RAMPAGE

Technical Whitepaper

Decentralized Truth Verification & Constitutional Journalism

Version 1.5.1 | March 2026 | Aligned to Constitution v1.5

TruthOracle.ai | RampageNews.com
Founder: Shea Patrick Kastl, JD

Constitution v1.5 Effective: February 27, 2026
Available Upon Request — Strategic Partner Distribution

Executive Summary

Rampage is constitutional infrastructure for verifiable truth in high-risk information environments. It combines three integrated components: a constitutionally-governed journalism platform (RampageNews.com), a live forensic-grade consensus verification engine powered by the Seven-Seal Protocol (TruthOracle.ai), and a blockchain roadmap designed to make truth attestation immutable, decentralized, and resistant to state capture.

This whitepaper, version 1.5.1, supersedes all prior versions and is aligned to Constitution v1.5, effective February 27, 2026. It expands the technical documentation of three core innovations: the Seven-Seal Protocol adversarial consensus architecture, the Strategic Validator Agreement Framework established by Amendment No. 7, and the Mempool Shield consensus-layer regulatory enforcement mechanism.

TruthOracle.ai and RampageNews.com are operational today. The Constitution is ratified. The Seven-Seal Protocol is live, processing real verification queries in active conflict zones as of the date of this document. The blockchain development sprint begins Q2 2026.

The Seven-Seal Protocol is designed for model-agnostic operation. The seven AI systems documented herein reflect the current configuration as of March 2026. As the AI landscape advances, and it will advance rapidly, the protocol's architecture ensures that model substitution and expansion can occur without compromising constitutional consensus standards. The selection criteria for Truth Bearer models are permanent; the specific models are not.

Key Innovations

- Seven-Seal Protocol: Forensic-grade adversarial consensus across seven independent AI systems from seven different developers, preventing monoculture bias and ensuring no single commercial, national, or architectural interest controls verification outcomes.
- Constitutional Governance: Immutable foundational principles (Articles I-IV) protect the platform's core mission from any actor. Amendable articles (V-XI) allow adaptive governance through supermajority community consensus.
- Mempool Shield: Pre-consensus AML/CFT screening embedded at the protocol layer, not the application layer, making prohibited-entity transactions physically impossible on the blockchain.
- Strategic Validator Agreement Framework (Amendment No. 7): Enables mission-aligned institutional partners to participate as validators during the pre-mainnet phase under reduced staking thresholds, without permanently compromising the constitutional floor.
- Level 4 Conflict Mode: Constitutional operational framework for active armed conflict zones, currently active for Iran following Operation Epic Fury (February 28, 2026).
- 95/5 Operational Split: 95% of resources dedicated to truth verification and journalism infrastructure; 5% reserved for humanitarian crisis response is activated only at Threat Level 2 or higher.

1. Constitutional Foundation

Rampage is governed by a written Constitution, not a terms of service, not a governance token whitepaper, not a DAO charter. A written constitution with immutable articles, ratified amendments, and a dissolution clause that requires the platform to close before it compromises its principles.

Constitution v1.5, effective February 27, 2026, is the final constitutional baseline for testnet launch. It incorporates seven ratified amendments and reflects a governance architecture specifically designed to resist capture: by states, commercial interests, or any single actor including the founder.

Constitutional Architecture

Immutable Core (Articles I-IV): These articles encode the platform's foundational values and cannot be amended under any circumstances. They establish human rights primacy, the legal framework and risk philosophy, operational boundaries, and truth verification standards. If any action, including compliance with an international legal obligation, would require violating Articles I-IV, the constitutional response is dissolution, not compromise.

Amendable Implementation (Articles V-XI): These articles govern how Rampage operates and can be amended through supermajority community consensus. Amendment requires a minimum 5,000-word justification, 30-day community deliberation, 60% quorum, 80% supermajority approval, and a 90-day implementation delay.

Amendment History — v1.5

Amendment	Description
No. 1 — Dec 2, 2025	95/5 Operational Split. Article V, Section 1.1.
No. 2 — Jan 17, 2026	Truth Bearer Designation. Validator requirements, staking, governance rights. Article VII.
No. 3 — Jan 18, 2026	Data Protection and Privacy Framework (GDPR-aligned). Article V, Section 5.
No. 4 — Jan 18, 2026	Regulatory Classification and Compliance. RPM utility token, MiCA positioning. Article II, Section 7.
No. 5 — Jan 18, 2026	Enhanced Internal Controls and Governance Documentation. Article VI, Section 5.
No. 6 — Jan 18, 2026	Enhanced AML/CFT Documentation. Mempool Shield, EDD requirements. Article V, Section 4.
No. 7 — Feb 27, 2026	Strategic Validator Agreement Framework. Pre-mainnet flexible staking with sunset provisions. Article VII, Sections 7.4-7.7.

2. The Seven-Seal Protocol: Forensic-Grade Adversarial Consensus

The Seven-Seal Protocol is the technical core of TruthOracle.ai. It is not a marketing label for a multi-LLM API call. It is a constitutionally-governed adversarial consensus architecture designed to produce forensic-grade verification verdicts that no single AI system, developer, government, or commercial interest can unilaterally control or corrupt.

The protocol takes its name from its seven independent verification seals: seven AI systems, seven different developers, seven distinct reasoning architectures that must reach constitutional consensus before TruthOracle issues a verification verdict. Breaking one seal does not break the chain. Compromising one model does not compromise the outcome. The architecture is designed for adversarial resilience from the ground up.

The Monoculture Problem

The central risk in AI-based verification is monoculture bias. This is the tendency of systems trained on similar data, by similar teams, using similar methodologies to produce correlated errors. A verification system that runs the same claim through seven instances of the same model has not achieved diversity. It has achieved the illusion of diversity while remaining vulnerable to a single point of architectural failure.

The Seven-Seal Protocol solves this by selecting models across four dimensions of diversity: developer geography (US, European, Chinese AI ecosystems represented), reasoning architecture (retrieval-augmented generation, chain-of-thought, real-time social scouting, standard completion), training data provenance (different internet corpora, different cutoff dates, different fine-tuning approaches), and institutional independence (no two models built by the same parent organization).

The Seven Seals; Current Configuration (March 2026)

The following seven AI systems constitute the current Seven-Seal Protocol configuration. This configuration reflects the best available adversarial diversity as of March 2026. The protocol is explicitly designed for model-agnostic operation, and as AI systems evolve, are deprecated, or are superseded, the constitutional selection criteria govern substitution, not attachment to any specific model version.

Model	Developer	Verification Role
Perplexity	Perplexity AI (US)	Primary news aggregator and source screener. Real-time web retrieval with citation. Acts as the first-pass evidence collector before consensus.
GPT-4o-mini	OpenAI (US)	Broad knowledge synthesis and cross-referencing. High-speed general reasoning across large context windows.
Claude Haiku	Anthropic (US)	Constitutional reasoning and structured claim analysis. Selected in part for Anthropic's demonstrated commitment to AI safety principles.

Gemini Flash	Google DeepMind (US)	Rapid multi-source fact-checking. Multimodal capability for image and document verification where applicable.
Mistral	Mistral AI (France / EU)	European AI ecosystem representation. GDPR-jurisdiction training and fine-tuning. Critical for operations in EU regulatory environments.
Grok-4.1 (X-Search)	xAI (US)	Real-time social corroboration layer. Scouts X (Twitter) for primary witness reports, on-ground accounts, and real-time social signal using the 2026 Responses API.
DeepSeek R1	DeepSeek (China)	Chain-of-thought reasoning architecture. Step-by-step logical decomposition of complex claims. Provides adversarial depth against the other models' breadth.

Grok-4.1 as Real-Time Social Layer

The integration of Grok-4.1 with X-Search represents a qualitatively different verification modality than the other six seals. While the remaining models reason from indexed knowledge, retrieved documents, and training data, Grok-4.1 scouts X in real time using the 2026 Responses API, identifying primary witness accounts, on-ground reports, and social signals that have not yet been indexed by any other source.

This matters in crisis verification. When BGP routing data shows a communications disruption in Tehran, Grok-4.1 cross-references that technical signal against live social accounts from within the affected area. When casualty figures are disputed, Grok-4.1 identifies first-hand accounts before they are filtered through state media or international outlets. This is not social media aggregation. It is constitutionally-governed primary source scouting with a chain-of-custody audit trail.

DeepSeek R1 as Adversarial Depth Layer

DeepSeek R1's chain-of-thought reasoning architecture provides a fundamentally different verification modality from the other six seals. Where retrieval-augmented models assess claims against retrieved evidence, DeepSeek R1 reasons through claims step by step, decomposing the logical structure of an assertion, identifying internal inconsistencies, and stress-testing the evidentiary chain before rendering an assessment.

In practical terms, if a claim contains a logical contradiction that retrieved evidence would not catch. a timeline inconsistency, a geographic impossibility, an attribution that conflicts with documented facts, DeepSeek R1's chain-of-thought process surfaces it. This is adversarial depth, not adversarial breadth. Together, the seven seals provide both.

Consensus Calculation

Each of the seven models independently analyzes a verification query and produces: a binary or graduated accuracy assessment, supporting evidence with source citations, a confidence score for its own assessment, and identification of ambiguities or insufficient information. The aggregation layer synthesizes these outputs, weights responses based on model confidence and source quality, and calculates a composite consensus score.

Level	Threshold / Sources	Application
Level 1	60% / 3+ sources	Standard claims, general information, routine verification
Level 2	67% / 5+ sources	Sensitive claims, public figures, policy matters, protest events
Level 3	80% / 7+ sources	Critical claims, conflict casualties, security implications, high-stakes decisions

Geographic diversity is constitutionally required across the source pool. Not all corroborating sources may originate from a single country or institution. State media alone is constitutionally insufficient; independent corroboration is mandatory. Every verification query and result is logged in a complete, immutable audit trail.

Model Selection Criteria — Constitutional Standards

The constitutional standards for Seven-Seal model selection are permanent even as specific models evolve. Any model operating as a Truth Bearer seal must satisfy:

- Developer Independence: No two active seals may be developed by the same parent organization.
- Geographic Diversity: The developer ecosystem must represent at least two distinct national AI development environments.
- Architectural Diversity: The active configuration must include at minimum one retrieval-augmented model, one chain-of-thought model, and one real-time corroboration model.
- Commercial Independence: No single seal may have a commercial relationship with Rampage that would create incentive to bias verification outcomes.
- Constitutional Alignment: Models developed by organizations with documented commitments to AI safety and responsible deployment are preferred.

3. Strategic Validator Agreement Framework (Amendment No. 7)

Amendment No. 7, ratified February 27, 2026, establishes the Strategic Validator Agreement (SVA) framework — the constitutional mechanism enabling Rampage to recruit mission-aligned institutional partners as validators during the pre-mainnet phase without permanently lowering the network's security and governance standards.

The Bootstrap Problem

The standard Truth Bearer staking minimum is 100,000 RPM tokens. This threshold is designed to ensure validators have meaningful economic skin-in-the-game before participating in governance and verification consensus. It is the right threshold for a mature, liquid network. It is prohibitive for the early recruitment of exactly the kind of institutional partners Rampage most urgently needs: journalism organizations, press freedom bodies, human rights monitors, and UN-affiliated NGOs.

These organizations cannot accumulate 100,000 RPM before the network launches. Their participation in the validator set is strategically essential for geographic distribution requirements, institutional credibility, and the third-party validation that makes TruthOracle trustworthy to the populations it serves. Amendment No. 7 resolves this tension constitutionally.

SVA Structure and Eligibility

During the Pre-Mainnet Phase (February 27, 2026 through mainnet genesis block, targeted Q4 2026), the Founder or Truth Bearer Council may execute Strategic Validator Agreements with qualified partners specifying a reduced minimum staking threshold. All SVAs must satisfy:

- Demonstrated institutional alignment with truth verification, journalistic integrity, or humanitarian crisis response.
- A written agreement specifying agreed minimum stake, duration, and transition conditions.
- Full compliance with all other Truth Bearer obligations including uptime, governance participation, and Mempool Shield AML/CFT compliance.
- No conflict with the immutable Articles I-IV.
- Public disclosure in the governance ledger within 30 days of execution.

Qualifying institutional categories include accredited journalism organizations and press freedom bodies; internationally recognized human rights organizations; UN agencies and affiliated NGOs; academic institutions with documented human rights or journalism research programs; and other mission-aligned entities approved by governance.

Sunset Provisions and Constitutional Safeguards

The 100,000 RPM floor is preserved absolutely. SVAs are temporary, bilateral, contractual arrangements, not precedent for permanent reduction. Upon mainnet launch, the standard threshold applies automatically and universally. SVA validators must reach 100,000 RPM within six months of mainnet deployment or lose validator status.

Any extension of the pre-mainnet SVA regime beyond mainnet launch requires a 67% supermajority of all staked tokens, must specify a defined duration not exceeding 12 months per extension, and mandates governance review at least once every six months. Pre-mainnet flexibility cannot drift into permanent institutional exception.

4. Mempool Shield: Consensus-Layer Regulatory Defense

The Mempool Shield is the most important technical compliance control in the Rampage architecture. It is not a content moderation policy. It is not an application-layer filter that can be administratively disabled. It is a hard-coded technical enforcement mechanism embedded at the blockchain's consensus layer making prohibited-entity transactions physically impossible on the Rampage network.

Architecture: Why Consensus Layer Matters

Most blockchain compliance controls operate at the application layer, a user interface or API wrapper that screens transactions before submitting them to the chain. Application-layer controls can be bypassed: by routing around the interface, by submitting transactions directly to the mempool, or by a governance decision to disable screening. They are policies, not physics.

The Mempool Shield operates differently. It is integrated at the consensus layer, the protocol level at which validators decide which transactions are valid and eligible for inclusion in a block. A transaction to a prohibited entity address is not merely rejected by an interface. It is rejected by the consensus mechanism itself. No validator can include it in a block. No governance vote can override it short of a full constitutional amendment requiring 60% quorum, 80% supermajority, and a 90-day implementation delay. The blockchain is physically incapable of executing prohibited transactions.

The 5-of-7 Oracle Signer Committee

The prohibited-entity oracle feed, drawing from OFAC SDN, the UN Consolidated List, EU sanctions lists, and FATF designations, is continuously updated from authoritative sources. Feed accuracy and timeliness are verified by a 5-of-7 trusted oracle signer committee appointed by governance.

The signer committee is a critical constitutional safeguard that is frequently overlooked in descriptions of the Mempool Shield. It addresses a second-order risk: that the prohibited-entity lists themselves could be weaponized, expanded to include legitimate humanitarian actors, political opponents of sanctioning governments, or organizations whose work conflicts with powerful state interests. The 5-of-7 threshold ensures that no single actor, including any sanctioning government, can unilaterally alter the feed that governs what the Mempool Shield blocks. The lists protect the network. The committee protects the lists.

Fail-Safe Default

If the oracle feed becomes unverifiable or cannot be authenticated by the signer committee, the Mempool Shield does not fail open. It defaults to a fail-safe state: all capital routing transactions are rejected until feed integrity is restored. This is a deliberate constitutional choice. The risk of temporarily halting legitimate humanitarian transactions is accepted as preferable to the risk of a single prohibited transaction clearing during a feed integrity failure.

This fail-safe architecture is what allows Rampage to make the strongest possible compliance claim to institutional partners: the network cannot route funds to prohibited entities even if governance were compromised or coerced. It is not a policy commitment. It is a technical guarantee.

Privacy-Preserving Compliance

The Mempool Shield uses zero-knowledge proofs for compliant transaction verification, confirming that a recipient address does not appear on prohibited-entity lists without revealing the transaction details or recipient identity to the screening mechanism. This preserves the constitutional commitment to civilian anonymity while maintaining the integrity of the compliance framework.

5. Operational Jurisdiction: Level 4 Conflict Mode

On February 28, 2026, the United States and Israel launched major combat operations against Iran , Operation Epic Fury. Iranian forces retaliated with strikes against US military assets across the Gulf region. This triggered the Rampage Constitution's maximum threat level classification for Iran: Level 4, Critical.

Iran's elevation to Level 4 is the first activation of maximum constitutional parameters in Rampage's operational history. It represents both the platform's most demanding test case and its most consequential use case, and TruthOracle was operational before the first strike landed.

What Level 4 Activates

Under Level 4 — Critical, the following constitutional parameters apply:

- Maximum verification activity: All Seven-Seal Protocol resources prioritized toward Iranian conflict theater verification.
- Full Level 3 operational authorities activated: Expanded capital routing to civilians, emergency NGO coordination, maximum information penetration efforts.
- Potential financial transaction suspension: Capital routing may be suspended pending enhanced screening given the extreme sanctions environment.
- Governance review required within 30 days: The Truth Bearer Council must formally assess and document the Level 4 designation.
- 95/5 enforcement maintained: Even at Level 4, 95% of resources remain dedicated to verification infrastructure. The humanitarian 5% does not expand beyond constitutional limits.

Information Warfare in Active Conflict

The information warfare dimension of the US-Iran-Israel conflict is precisely what the Seven-Seal Protocol was architected to address. Both sides in an active armed conflict have strategic incentives to control the information environment: to minimize reported civilian casualties on their own side, maximize reported casualties on the opposing side, justify targeting decisions, and shape international public opinion.

TruthOracle's constitutional constraints make it structurally resistant to this pressure. It verifies observable facts: casualties documented through hospital records and witness testimony, communications disruptions measured via BGP routing data, infrastructure damage confirmed through satellite imagery and multi-source corroboration. It cannot determine which side is morally justified. It cannot attribute responsibility. It publishes facts, not verdicts.

Grok-4.1's real-time X-Search scouting is particularly valuable in this environment by identifying primary witness accounts from within Iran before they are filtered through state media on either side. DeepSeek R1's chain-of-thought reasoning stress-tests casualty claims and timeline assertions for internal logical consistency. The remaining five seals provide breadth across retrieved evidence from multiple geographic and linguistic source pools.

Constitutional Prohibitions That Apply in Level 4

The Constitution's prohibitions apply with full force regardless of threat level. At Level 4, Rampage explicitly does not:

- Support, coordinate with, or provide resources to any armed group — US, Israeli, Iranian, or otherwise.
- Route capital to IRGC, Iranian government entities, US military, Israeli military, or any armed actor.
- Take sides in the conflict or issue political or moral judgments about any party's conduct.
- Accept direction from any government regarding what to verify or how to report it.

These are not editorial policies subject to crisis exception. They are constitutional architecture. Level 4 expands operational capacity for civilian verification and humanitarian routing. It does not relax the nonviolence commitment or the civilian-only focus by a single provision.

6. Mission Model and Operational Split

Rampage operates under a constitutionally mandated 95/5 resource allocation framework enforced through on-chain treasury controls and quarterly governance reporting. This is not a policy preference, it is constitutional law for the network.

- Core Function (95%): Truth Verification via TruthOracle operations and passive Information Access infrastructure. Active at all Threat Levels.
- Crisis Function (5%): Active Information Penetration and Humanitarian Economic Access. Activates only at Threat Level 2 and higher. Strictly limited to civilian, nonviolent, humanitarian purposes.

The Governance Module calculates and publishes the actual allocation ratio on-chain every calendar quarter. Persistent deviation exceeding 18 consecutive months triggers a mandatory governance proposal to amend the framework.

7. Legal, Compliance, and Risk Posture

Rampage does not claim legal immunity. This is a constitutional value, not a disclaimer. All participants are responsible for understanding and complying with applicable law in their jurisdictions. Rampage provides legal risk information, constitutional defenses, and a legal defense fund, but cannot prevent prosecution or regulatory action.

AML/CFT Framework

- Mempool Shield: Pre-consensus screening embedded at the protocol layer. Cannot be bypassed by governance, emergency powers, or administrative action.
- 5-of-7 Oracle Signer Committee: Governs the integrity of prohibited-entity feed sources.
- Multi-signature treasury controls: Scales from 5-of-9 (under $5M) to 21-of-39 (over $100M) with geographic distribution requirements.
- Enhanced Due Diligence: Mandatory for all transfers over $5,000 and all recipients in Level 3+ crisis zones.
- Quarterly third-party audits: Results published in transparency reports.

Validator Protection

The constitutional Legal Defense Fund reserves 5% of the community treasury of 1,050,000 RPM which is permanently for validator legal defense. Validator anonymity is structured in three tiers: mandatory in high-risk jurisdictions (Tier 1), optional in medium-risk jurisdictions (Tier 2), and standard transparency in low-risk jurisdictions (Tier 3). Zero-knowledge proofs verify tier membership without revealing specific location.

Regulatory Classification

RPM is a utility token — governance and validator staking, not investment return. Rampage takes the position that RPM is not a security, asset-referenced token (ART), or e-money token (EMT) under EU MiCA. Rampage is a decentralized protocol, not a traditional Crypto-Asset Service Provider.

8. RPM Token Economics and Governance

- Total Supply: 21,000,000 RPM (fixed, no inflation)
- Legal Defense Fund: 1,050,000 RPM (5% of treasury, constitutionally reserved)
- Standard Truth Bearer Threshold: 100,000 RPM
- Governance Proposal Threshold: 10,000 RPM minimum stake
- Treasury Self-Custody: Ledger + Trezor hardware wallets, diversified across BTC and five stablecoins

Constitutional governance voting thresholds:

- Standard proposals: 51% quorum, 67% approval
- Constitutional amendments: 60% quorum, 80% supermajority
- Humanitarian escalation: 75% of total voting power, 50% of all circulating RPM, 30-day vote, 90-day delay
- Emergency legal compliance: 40% quorum, 75% approval, 7-day process

9. Blockchain Roadmap

Current Phase: Pre-Blockchain MVP (Q1 2026)

- TruthOracle.ai is live, Seven-Seal Protocol operational
- RampageNews.com is live, constitutional journalism workflow
- Constitution v1.5 is ratified, seven amendments incorporated
- n8n workflow automation for verification pipeline
- PHP proxy security layer and SSL infrastructure
- Self-custodied treasury across hardware wallets

Phase 2: Testnet (Q2-Q3 2026)

- Deploy Rampage blockchain testnet, a custom L1, EVM-compatible, Cosmos SDK
- Implement RPM token smart contracts with constitutional governance hooks
- Release Truth Bearer node software
- Integrate Mempool Shield at consensus layer for testnet validation
- Execute Seven-Seal consensus protocol trials across distributed validator set
- Commission independent security audits and penetration testing
- Onboard initial SVA partners under Amendment No. 7

Phase 3: Mainnet Launch (Q4 2026)

- Mainnet deployment with genesis block and constitutional bootstrap
- RPM token distribution and treasury activation
- Full TruthOracle decentralization for verification results written to chain
- On-chain governance activation
- Mempool Shield production deployment with live oracle feed
- SVA validators required to reach 100,000 RPM within 6 months

Phase 4: Ecosystem Expansion (2027+)

- Third-party verification integrations
- Developer SDK and API releases
- Mobile applications for high-risk environments
- Global Truth Bearer network expansion

Technology Stack

Layer	Technologies
Frontend	React, TailwindCSS, responsive design optimized for low-bandwidth environments
Backend	Node.js, n8n workflow automation, PHP proxy security layer
AI / ML	Seven-Seal Protocol: Perplexity, GPT-4o-mini, Claude Haiku, Gemini Flash, Mistral, Grok-4.1 (X-Search), DeepSeek R1 (chain-of-thought)
Blockchain	Custom L1 (in development), EVM compatibility, Cosmos SDK infrastructure, IBC interoperability
Security	End-to-end encryption, zero-knowledge proofs (zk-SNARK for anonymity tier and compliance verification), Mempool Shield consensus-layer screening
Privacy	Tor and VPN for Tier 1 validators, hashed on-chain identifiers, shielded pool option (Zcash-style), zero-knowledge compliance proofs

10. Strategic Partner Integration Model

Rampage is built for partnership — not as a commercial proposition but as a mission alignment. All partnerships are subordinated to constitutional guardrails. No partnership can modify, waive, or circumvent Articles I-IV.

Partnership Categories

- Verification Partners: Organizations integrating TruthOracle into editorial workflows or contributing source material to the consensus process.
- Technical Infrastructure Partners: Organizations providing satellite internet, mesh networking, or secure communications extending Rampage's reach into high-risk environments.
- Validator Partners (SVA Framework): Mission-aligned institutional organizations participating under Amendment No. 7 during the pre-mainnet phase. All SVA terms publicly disclosed.
- Humanitarian and Legal Support Partners: UN agencies, ICRC, press freedom organizations, and human rights bodies providing legal defense resources, jurisdiction risk assessments, and humanitarian corridor access.

Active Partnership Conversations — March 2026

Strategic partnership letters have been transmitted to Anthropic (Dario Amodei, CEO), UN OCHA, ICRC, Reporters Without Borders, and additional organizations throughout March 2026. All partnership engagements are subordinated to constitutional requirements and subject to public disclosure under the SVA framework where applicable.

11. Conclusion

Rampage Whitepaper v1.5.1 documents a platform that is no longer theoretical. TruthOracle.ai processed real verification queries during active US-Iran-Israel combat operations on February 28, 2026. The Seven-Seal Protocol ran seven independent AI systems in adversarial constitutional consensus while state actors on multiple sides generated competing, unverifiable information claims. The Constitution held. The architecture performed.

The immutable core — Articles I-IV — cannot be amended. The Mempool Shield cannot be bypassed. The civilian-only focus cannot be voted away. The Seven-Seal Protocol's model-agnostic design ensures that as AI systems evolve, the constitutional verification standards remain constant even as the specific seals change.

Testnet launches Q2-Q3 2026. Mainnet Q4 2026. The window for foundational partnership is open now.

"Truth is not determined by majority vote, but it can be verified through consensus."

Technical Appendix: Seven-Seal Protocol Implementation

This appendix documents the technical implementation of the Seven-Seal Protocol for developer review, institutional due diligence, and partner technical evaluation. All schemas represent the current pre-blockchain MVP implementation. Blockchain integration will extend these schemas with on-chain attestation fields at testnet launch.

A. Verification Query JSON Schema

The following schema defines the structure of a TruthOracle verification request. All queries entering the Seven-Seal Protocol conform to this schema:

```
{
  "query_id": "string (UUID v4)",
  "timestamp": "string (ISO 8601 UTC)",
  "claim": "string (natural language claim for verification)",
  "jurisdiction": "string (ISO 3166-1 alpha-2 | 'GLOBAL')",
  "threat_level": "integer (0-4)",
  "verification_level": "integer (1-3)",
  "required_threshold": "float (0.60 | 0.67 | 0.80)",
  "minimum_sources": "integer (3 | 5 | 7)",
  "seals_requested": [
    "perplexity", "gpt4o_mini", "claude_haiku",
    "gemini_flash", "mistral", "grok_41", "deepseek_r1"
  ],
  "geographic_diversity_required": "boolean",
  "audit_trail_required": "boolean (always true)",
  "requestor_tier": "string (PUBLIC | VALIDATOR | INSTITUTIONAL)"
}
```

B. Seal Response JSON Schema

Each of the seven seals returns an independent response conforming to the following schema. Responses are collected by the aggregation layer before consensus calculation:

```
{
  "seal_id": "string (model identifier)",
  "query_id": "string (UUID v4 — matches request)",
  "timestamp": "string (ISO 8601 UTC)",
  "assessment": {
    "verdict": "string (VERIFIED | UNVERIFIED | DISPUTED | INSUFFICIENT_DATA)",
    "confidence_score": "float (0.00 - 1.00)",
    "reasoning": "string (model reasoning — chain-of-thought for DeepSeek R1)"
  },
  "evidence": [
    {
      "source_id": "string (UUID v4)",
      "source_url": "string",
      "source_type": "string (PRIMARY | SECONDARY | TECHNICAL | SOCIAL)",
      "source_geography": "string (ISO 3166-1 alpha-2)",
      "source_language": "string (ISO 639-1)",
      "retrieval_timestamp": "string (ISO 8601 UTC)",
      "relevance_score": "float (0.00 - 1.00)"
    }
  ],
  "ambiguities": ["string array — identified uncertainties"],
  "social_signals": {
    "applicable": "boolean (true only for grok_41)",
    "x_search_results": "integer",
    "corroborating_accounts": "integer",
    "conflicting_accounts": "integer"
  }
}
```

C. Consensus Aggregation Schema

The aggregation layer collects all seven seal responses and calculates the consensus verdict according to constitutional thresholds:

```
{
  "consensus_id": "string (UUID v4)",
  "query_id": "string (UUID v4 — matches request)",
  "timestamp": "string (ISO 8601 UTC)",
  "seals_received": "integer (max 7)",
  "seals_agreeing": "integer",
  "consensus_percentage": "float (0.00 - 1.00)",
  "threshold_required": "float (0.60 | 0.67 | 0.80)",
  "threshold_met": "boolean",
  "composite_confidence": "float (weighted average of seal confidence scores)",
  "geographic_diversity_met": "boolean",
  "verdict": "string (VERIFIED | UNVERIFIED | DISPUTED | INSUFFICIENT_DATA)",
  "verification_level": "integer (1-3)",
  "source_count": "integer (total unique sources across all seals)",
  "source_geography_distribution": {
    "country_code": "integer (source count per country)"
  },
  "audit_trail_hash": "string (SHA-256 hash of complete query + responses)",
  "on_chain_attestation_ready": "boolean (true when blockchain active)"
}
```

D. PHP Proxy Security Layer

TruthOracle.ai's PHP proxy layer provides API key security, rate limiting, and request authentication for the pre-blockchain MVP. In production blockchain deployment, this layer transitions to a validator node authentication system with on-chain key management. The proxy handles:

- API Key Isolation: All LLM API keys are stored server-side. Browser-facing interfaces never receive or transmit API credentials. The proxy authenticates requests before routing to LLM endpoints.
- Rate Limiting: Constitutional query tiers enforce rate limits: PUBLIC tier (standard rate), VALIDATOR tier (elevated rate with Truth Bearer authentication), INSTITUTIONAL tier (dedicated capacity with SVA authentication).
- Request Sanitization: All incoming queries are sanitized and validated against the Query Schema (Section A) before routing to the Seven-Seal Protocol.
- Audit Logging: Every request, response, and consensus calculation is logged with timestamp, query ID, and SHA-256 hash for the immutable audit trail. Logs are append-only and cannot be modified after creation.
- Fail-Safe Routing: If fewer than three seals return valid responses, the query is flagged INSUFFICIENT_DATA and does not proceed to consensus calculation. This prevents artificially high consensus scores from small sample sizes.

Core proxy request routing logic:

```
function routeToSeals(query, sealsRequested) {
  const responses = [];
  const minRequired = getMinSeals(query.verification_level);

  for (const seal of sealsRequested) {
    try {
      const response = await callSealEndpoint(seal, query);
      responses.push(validateSealResponse(response));
    } catch (error) {
      logSealFailure(seal, query.query_id, error);
    }
  }

  if (responses.length < minRequired) {
    return { verdict: 'INSUFFICIENT_DATA',
             seals_received: responses.length };
  }

  return calculateConsensus(responses, query.required_threshold);
}
```

E. On-Chain Attestation Schema (Testnet Target)

At testnet launch, the consensus output will be written to the Rampage blockchain as an immutable attestation record. The on-chain schema extends the consensus schema with blockchain-specific fields:

```
{
  "attestation_id": "string (UUID v4)",
  "block_height": "integer",
  "block_hash": "string (SHA-256)",
  "consensus_id": "string (UUID v4 — links to consensus record)",
  "validator_signatures": [
    {
      "validator_id": "string (hashed — anonymity preserved)",
      "validator_tier": "integer (1-3)",
      "signature": "string (validator cryptographic signature)",
      "timestamp": "string (ISO 8601 UTC)"
    }
  ],
  "mempool_shield_cleared": "boolean",
  "correction_history": [
    {
      "correction_id": "string (UUID v4)",
      "correction_timestamp": "string (ISO 8601 UTC)",
      "original_verdict": "string",
      "corrected_verdict": "string",
      "correction_reasoning": "string",
      "correction_block_height": "integer"
    }
  ],
  "legal_disclaimer_hash": "string (SHA-256 hash of applicable disclaimer)",
  "immutable": true
}
```

All on-chain attestations are permanent and cannot be deleted. Corrections are recorded as additional entries in the correction history array, the original attestation is preserved alongside the correction, maintaining complete transparency about what was published and what was subsequently found to be inaccurate. This is a constitutional requirement, not a technical preference.

Contact and References

- TruthOracle Verification Engine: truthoracle.ai
- RampageNews Platform: rampagenews.com

- Governing Document: Rampage Constitution v1.5 (effective February 27, 2026)
- Founder: Shea Patrick Kastl, JD
- Technical and Partnership Inquiries: Contact through official channels

Rampage Whitepaper v1.5.1 | March 2026 | Aligned to Constitution v1.5 | Available Upon Request — Strategic Partner Distribution