- **Team dashboards**
- **Usage-based scan credits**
- **Webhook notifications**
- **Integration with developer workflows**

Phase 5: Advanced Validation

Focus:

- **Deeper exploitability validation**
- **Generated test outlines**
- **Foundry-style test suggestions**
- **Invariant generation**
- **Benchmark-driven model evaluation**
- **More specialized auditor agents**

Phase 6: Ecosystem Expansion

Focus:

- **Security research grants**
- **Auditor tooling**
- **Benchmark contributions**
- **Partner integrations**
- **Launchpad integrations**
- **Protocol security workflows**
- **Future auditor network exploration**

---

17. Multi-Chain and Multi-Language Vision

One AI Security Layer for a Multi-Chain World

Smart-contract development is no longer limited to one chain or one language. Protocols increasingly span:

- **EVM chains**
- **Solana**
- **Cosmos**
- **Move-based ecosystems**
- **Starknet**
- **Layer 2 networks**
- **Appchains**
- **Emerging execution environments**

AI Consensus Audit is designed as a code-reasoning audit layer, not a scanner locked to one virtual machine. The long-term goal is to provide a unified security workflow across smart-contract ecosystems. **One panel. One report. One risk view.**

18. Role of Human Auditors

AI-Assisted Security, Not Human Replacement

AI Consensus Audit is built to complement human auditors, not replace them. Human auditors remain essential for:

- **Protocol design review**
- **Economic modeling**
- **Complex system assumptions**
- **Governance analysis**
- **Manual exploit development**
- **Deep adversarial reasoning**
- **Final deployment assurance**

AI Consensus Audit helps teams prepare for that process by surfacing likely risks earlier. The platform is designed to:

- **Catch obvious and moderate issues earlier**
- **Highlight high-risk code paths**
- **Reduce scanner noise**
- **Generate useful context**
- **Help teams fix issues before manual review**
- **Give human auditors a clearer starting point**

Security improves when AI and human expertise work together.

---

19. Security Philosophy

More Alerts Are Not the Goal

Many automated tools produce long lists of warnings. That is not the goal of AI Consensus Audit. The goal is to produce fewer, better, more actionable findings. The platform is built around several principles:

- **Evidence over popularity**
- **Exploitability over theory**
- **Specificity over generic warnings**

- **Disagreement as signal**
- **High-impact minority findings preserved**
- **Severity separated from confidence**
- **Human-readable reporting**
- **Continuous improvement**
- **Product utility first**

Every serious finding should earn its place in the final report.

---

20. Compliance-Oriented Token Design

Utility-First by Design

$AICA is designed for platform usage. The token is not designed to provide:

- **Equity**
- **Ownership rights**
- **Revenue share**
- **Profit participation**
- **Passive yield**
- **Dividends**
- **Claims on assets**
- **Governance over company operations**

$AICA is intended to be used for:

- **Paying for scans**
- **Accessing the preferred scan rate**
- **API usage**
- **Continuous re-scans**
- **Advanced platform features**
- **Security research grants**
- **Ecosystem participation**

The project's messaging should remain focused on utility, product access, and smart-contract security.

---

21. Why $AICA Matters

The Convergence of Crypto and AI Security

Crypto teams build on-chain.
AI reviews code at machine speed.
$AICA connects the two through a native payment layer for smart-contract security.
AIConsensusAudit.com gives builders access to independent AI auditors, evidence-weighted consensus, exploitability review, and ranked vulnerability reports. $AICA gives those builders a crypto-native way to use the platform. **Connect wallet. Pay with $AICA. Run scan. Get report.** That is the core utility. No subscriptions.
No credit cards.
No sales calls. Just fast access to AI-powered smart-contract security review.

---

22. Summary

AI Consensus Audit is building a practical AI security layer for smart-contract teams. The platform combines:

- **Independent AI auditors**
- **Specialized review perspectives**
- **Evidence-weighted consensus**
- **Disagreement review**
- **Exploitability validation**
- **Ranked security reports**
- **Multi-chain and multi-language support**
- **Continuous development**
- **Crypto-native payments through $AICA**

$AICA is the official payment token of AIConsensusAudit.com. It allows users to access the preferred scan rate :**Standard full scan: $100**
Full scan with AICA: 50% OFF. The project's mission is simple: **Make smart-contract security faster, clearer, more accessible, and more crypto-native.** AI Consensus Audit helps builders find risk before attackers do. $AICA powers access to that workflow.