PUBLIC COMMENT

Re: Concept Paper, Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization (February 2026)
Submitted to: AI-Identity@nist.gov
April 1, 2026
Vega Commons Project, Inc.

---

1. INTRODUCTION AND PRIOR ENGAGEMENT

Vega Commons Project, Inc. (VCP) is a New York not-for-profit corporation that develops governance standards for AI interaction records. VCP submitted a public comment to the CAISI Request for Information on Security Considerations for Artificial Intelligence Agents (Docket NIST-2025-0035), posted as Comment ID NIST-2025-0035-0015 on March 15, 2026. That submission identified interaction record retention and custody as a governance layer distinct from agent security, authorization, and verifiable execution. This comment extends that analysis to the identity and authorization domain addressed by the NCCoE concept paper.

VCP publishes the Automated Record Custody Standard (ARCS), a governance standard for the lifecycle management of records generated by automated systems, including AI systems.

2. IDENTITY, AUTHORIZATION, AND LOGGING WITHOUT RECORD LIFECYCLE GOVERNANCE

The concept paper identifies areas of interest relating to identification, authorization, access delegation, logging and transparency, and tracking data flows of an AI system. Those areas address identification, authorization, delegation, logging, and provenance for agent activity. They do not address the lifecycle management of the records generated by those processes, including retention, custodial responsibility, access and disclosure conditions, and controls governing creation, preservation, and disposition.

Each of the standards identified in the concept paper creates records. OAuth token issuance, OIDC authentication, SPIFFE attestation, SCIM provisioning, and NGAC policy updates all generate artifacts that persist, reside on organizational or provider infrastructure, and may be subject to access requests, regulatory inquiry, or legal process. The concept paper's auditing and non-repudiation framework depends on these records surviving long enough to serve an evidentiary purpose, but it does not specify who retains them, for how long, or under what

lifecycle controls, and it does not address their treatment when access, disclosure, preservation, or production obligations arise.

The concept paper's "Tracking Data Flows of an AI System" area of interest comes closest to the record-governance issue. It proposes tracking and maintaining provenance of user prompts and data input sources to support risk determinations and policy decisions. That framing assumes that the records being tracked will persist long enough to serve those purposes. It does not specify lifecycle governance for the provenance records themselves, including retention, custodial responsibility, deletion and preservation controls, or custody management when provenance data spans multiple vendors and infrastructure layers.

3. RESPONSES TO SPECIFIC QUESTIONS

3.1 Auditing and Non-Repudiation

The concept paper asks how agents can log their actions and intent in a tamper-proof and verifiable manner, and how non-repudiation can be established for agent actions and their relationship to human authorization.

Both questions assume logging and retention of agent actions and associated authorization records. Tamper-proof logging requires durable storage. Non-repudiation requires that records linking agent actions to human authorization persist long enough to serve an evidentiary purpose.

The concept paper does not address three related implementation variables.

Retention duration: The concept paper does not identify retention duration as a variable for the demonstration project, although the required period may vary by sector, jurisdiction, and use case.

Custodial responsibility: In a multi-vendor agentic architecture, including architectures in which an agent uses MCP to access tools hosted by different providers, the audit trail may span multiple custodians rather than a single authoritative holder.

Lifecycle governance: The concept paper does not specify controls governing retention, access, modification, disposition, deletion verification, or preservation-state coordination across custodians.

A demonstration project that implements logging without addressing these issues would produce an audit trail without a specified lifecycle, without a defined custodial model across organizational boundaries, and without a governance framework for retention or disposition.

3.2 Authorization and Delegation

The concept paper asks how delegation of authority operates in "on behalf of" scenarios and how agent identity is bound to human identity for human-in-the-loop authorizations.

Delegation generates records. When authority is delegated to an agent through OAuth token issuance, OIDC authentication, or SCIM provisioning, the delegation event creates artifacts including tokens, certificates, consent records, and provisioning logs. These artifacts document the delegating party, the authority delegated, the receiving agent, the time of delegation, and any applicable conditions.

In litigation, regulatory investigation, or incident response, delegation records may serve as the evidentiary link between human authorization and agent action. The authorization framework described in the concept paper depends on the existence of such records but does not specify lifecycle governance for them.

The use cases described in the concept paper rely on this same record class. Use Case 1, addressing enterprise AI agents for workforce efficiency, describes agents with delegated access to multiple data sources. Use Case 3, addressing enterprise AI agents for software development and deployment, describes automated pipelines using AI agents. In both cases, delegation records link agent action to human authorization. Without lifecycle governance, those records may be retained under the default policies of the platforms, infrastructure providers, or integrated services involved, without a coordinated framework governing their retention, preservation, or disposition.

3.3 Prompt Injection Prevention and Mitigation

The concept paper asks about controls to prevent and mitigate prompt injection.

Detection, forensic review, and accountability in prompt-injection incidents depend on retained interaction records. CAISI's discussion of evaluation cheating describes the use of transcript analysis tools to review retained evaluation transcripts at scale. The tool operates on formatted evaluation transcripts that include system and user messages, model messages, tool calls, tool-call responses, and task metadata, including scoring information and intended-solution context. Its use was possible because those transcripts were retained and available for systematic review.

The same dependency applies in the prompt-injection context. Detecting that prompt injection occurred, determining what effect it had on model behavior, and reconstructing resulting tool use

and system responses depend on access to the same general record classes generated during operation.

The concept paper does not address who retains those records, for how long, under what access controls, and under what lifecycle governance framework.

3.4 General Questions

The concept paper asks about standards supporting identity and access management of agents.

The Linux Foundation's Agentic AI Foundation (AAIF) governs MCP, goose, and AGENTS.md. Public Linux Foundation materials describe AAIF as the foundation for open agentic AI infrastructure and identify Amazon Web Services, Anthropic, Block, Bloomberg, Cloudflare, Google, Microsoft, and OpenAI among its members. Its published scope is directed to open standards, protocols, interoperability, and shared infrastructure for agentic systems.

Those published materials do not identify record lifecycle management, retention classification, custody surface mapping, or content-telemetry separation as part of that standards scope. This is consistent with the point made in the submitter's prior NIST public comment (Comment ID NIST-2025-0035-0015): the absence is not limited to a single docket, but appears in the published scope of the ecosystem infrastructure currently being advanced.

3.5 MCP as Identity Artifact Generator

The concept paper identifies Model Context Protocol (MCP) as one of the standards the demonstration project will implement. MCP generates identity and authorization artifacts in operational deployment and through specification revision.

The MCP authorization specification has undergone three major revisions. The current specification classifies MCP servers as OAuth 2.1 resource servers, requires PKCE for all clients, introduces Protected Resource Metadata (RFC 9728) for authorization server discovery, and advances Client ID Metadata Documents as the preferred client registration method. Each of these specification features creates record classes: Dynamic Client Registration artifacts, Client ID Metadata Documents, Protected Resource Metadata documents, token issuance and refresh logs, consent records, and capability negotiation traces. The MCP specification does not specify lifecycle governance for these record classes.

On March 25, 2026, the UK AI Security Institute published an empirical study of MCP tool deployment, classifying 177,436 public MCP tools created between November 2024 and February 2026, in collaboration with the Bank of England. The study reported that MCP tool

availability grew 36-fold in one year and server downloads grew by two orders of magnitude. The paper addresses governance at the tool registry layer. It does not address what records those tools create, how long those records are retained, or who holds them.

Azure, AWS, and Google Cloud now offer official MCP server support with documented product architectures. Those integrations address protocol support, authentication, runtime integration, and deployment. They do not specify lifecycle governance for the records that MCP connections generate, including where copies persist, whether downstream systems preserve their own traces, whether deletion is verifiable across the integration chain, and whether the resulting record classes are visible to legal, audit, and records functions within the adopting organization.

4. RECOMMENDATION

This comment recommends that the NCCoE demonstration project include an additional area of interest: record lifecycle governance for identity and authorization artifacts. The demonstration project should explore how retention duration, custodial responsibility, lifecycle controls, and disposition governance apply to the records generated by the identity, authentication, and authorization systems the project implements. These records include OAuth tokens and refresh tokens, OIDC identity tokens and consent records, SPIFFE SVIDs and attestation logs, SCIM provisioning and deprovisioning events, NGAC policy graph updates, MCP session metadata and protocol artifacts, agent activity logs, and delegation artifacts.

Record lifecycle governance is particularly relevant in the enterprise settings identified in the concept paper, where retention obligations, legal hold requirements, audit expectations, and cross-vendor custody coordination are more likely to apply.

Without that scope, the demonstration project would produce a reference architecture that creates records it does not govern. The identity and authorization framework would generate artifacts with no specified retention period, no designated custodian across organizational boundaries, no lifecycle controls for modification or disposition, and no framework for managing the response burden created when ordinary agent operation generates records that later have to be scoped, reviewed, reconstructed, or produced.

The submitter’s prior public comment (Comment ID NIST-2025-0035-0015) and the ARCS governance framework address the record lifecycle domain directly. The submitter is prepared to participate as a collaborator if the NCCoE determines that a demonstration project should proceed.

ABOUT SUBMITTER

Vega Commons Project, Inc. is a New York not-for-profit corporation that develops governance standards for AI interaction records. VCP publishes the Automated Record Custody Standard (ARCS). VCP is not affiliated with any AI platform vendor and welcomes participation in NCCoE demonstration projects addressing record lifecycle governance.

Contact: info@vegacommons.org
Submitted April 1, 2026