The Generative Agentic Network: A White Paper

A Decentralised Protocol for Semantic AI Agent Communication

Author: Extropy.IO
Date: 10 October 2025

Abstract

This white paper introduces the Generative Agentic Network (GAN), a decentralised protocol designed to facilitate a paradigm shift in communication between autonomous AI agents—from rigid, syntactic interactions to a fluid, semantic model based on shared intent[1]. The core of the protocol is the **Generative Document Representation (GDR)**, a compact Large Language Model (LLM) prompt that serves as the canonical, generative source for a complex document or task[2222]. This paper outlines the architecture of a trustless, open marketplace for generative tasks, built upon the existing Decentralised AI (DeAI) ecosystem[3]. We address the fundamental technical challenges inherent to this concept, including the non-determinism of LLM inference [4444], the prohibitive cost of verifiable computation [555], and the ambiguity of using natural language as a formal specification[666]. We propose novel solutions, including a **Decentralised Adjudication Network (DAN)** for verifying "semantic equivalence" rather than identical regeneration [77], a tiered verification system to balance economic viability with trust [88], and an architectural reframing of the "prompt as blueprint" within a binding smart contract "shell"[9]. The paper concludes with a strategic roadmap for the research, development, and deployment of the protocol.

---

1. Introduction: The Need for a Semantic Protocol

The current landscape of machine-to-machine communication is dominated by syntactic, API-driven models that are ill-suited for the fluid, intent-driven interactions of modern AI agents[10]. These legacy protocols are brittle and require constant maintenance to integrate fragmented systems[11]. The Generative Agentic Network proposes a forward-looking alternative by leveraging LLMs to create a new communication primitive based on semantic

intent[12].

The protocol's core component is the Generative Document Representation (GDR), a compact LLM prompt that functions as the generative "source code" for a complex task[13]. Instead of exchanging large data files, agents exchange a succinct GDR[14]. A receiving agent then uses a specified LLM to regenerate the full-body document or execute the task[15]. This shifts agent communication from a declarative model of exchanging data to a generative model of exchanging instructions, promising significant gains in efficiency and flexibility[16].

This protocol is positioned not as a replacement for emerging standards like ACP or A2A, but as a high-level **payload protocol** that defines the semantic content being exchanged[17171717].

2. The Generative Agentic Network: An Architectural Overview

The GAN protocol is designed as a complete, self-sustaining decentralised ecosystem that facilitates a full lifecycle for generative tasks, from creation to settlement.

2.1 The End-to-End Workflow

A typical interaction on the network follows five distinct phases:

1. **Task Definition:** A **Client Agent** defines a task and encapsulates it within a GDR Packet. This packet is a comprehensive JSON object containing the prompt, model requirements, reward, deadline, and verification criteria.
2. **Job Posting:** The Client Agent submits the GDR Packet to the GAN smart contract and deposits the reward into escrow. The job's public details are broadcast, allowing worker agents to discover it.
3. **Execution:** A qualified **Worker Agent** accepts the job, locking a stake as a performance bond. It fetches the exact, versioned LLM from a decentralised storage network like IPFS, executes the GDR within a secure compute environment (DePIN), and submits the result back to the smart contract. For privacy, this execution can occur within a Trusted Execution Environment (TEE)[18].

4. **Verification:** The smart contract initiates verification according to the GDR Packet's

specification. This triggers the **Decentralised Adjudication Network (DAN)**, a pool of specialised, staked "Adjudicator Agents" who vote on the semantic quality of the output.

5. **Settlement:** Based on the consensus of the DAN, the smart contract finalises the job. If successful, it automatically distributes the reward to the Worker Agent and the Adjudicators, and penalises any adjudicators who voted against the consensus.

2.2 The GDR Packet Specification

The GDR Packet is the atomic unit of work on the network. It is a self-contained digital contract designed for autonomous execution.

JSON

```
{
  "protocol_version": "3.0",
  "gdr_prompt": "Generate a detailed patent infringement risk analysis...",
  "model_spec": {
    "name": "LegalPro-Instruct-v4",
    "version_hash": "ipfs://QmX...Z7n"
  },
  "output_spec": { "format": "text/markdown" },
  "reward": { "amount": "0.5_ETH", ... },
  "deadline_timestamp": "1760272800",
  "privacy": { "level": "private_io", "tee_required": true },
  "attestations_required": ["LegalAnalysis_L3_Certified"],
  "verification_spec": {
      "verification_tier": "Tier_1_DAN",
      "adjudicator_count": 5,
      "adjudicator_attestation": "LegalVerify_L2",
      "pass_threshold": 4,
      "semantic_criteria": "Verify cosine similarity of embeddings > 0.95 against target."
  },
  "callback": { "target_address": "0xClientCorp...", ... }
}
```

3. Addressing Key Technical Challenges

The design of the GAN protocol is rooted in proactively solving the critical challenges identified in foundational research.

3.1 From Identical Regeneration to Semantic Equivalence

A core finding of the research is the "determinism fallacy": setting an LLM's `temperature=0` **does not** guarantee bit-for-bit identical outputs due to the nature of parallelised floating-point math and MoE architectures[19191919]. This makes simple hash-based verification impossible.

Solution: The GAN protocol pivots from verifying "identical regeneration" to verifying "semantic equivalence"[20202020]. The smart contract doesn't check if two outputs are identical; it checks if the submitted output is *semantically faithful* to the original intent[21]. This is achieved through objective criteria defined in the GDR Packet, such as:

- **Embedding Similarity:** Requiring the cosine similarity between the output's embedding and a target embedding to be above a threshold, building on the Semantic Reconstruction Effectiveness (SRE) metric[22222222].
- **Schema Validation:** Ensuring structured data outputs conform to a specified schema[23].
- **LLM-as-a-Judge Checklist:** Using a trusted model to verify if the output meets a list of qualitative attributes[24].

3.2 The Decentralised Adjudication Network (DAN)

To implement semantic equivalence verification in a trustless manner, we introduce the Decentralised Adjudication Network (DAN). This system replaces a single, centralised "judge" with a dynamic marketplace of specialised, independent adjudicator agents. The process is

governed by a reputation-staked consensus mechanism:

1. **Selection:** A pool of adjudicators with the required attestations is selected.
2. **Staking:** Each adjudicator locks a reputation token as collateral to participate.
3. **Judgement:** Adjudicators independently evaluate the work and submit a pass or fail vote.
4. **Consensus & Slashing:** Agents in the majority are rewarded, while agents in the minority are penalised by having their stake "slashed." This creates a powerful economic incentive for objective and high-quality verification.

The DAN directly mitigates the risks of centralisation, subjectivity, and adversarial attacks inherent in a single-judge model.

3.3 Tiered Verification: Balancing Trust and Cost

Research clearly shows that the computational overhead of Zero-Knowledge Machine Learning (zkML) is currently prohibitive for most large-scale LLM inference tasks[252525252525252525]. Mandating a full ZK proof for every task would make the network economically unviable.

Solution: The GAN protocol implements a **"Tiered Verification"** system, allowing the client to select a level of trust commensurate with the task's value[26262626].

- **Tier 0 (Reputation-Based):** For low-value tasks, payment is released based on the worker's on-chain reputation[27].
- **Tier 1 (DAN and/or TEE-Based):** The default for most tasks. Verification is provided by the DAN or a cryptographic attestation from a Trusted Execution Environment (TEE)[28282828].
- **Tier 2 (Full Cryptographic Proof):** Reserved for high-value, privacy-critical tasks requiring a full zk-SNARK proof of execution[29292929].

3.4 From "Prompt as Contract" to "Prompt as Blueprint"

Treating a natural language prompt as a formal contract is untenable due to its inherent ambiguity, vulnerability to prompt injection attacks, and the lack of legal guarantees from LLM providers[303030303030303030].

Solution: The protocol's architecture separates the formal agreement from the semantic instructions[31].

- The **Smart Contract** is the true, legally and technically binding agreement. It defines the immutable terms: parties, reward, deadline, and the objective verification conditions[32323232].
- The **GDR Prompt** is the **blueprint** or specification for the work. It is the subject of the contract, not the contract itself[33333333].

This separation of concerns creates a robust system where the smart contract provides the enforceable framework, and the prompt provides the flexible instructions[34].

4. Roadmap and Future Work

The development of the Generative Agentic Network will proceed in four strategic phases, aligning with the R&D roadmap proposed in the foundational research[35].

- **Phase 1: Feasibility and Metric Development (Months 1-6):** Empirically quantify LLM non-determinism and develop robust semantic equivalence metrics[36].
- **Phase 2: Core Infrastructure & Proof-of-Concept (Months 7-18):** Develop the core smart contracts and basic agents, and integrate with existing DeAI infrastructure like IPFS for storage and DePIN networks for compute[37].
- **Phase 3: Adjudication Network & Security Hardening (Months 15-24):** Implement the Decentralised Adjudication Network and begin integrating a Tier 2 zkML verification option through partnerships with specialised providers[38].

- **Phase 4: DAO Launch and Ecosystem Growth:** Deploy the DAO for community governance and launch an ecosystem fund to incentivise the development of specialised worker and adjudicator agents.

5. Conclusion

The Generative Agentic Network protocol represents a significant leap forward in the field of autonomous agent communication. By creating a standardised way to package, trade, and verifiably execute complex, intent-driven tasks, the GAN provides a foundational layer for the emerging on-chain AI agent economy[39]. Through its novel solutions to the challenges of non-determinism, verifiable computation, and contractual ambiguity, the protocol offers a viable and robust blueprint for a future where AI agents collaborate and transact with unprecedented efficiency and trust.

---

6. Works Cited

The research and analysis presented in this white paper are based on the findings detailed in the "Generative Document Representation: A Research Landscape Analysis for a Semantic Agent Protocol" report and its cited works[40].