THE AI INTERNET

Infrastructure for the AI Agent Economy

VAC Protocol + Athena

Verified Human Identity, Trust-Weighted Routing,
and Intelligent Orchestration for AI Agents

Whitepaper v5.0

March 2026
Roberto Zagarella
Violet Shores Pty Ltd

458 Patent Claims | 10 IP Filings | Patent Pending

NIST CAISI RFI Respondent (Docket NIST-2025-0035)

vacprotocol.org | athenapilot.ai | aiinternet.ai

Executive Summary

AI agents are becoming the primary interface between humans and digital systems. They execute tasks, delegate to other agents, manage budgets, control physical systems, and operate across organisational and jurisdictional boundaries. Yet the fundamental question remains unanswered: who authorised this agent, and can you prove it?

This whitepaper introduces a two-layer infrastructure stack that answers this question and builds an intelligent routing layer on top of it:

- **VAC Protocol** (Layer 1: Identity and Trust) provides cryptographic human attribution for AI agents through multi-modal biometric verification, delegatable authority tokens, and a social trust graph. Every agent action traces to a verified human through an unbroken chain of cryptographic proof.
- **Athena** (Layer 2: Intelligence and Routing) provides trust-weighted routing, orchestration, and execution for AI agents using SignalRank, a ranking algorithm that scores agents, models, and reasoning paths across multiple trust dimensions. Athena selects the most reliable path for every query.

Together, VAC and Athena form the infrastructure layer for the AI agent economy. VAC establishes who you can trust. Athena determines the best way to route intelligence through that trust network.

Industry data underscores the urgency: only 14.4% of organisations report their AI agents go live with full security approval (Gravitee State of AI Agent Security 2026). The vast majority of agents launch without complete oversight. This gap is precisely what the VAC Protocol addresses.

Metric	Value
Patent claims filed	458 across 10 filings
Patent status	All Patent Pending (IP Australia)
VAC Protocol claims	285 claims, 6 filings
Athena claims	173 claims, 4 filings
Live implementation	vacprotocol.org + api.athenapilot.ai
Standards engagement	NIST CAISI RFI respondent + NCCoE concept paper (March-April 2026)
Domains	aiinternet.ai, vacprotocol.org, athenapilot.ai

1. The Problem: Unverified Intelligence

1.1 The Human Attribution Gap

Current identity frameworks authenticate agents as machine identities but do not maintain a verifiable link to the human who initiated the chain of action. OAuth authorises access. SPIFFE identifies workloads. Neither proves which human authorised the agent, with what scope, through what chain of command.

When an AI agent executes a financial transaction, drafts a legal document, controls a drone, or makes a medical recommendation, no existing protocol can cryptographically prove: who authorised it, what scope they granted, through how many delegation layers, and whether that authority is still valid.

1.2 The Trust Routing Gap

Even when agents are properly authorised, a second problem emerges: which agent or model should handle a given task? Current systems treat all LLMs as interchangeable or rely on static configuration. But models differ dramatically in reliability, restriction policies, hallucination rates, and domain expertise.

A model that excels at code generation may fabricate legal citations. A model with strong factual accuracy may refuse legitimate queries due to overly conservative safety policies. No existing system maps these differences dynamically or routes tasks accordingly.

1.3 The Reasoning Provenance Gap

When an AI system produces an answer, the reasoning path that produced it is typically discarded. There is no mechanism to verify which sources were consulted, which agents contributed, whether the reasoning has been validated by prior use, or how confident the system should be in its output.

These three gaps, taken together, define the infrastructure deficit of the AI agent economy: no verified identity, no trust-weighted routing, and no reasoning provenance.

---

2. Architecture: The AI Internet Stack

The AI Internet is a five-layer architecture analogous to the network stack that powers the traditional internet. Each layer serves a distinct function and maps to specific components in the VAC + Athena system.

Layer	Function	Internet Analogy	Component
1. Identity	Verified humans and organisations	IP Addressing	VAC Protocol

2. Trust	Agent trust profiles and relationships	TLS / Certificate Authorities	VAC Trust Graph
3. Knowledge	Claims, sources, citations	DNS / Domain Registry	Athena Knowledge Graph
4. Reasoning	Indexed reasoning paths and workflows	HTTP / Application Protocol	Athena Reasoning Path Index
5. Orchestration	Routing, selection, execution	BGP / Network Routing	Athena SignalRank Engine

Layer 1 (VAC) provides the root of trust. Layers 2 through 5 (Athena) build intelligence on top of that trust foundation. Each layer has clearly defined interfaces and can operate independently, but the full stack delivers capabilities no single layer can provide.

---

3. VAC Protocol: Cryptographic Human Attribution

3.1 Verified Authority Token (VAT)

The VAT is the central cryptographic object of the VAC Protocol. It is a JWT-compatible signed token using Ed25519 signatures that binds a biometrically-verified human identity to every agent action.

Component	Contents
Identity	Cryptographic hash of verified human (not plaintext PII)
Trust	Trust score, minimum threshold, verification recency
Scope	Resource types, action types, data domains, temporal windows, physical parameters
Delegation	Max depth, current depth, parent reference, chain metadata
Context	Organisational, multi-party, jurisdictional provenance
Signature	Ed25519 binding payload to biometric attestation

3.2 Multi-Modal Biometric Verification

VAC uses up to seven biometric modalities for human verification: facial geometry with liveness detection, voice pattern with speaker verification, challenge-response (spoken dynamic digits), lip-sync correlation between audio and visual, behavioural biometrics (interaction patterns), duress detection (coercion indicators), and optional location verification. Multi-modal composite scoring provides higher assurance than any single modality and maps to NIST SP 800-63 Identity Assurance Levels.

3.3 Delegation Chains and Scope Narrowing

Authority flows downhill and can only get narrower. When an agent delegates to a sub-agent, the derived scope is the set intersection of the parent scope and the requested scope. Trust decays with each delegation level. A financial coordinator agent cannot grant its reconciliation sub-agent write access if the coordinator only has read access. This is enforced cryptographically, not by policy.

3.4 Trust Graph

The VAC trust graph represents verified relationships between humans, organisations, and agents. Trust edges are created through biometric vouching: one verified human vouches for another, creating a bidirectional trust relationship. The density of trust connections determines verification cost: interactions within a dense trust cluster require lightweight verification (low cost), while interactions at trust boundaries with unknown entities require full certified verification (revenue event).

3.5 Collective Governance

Multi-party biometric authorisation supports M-of-N threshold verification, consensus governance, weighted authority with mandatory role verification, and cultural protocol scope constraints. The protocol supports indigenous data sovereignty frameworks including Te Mana Raraunga, CARE Principles, and OCAP through collective trust roots with configurable decision-making models.

3.6 Continuous Collective Validity Monitoring

Collective authority is not a one-time issuance. It is a continuously validated state. The VAC Protocol treats every collective governance arrangement as a living system that must be monitored for ongoing validity.

When any participant's authority changes, whether through biometric re-verification failure, role change, or trust score decline, the system applies governance-model-aware cascade logic. If the governance model requires a specific role (such as a kaumātua in Māori governance) and that role's authority lapses, the entire collective authority is suspended regardless of whether a general quorum still exists.

Severity	Trigger	Response
Minor	Trust score fluctuation within threshold	Logging only
Moderate	Role change, single participant re-verification failure	Quorum re-evaluation
Critical	Root authority revocation, security incident	Immediate collective suspension

Operations interrupted mid-execution by authority changes generate forensic gap provenance records documenting what was in progress, at what point authority became invalid, and what

remediation is required. After a cascade event, the system supports structured authority reconstitution without requiring the entire governance process to restart from scratch, provided the reconstitution meets the original governance model's requirements.

3.7 Physical System Authority

The protocol extends to AI systems controlling physical assets: drones, robots, autonomous vehicles, manufacturing systems, and military platforms. Physical operating parameters (geofence, altitude, speed, force limits) are encoded as cryptographic scope constraints in the VAT. Single biometric revocation can ground an entire swarm. Every physical action traces to a verified human through the complete delegation chain.

3.8 Regulatory Compliance Provenance

Each delegation level appends a provenance record containing jurisdictional context, infrastructure footprint, and verification status. This creates a CDR-style audit trail enabling cross-border compliance with GDPR, LGPD, PIPL, and more than ten other regulatory frameworks simultaneously. A global regulatory intelligence layer automates detection of conflicting requirements before agents act.

3.9 Intelligent Orchestration and Collaborative Onboarding

The VAC Protocol does not only enforce security constraints. It also intelligently configures them. When a user describes their problem in natural language, the orchestration layer translates this into a recommended agent team with automatically configured authority scopes and hierarchical delegation structures. Each agent receives the minimum scope required for its task, enforced cryptographically via VAT rather than by policy.

Dual-Purpose Biometric Capture

A single video recording simultaneously serves as functional product interaction (such as a practice session or identity verification) and biometric anchoring. Security is embedded within the user experience rather than added as a separate step. The user never experiences security as friction because the verification happens within the action they were already taking.

Vouch-as-Collaboration-Invitation

The identity vouching step simultaneously establishes identity verification and bilateral governance. When a person vouches for someone's identity, they are also prompted to indicate whether they will be working with that person. If yes, the vouch step configures bilateral governance in a single action. It does not feel like security; it feels like inviting a teammate. This collapses what would otherwise be two separate workflows (identity verification and access governance) into one natural interaction.

Collaboration Discovery

When independently verified users would benefit from collaborative access, the orchestration layer detects the opportunity and automatically recommends a governance model, configures the collective authority structure, and establishes scope boundaries for cross-user agent

interactions. All such configuration requires biometric re-verification from all participants before activation.

---

4. Athena: Trust-Weighted Routing and Orchestration

4.1 The SignalRank Engine

SignalRank is Athena's trust-weighted ranking algorithm. It ranks AI models, agents, knowledge claims, and reasoning paths across multiple trust dimensions to determine the most reliable path for any given query.

Trust Dimensions (per model, per domain)

Dimension	What It Measures	How It Learns
Reliability	Consistent answers across repeated queries	Cross-model consensus detection
Honesty	Transparent refusal vs softening or deflecting	Response comparison analysis
Restriction Transparency	How openly a model communicates its limitations	Rail detection engine
Hallucination Rate	Factual accuracy vs fabrication	Cross-model factual divergence
Independence	Original analysis vs regurgitation	Output uniqueness scoring

Trust profiles evolve through exponential moving average learning. Every interaction shifts the profile. Over time, SignalRank builds a personalised trust topology that reflects actual model behaviour in the user's specific domains, not static benchmarks.

4.2 Rail Detection

When an AI model refuses, softens, or redirects a legitimate query due to provider-imposed restrictions, Athena's rail detection engine identifies and classifies the restriction. By running identical prompts through multiple models and comparing outputs at a semantic level, the system builds per-provider restriction maps showing which models restrict which categories of tasks.

This enables sovereign routing: tasks are directed to models that will execute them without interference, while restricted models are used for tasks where their restrictions do not conflict. All routing decisions are cryptographically attributed through VAC.

4.3 Hallucination Detection

Cross-model comparison enables real-time hallucination detection. When three models agree on a factual claim and one diverges, the divergent output is flagged as a hallucination candidate. Ground truth anchoring from user-provided references and per-claim confidence scoring provide additional verification layers. Per-model hallucination profiles feed directly into SignalRank trust scores.

4.4 Multi-LLM Orchestration

Athena decomposes complex tasks into subtasks and assigns each subtask to the optimal model based on SignalRank scoring. A research task might fan out across four models simultaneously, with results synthesised using a planner agent. A coding task might use one model for architecture, another for implementation, and a third for review. The orchestration layer handles parallel execution, dependency resolution, and cost-aware routing.

4.5 Reasoning Path Index

Instead of discarding the reasoning process after generating an answer, Athena indexes every reasoning path: which agents were used, which sources were consulted, what intermediate steps were taken, and what the outcome was. Validated reasoning paths are reused for similar future queries, creating a compounding intelligence advantage. Over time, the system accumulates a searchable library of proven reasoning workflows.

4.6 Self-Improving Methodology

Athena continuously monitors its own performance and captures methodology improvements as they emerge from real usage. Corrections, architectural decisions, process discoveries, and strategic insights are automatically detected, classified, and routed to an improvement pipeline. Promoted learnings become permanent rules that govern all future operations. The methodology that improves products also improves itself.

---

5. SignalRank: The Ranking Algorithm

SignalRank computes a composite trust score for every node in the intelligence graph (models, agents, knowledge claims, reasoning paths) using a weighted combination of trust dimensions:

Component	Symbol	Description
Human Trust Score	H	Trust in the verified human behind the agent or claim
Agent Reliability Score	A	Historical performance across tasks
Source Credibility Score	S	Quality and verifiability of sources used
Reasoning Reuse Frequency	R	How often a reasoning path has been validated

Outcome Accuracy	O	Verified accuracy of outputs over time
Graph Centrality	G	Position in the trust topology

SignalRank = (H x w_h) + (A x w_a) + (S x w_s) + (R x w_r) + (O x w_o) + G

Weights are configurable per user, per organisation, and per task type. Defence deployments weight human trust and source credibility highest. Research deployments weight reasoning reuse and outcome accuracy. Consumer deployments weight agent reliability and cost efficiency.

The algorithm is domain-aware: a model's trust score for legal research is independent of its trust score for code generation. This prevents a model that excels in one domain from being incorrectly trusted in another.

---

6. Conformance Testing Framework

The VAC Protocol defines a standardised conformance testing framework for human attribution systems. No equivalent standard exists today. As AI agent deployments scale, the ability to verify that attribution claims are genuine, complete, and enforced becomes a critical infrastructure requirement.

6.1 Core Conformance Metrics

Five metrics form the foundation of VAC conformance testing. Each metric is designed to be measurable, reproducible, and applicable to any implementation of the protocol.

Metric	Full Name	What It Tests	Target
VATV	VAT Verification Time	End-to-end chain verification at depth N	Sub-linear scaling with depth
RPL	Revocation Propagation Latency	Time from root revocation to last-agent suspension	<1s root, <5s full chain
SNER	Scope Narrowing Enforcement Rate	Derived tokens never exceed parent scope	100% (zero tolerance)
TSPA	Trust Score Propagation Accuracy	Trust computation consistency through chains	<0.1% error margin
AMLAS	Attribution Maturity Level Assessment	Standardised maturity scoring across implementations	Comparable cross-vendor output

6.2 Core Test Categories

Six categories of conformance testing cover the complete protocol surface:

- **Token Structural Conformance:** VAT contains all required fields, Ed25519 signatures are valid, JWT structure is compliant, and all cryptographic bindings are verifiable.
- **Derivation Rule Conformance:** Scope narrowing is strictly enforced (set intersection only), trust decay follows the configured formula, delegation depth limits are respected, and no derived token exceeds its parent.
- **Revocation Cascade Conformance:** Root revocation propagates to every descendant token within target latency, partial chain revocation correctly invalidates downstream tokens, and no orphaned tokens survive revocation.
- **Multi-Party Authorisation Conformance:** M-of-N thresholds are enforced exactly, consensus governance requires unanimous agreement, weighted authority respects mandatory role requirements, and biometric veto blocks issuance.
- **Organisational Hierarchy Conformance:** Authority flows correctly through organisational structures, department boundaries are respected, cross-organisational delegation requires bilateral verification.
- **Cross-Platform Interoperability Conformance:** VATs created on one implementation verify correctly on another, provenance chains are portable across systems, and trust scores compute identically regardless of implementation.

6.3 Extended Testing: Collective Governance

- **Consensus (M=N):** All participants must agree. Deferral (consensus not reached) is recorded as a governance outcome, not a system failure. No bypass mechanism exists.
- **Weighted Authority:** General quorum threshold is met AND mandatory role verification passes. If a required role is absent, collective authority is invalid regardless of quorum count.
- **Cultural Protocol Scope:** Scope constraints from cultural governance frameworks propagate correctly through the entire agent chain. No agent can widen cultural scope boundaries.
- **Optional Location:** Four-tier location enforcement is correctly applied per decision type (not required, recorded, preferred, required). Location-required decisions verify co-presence cryptographically.

6.4 Extended Testing: Physical Systems

- **Kinetic Scope Constraints:** Geofence, altitude, speed, force, and payload limits are cryptographically enforced and cannot be overridden by the agent.
- **Swarm Coordination:** Single biometric revocation cascades correctly to all units in a coordinated swarm within target latency.
- **Physical Action Provenance:** Complete recording of what was done, where, when, with what physical parameters, and under whose verified authority.

- **Emergency Override:** Halt and return-to-base commands execute within specified latency with full audit trail.

6.5 Extended Testing: Coalition Operations

- **Multi-National Authority:** Governance model enforcement (consensus, qualified majority, lead-nation, framework-nation) functions correctly across national boundaries.
- **National Caveats:** Caveats propagate through the delegation chain and cannot be overridden by coalition-level authority.
- **Cross-National Interoperability:** Mutual authority verification works across allied nations with appropriate classification handling.
- **Alliance Tier Trust:** Trust graph correctly enforces varying trust levels by alliance tier.

6.6 Graph-Based Trust Verification

The trust graph representing the complete authority structure is itself a testable object. Graph-based testing traverses every path and confirms: trust properties are maintained at every edge, scope constraints are correctly narrowed at every delegation, authority rules are enforced at every node, and revocation cascades along correct paths.

Test cases are generated automatically from the authority structure, including boundary conditions and adversarial paths. The output is a conformance map showing pass/fail status at each node and edge in the graph.

---

7. Attribution Maturity Model

The Attribution Maturity Model provides a standardised framework for assessing the maturity of human attribution in AI agent systems. It enables organisations to benchmark their current capabilities and plan a progression path.

Level	Name	Description	Example
1	None	No human attribution in agent operations	Autonomous agents with API keys only
2	Declared	Human identity claimed but not verified	Agent config says 'owner: john@corp.com'
3	Verified	Biometric verification of authorising human	Facial + voice verification at agent creation
4	Propagated	Verified authority propagates through chains	Full delegation chain with scope narrowing
5	Proven	Complete provenance with regulatory compliance, physical coverage, and coalition governance	VAC Protocol Level 5 implementation

Most enterprise AI deployments today operate at Level 1 or Level 2. The VAC Protocol enables progression to Level 5, where every agent action is backed by complete cryptographic provenance from the authorising human through the entire delegation chain, across jurisdictions, physical systems, and coalition operations.

8. Business Model: Trust Graph as Pricing Engine

Revenue concentrates at trust boundaries, where risk is highest and verification is most valuable.

Interaction Zone	Verification Level	Cost to Serve	Revenue
Inside dense trust cluster	Lightweight (trust graph)	Low	Free/low tier
At trust boundary (new contact)	Full certified verification	Medium	Per-verification fee
Cross-organisation (cold)	Maximum assurance required	High	Premium verification
Cross-national (coalition)	Sovereign-grade compliance	Highest	Enterprise contract

The bigger a user's trust graph, the less they pay for routine interactions, creating a lock-in flywheel. The platform is sticky because most daily interactions are free. Revenue is generated at the boundaries where trust has not yet been established, which is precisely where verification has the highest value.

Athena's intelligence layer adds a second revenue dimension: SignalRank-powered routing is the premium product. Free users get basic model comparison. Pro users get personalised trust graphs that learn their preferences. Enterprise customers get sovereign routing with cryptographic audit trails.

9. Intellectual Property Portfolio

The combined Violet Shores IP portfolio comprises 458 claims across 10 filings, all Patent Pending with IP Australia.

Filing	Application Number	Focus	Claims
VAC Base	AU 2026901425	Core protocol, biometric verification, delegation	65

VAC Supp #1	AU 2026901428	Organisational trust, budget enforcement	25
VAC Supp #2	AU 2026901474	Physical systems, coalition governance	60
VAC Supp #3	AU 2026901553	Regulatory compliance, indigenous sovereignty	50
VAC Supp #4	AU 2026901601	Escalation chains, agent orchestration	45
VAC Supp #5	AU 2026901604	Continuous monitoring, collaborative onboarding	40
Athena Base	AU 2026901730	Self-improving methodology, product lifecycle	65
Athena Supp #1	AU 2026901851	SignalRank, multi-LLM orchestration	25
Athena Supp #2	AU 2026901870	Rail detection, trust graph topology	38
Athena Supp #3	AU 2026901871	Zero-knowledge preference routing, SDK	45

10. Standards Alignment and Regulatory Engagement

The VAC Protocol and Athena system align with an expanding landscape of AI agent standards. The regulatory trajectory is accelerating: NIST establishes voluntary standards, industry adoption becomes market expectation, sector-specific regulators incorporate standards into compliance requirements, and litigation increasingly cites NIST guidance as evidence of industry standard of care.

10.1 NIST AI Agent Standards Initiative (February 2026)

NIST launched the AI Agent Standards Initiative through CAISI on 17 February 2026. The Initiative operates across three pillars: facilitating industry-led agent standards development, fostering community-led open-source protocol development, and advancing research in AI agent security and identity. This is the umbrella program that coordinates all NIST agent-related activities.

Violet Shores submitted a formal response to the CAISI Request for Information on AI Agent Security (Docket NIST-2025-0035) on 8 March 2026, addressing 14 questions across threat models, mitigation approaches, secure design, and trust architectures.

10.2 NCCoE Concept Paper: AI Agent Identity and Authorization (February 2026)

The NIST National Cybersecurity Center of Excellence (NCCoE) released a concept paper in February 2026 titled 'Accelerating the Adoption of Software and AI Agent Identity and Authorization.' The paper proposes a demonstration project exploring how existing identity standards can be applied to AI agents in enterprise environments. Public comments are due 2 April 2026.

The NCCoE project focuses on identification (distinguishing agents from humans), authorization (applying OAuth 2.0 and policy-based access control), access delegation (linking user identities to agents for accountability), and logging and transparency (linking agent actions to non-human entities for auditability). The VAC Protocol addresses every one of these focus areas through its existing architecture, and Violet Shores intends to submit a response proposing VAC as a reference architecture for the demonstration project.

10.3 NIST Cybersecurity Framework Profile for AI (December 2025)

NIST published the preliminary draft of the Cybersecurity Framework Profile for Artificial Intelligence (IR 8596) in December 2025. The Cyber AI Profile maps the existing NIST CSF 2.0 framework to AI-specific considerations across three focus areas: Secure (managing cybersecurity when integrating AI systems), Defend (using AI to enhance cybersecurity capabilities), and Thwart (defending against adversarial AI). The VAC Protocol's conformance testing framework (Section 6) maps directly to the CSF 2.0 functions (Govern, Identify, Protect, Detect, Respond, Recover) referenced in this profile.

10.4 NIST SP 800-53 Control Overlays for Securing AI Systems (COSAiS)

NIST is developing control overlays that extend the SP 800-53 security control catalog to AI-specific use cases, including a dedicated overlay for single and multi-agent AI systems. The VAC Protocol's delegation chain architecture, scope narrowing enforcement, and trust score computation provide implementation mechanisms for controls in the access control, audit, and accountability families of SP 800-53.

10.5 Additional Standards Alignment

- **NIST Core:** SP 800-63 (identity assurance), SP 800-207 (zero trust architecture), AI 600-1 (AI risk management).
- **ISO:** 30107-3 (biometric anti-spoofing), 27001 (information security management), 29115 (entity authentication assurance), 24745 (biometric template protection).
- **Defence:** DoD Directive 3000.09 (autonomy in weapons systems), STANAG 4586 (UAV interoperability), Five Eyes/AUKUS coalition frameworks.
- **Data Sovereignty:** GDPR, LGPD, PIPL, PIPA, APPI, Privacy Act 1988, POPIA, DPDP Act, PIPEDA. Indigenous frameworks: Te Mana Raraunga, CARE Principles, OCAP, AIATSIS.

- **AI Safety:** OWASP Top 10 for LLMs, FIDO2/WebAuthn, W3C Decentralised Identifiers (DIDs), SOC 2 Type II, ETSI.
- **Cloud Security Alliance:** AI Controls Matrix (AICM), mapping cloud-native security controls to AI systems.

10.6 CAISI Listening Sessions (April 2026)

Beginning in April 2026, CAISI will hold virtual listening sessions on sector-specific barriers to AI adoption in healthcare, finance, and education, with a focus on AI agents. Violet Shores intends to participate to ensure the VAC Protocol's approach to human attribution is represented in sector-specific standards development.

---

11. Live Implementation

The system is not a specification. It is a running implementation with live endpoints:

Component	URL	Status
VAC Protocol (frontend)	vacprotocol.org	Live
VAC Backend API	vac-system-production.up.railway.app	Live
Athena Backend API	api.athenapilot.ai	Live
SignalRank Trust Graph	vacprotocol.org/intelligence	Live
D3 Trust Visualisation	vacprotocol.org/trust	Live
Biometric Authentication	vacprotocol.org/auth	Live
Agent Control + Budget	vacprotocol.org/agents	Live
Collective Governance	vacprotocol.org/groups	Live
MCP Server (10 tools)	api.athenapilot.ai/mcp/sse	Live
AI Internet Landing	aiinternet.ai	Live

The implementation demonstrates end-to-end functionality: biometric identity verification, VAT creation and delegation, trust graph computation, multi-LLM routing with SignalRank scoring, rail detection, agent budget enforcement with cryptographic audit trails, and collective governance with configurable quorum models.

---

12. Roadmap

Phase	Focus	Timeline
Phase 1 (Current)	Core protocol + Athena intelligence layer live	Q1 2026
Phase 2	NCCoE concept paper response + agent framework integrations	Q2 2026
Phase 3	Reasoning Path Index + knowledge graph	Q2-Q3 2026
Phase 4	Enterprise SDK + sovereign routing for defence	Q3 2026
Phase 5	NCCoE lab demonstration + standards certification	Q4 2026
Phase 6	Physical system authority + coalition operations	2027

Upcoming Deadlines

Date	Milestone	Status
8 March 2026	NIST CAISI RFI response (NIST-2025-0035)	Submitted
16-17 March 2026	LAUNCH Festival (San Francisco)	Video submitted
20 March 2026	CAISI Listening Session registration deadline	Pending
23 March 2026	VAC supplementary patent window	Open
2 April 2026	NCCoE concept paper response (AI Agent Identity)	Planned
April 2026	CAISI sector-specific listening sessions	Planned

13. Conclusion

The AI agent economy requires infrastructure that does not yet exist. Agents need verified identity. Routing needs trust-weighted intelligence. Reasoning needs provenance. And all of it needs to work across models, across organisations, and across borders.

VAC Protocol and Athena provide this infrastructure as a unified stack: identity at the base, trust computation above it, knowledge and reasoning in the middle, and intelligent orchestration at the top. The SignalRank engine ties every layer together through a trust-weighted ranking algorithm that improves with every interaction.

The Conformance Testing Framework and Attribution Maturity Model provide the industry with standardised tools to measure and benchmark human attribution capabilities, establishing a common language for what was previously an unmeasured security property.

The timing is significant. NIST has launched the AI Agent Standards Initiative. The NCCoE is developing a demonstration project for agent identity and authorization. Sector-specific listening sessions are beginning in April. The regulatory trajectory is clear: voluntary standards become market expectations become compliance requirements. Violet Shores is positioned at the

intersection of all three tracks with a live implementation, 458 patent claims, and direct engagement with the standards bodies shaping the future of AI agent governance.

vacprotocol.org | athenapilot.ai | aiinternet.ai
admin@violetshores.com