Actor Governance: What Does the Audit Record Actually Establish?

CoreBond Research Series — CBRS-001
August 2026
Todd Kirton
DOI: 10.5281/zenodo.22027092

Then the accountability question changes. An investigation asks which person initiated the operation. A compliance review asks who approved the privilege. A service owner asks whether an application acted for itself or under delegated user authority. A forensic review asks which workflow triggered the request. The original record can remain technically correct while being insufficient for the new proposition.

The practical failure is semantic over-attribution. A later reviewer can take evidence that establishes an executing principal and silently treat it as evidence of an initiator, delegator, approver, responsible human, or intentional participant. Those roles can coincide, but modern systems give no reason to assume that they always do.

This paper uses actor governance as descriptive shorthand for one bounded problem: identifying which actor-related proposition a governance or accountability decision needs and determining whether preserved provenance supports that proposition. It is not proposed as a new standards category or security layer. The question is narrower: when governance records that an identity performed an action, what has actually been established about the actor relationships behind that action?

The Easy Thesis Fails

A tempting answer is that governance systems record credentials or accounts while silently assuming those records identify a human. Current identity architecture makes that claim difficult to defend.

Modern platforms intentionally represent different kinds of principals. Databricks documents service principals as specialized identities for automation and programmatic access, including jobs, scripts, applications, and CI/CD platforms (Databricks, 2026). These are not examples of software being mistaken for people; they are intentional non-human identities.

The first correction is therefore straightforward: actor cannot silently mean natural person. Governance records are most useful when they distinguish the principal that exercised authority from deeper actor relationships that require separate evidence. Human-specific attribution becomes necessary only when the downstream accountability proposition itself is human-specific. This paper concerns that identity-and-accountability dimension of governance; it does not attempt to describe every function normally included under data governance.

Authentication Already Establishes More Than Possession

The easy thesis also understates authentication. NIST's current digital-identity guidance defines authentication around a claimant demonstrating control of authenticator(s) bound to a subscriber account (Temoshok et al., 2025). Authentication is therefore not inherently an anonymous credential check detached from an account relationship. Identity proofing and enrollment, when required by the applicable assurance model, are separate parts of the broader digital-identity process rather than propositions established by the authentication event itself.

WebAuthn is another strong counterexample to the claim that a hardware authenticator proves possession and nothing more. WebAuthn ceremonies include user-presence and user-verification semantics. Depending on authenticator policy, user verification can require a local PIN, biometric recognition, or another verification method before the authenticator produces the assertion expected by the relying party (World Wide Web Consortium [W3C], 2026b).

The propositions must remain separate. A WebAuthn user-verification result establishes that the authenticator satisfied its user-verification requirement for that ceremony. Within a relying party's account model, the resulting assertion can contribute strong evidence that the registered credential was exercised with the required user participation. It does not, by itself, establish an arbitrary civil-identity claim, prove the user's mental intent, or establish that the same person controls every downstream action after the ceremony.

NIST's separate treatment of session monitoring and reauthentication reinforces the boundary between an authentication event and later session risk (Temoshok et al., 2025). Existing authentication is serious actor-related evidence. The mistake is not trusting it; the mistake is attaching propositions to it that the mechanism did not establish.

The Actor Is Not Always a Human

Automation makes a human-at-every-action model fail more decisively. Service principals, managed identities, applications, scheduled jobs, CI/CD systems, and machine-to-machine services can exercise legitimate authority without a person being present when each operation occurs.

A human may have created the workload, approved its deployment, granted its permissions, or triggered a workflow earlier. None of those facts turns the workload's later execution into a contemporaneous human action. For some accountability questions, the workload identity is the correct answer.

For others, it is not enough. An investigation may need the engineer who deployed the job, the application version that triggered the request, the administrator who granted the permission, or the user whose delegated authority was exercised. These are not competing candidates for one universal 'real actor.' They are different relationships that answer different questions.

Delegation Can Preserve More Than the Immediate Principal

Delegated authorization provides perhaps the strongest evidence against the idea that modern systems can see only the immediate credential. OAuth 2.0 Token Exchange explicitly distinguishes

impersonation from delegation. In delegation, one principal acts as an agent for another rather than simply becoming that other principal (Jones et al., 2020).

RFC 8693's actor semantics can identify a current actor separately from the subject, and nested actor claims can represent prior actors as a delegation history trail (Jones et al., 2020). This is not a minor exception to the argument; it is a substantial standards-based mechanism for carrying actor and delegation provenance.

Microsoft's On-Behalf-Of flow provides a concrete implementation example. A middle-tier API can receive a user's delegated token and obtain a downstream token that continues to represent the user's delegated identity and permissions instead of replacing the user context with an unrelated application-only identity (Microsoft, 2026a).

These mechanisms solve real portions of the provenance problem. They do not establish that every deployment propagates the same context, that upstream approvals or triggering events are represented, or that every later accountability proposition is encoded in the token. App-only authorization also intentionally operates without a signed-in user. Their importance is precisely that they force the paper away from claiming absence and toward asking what a particular architecture actually preserved.

A Provenance Trail Is Not the Same as Current Authority

Preserving history creates another distinction. Information about who previously participated in a delegation chain is not necessarily information that should control the current authorization decision.

RFC 8693 makes this boundary unusually clear. When nested actor claims represent prior actors, the token's top-level claims and current actor are relevant to access-control policy, while earlier actors in the history are informational (Jones et al., 2020). A prior actor can matter to later audit or accountability without retaining current authorization significance.

The reverse is also possible. A principal can be properly authorized at the moment of execution while the resulting audit trail remains too limited for a later accountability question. Authorization sufficiency asks whether the principal may act now. Accountability sufficiency asks whether the available record later establishes the relationship a reviewer is trying to attribute. The questions overlap, but they are not interchangeable.

Consequential Actions Can Receive Stronger Ceremony

Existing systems also challenge the claim that authentication must be a one-time login followed by undifferentiated reliance on every later action. Privileged-access management can add fresh ceremony when authority rises. Microsoft Entra Privileged Identity Management, for example, can be configured for just-in-time or time-bounded privileged access and can require controls such as multifactor authentication, justification, and approval during activation (Microsoft, 2026b).

A different domain shows tighter binding between user confirmation and a particular action. W3C Secure Payment Confirmation presents transaction information to a user and incorporates that confirmation into a WebAuthn-based authentication ceremony that produces a signed result (W3C,

2026a). It is a payment-specific counterexample, not evidence that the same mechanism is appropriate or generally available for data-governance actions.

Together, these mechanisms narrow the problem again. Mature systems can demand fresher, more specific, or more explicit actor evidence when consequence warrants it. The remaining question is whether the evidence preserved for a particular operation supports the particular accountability claim later made about it.

Where the Accountability Gap Actually Appears

The surviving gap appears when a later attribution reaches beyond the actor relationship that the preserved provenance actually establishes.

A database audit record may correctly identify a service principal. If the later question is which principal executed the query, the record may be sufficient. If the question becomes which user delegated authority to the application, which administrator approved its privilege, which workflow triggered the operation, or which human explicitly confirmed a consequential action, the same record may no longer answer the question.

That does not make the original record wrong. It means the accountability proposition changed.

Modern governance systems can accurately record human and non-human principals, and mature identity systems can preserve user-verification and delegation context. The accountability gap appears when a downstream decision attributes an initiator, delegator, approver, or human actor beyond the relationship the available provenance actually establishes.

The accountable subject is therefore consequence-dependent. It may be the executing principal, delegated subject, current actor, initiator, approver, responsible owner, or a present human participant. These are different provenance dimensions rather than positions on one universal ladder. The architecture needs the additional provenance required by the proposition that matters, not maximum identity history for its own sake.

Logging Preserves Evidence; It Does Not Create Missing Semantics

Logs remain essential to accountability. NIST log-management guidance describes logs as records of events used for security operations, incident investigation, and compliance activities (Kent & Souppaya, 2006). Identity platforms and applications can record principal identifiers, application context, role activations, approvals, delegated users, transaction details, and other evidence useful to later review.

The downstream record alone, however, cannot establish identity or delegation context that it never received or generated. Independent evidence may later permit forensic correlation: scheduler records, upstream application logs, network telemetry, deployment history, or other sources can sometimes reconstruct part of the chain. That reconstruction is a separate evidentiary process with its own assumptions and uncertainty.

Consider two service calls. In one, an on-behalf-of flow carries delegated user context into the downstream service. In another, an application uses its own app-only identity. Both downstream

records may be correct. Assuming the delegated context was correctly issued and preserved, only the first contains protocol context directly supporting attribution to a delegated user. Later investigators may correlate the second with other records, but the downstream entry itself does not prove a user relationship that was absent from the call.

The practical boundary is narrower than 'logs cannot reconstruct the past.' Logs preserve assertions and events available to them; deeper reconstruction requires additional evidence.

Provenance Is Evidence, Not Omniscience

Even a rich provenance record inherits the trust properties of the systems that produced it. Authentication can be attacked, tokens can be stolen, applications can misstate context, and logging or storage systems can be compromised. A protected or signed record can provide strong integrity for what a system asserted without proving that every underlying real-world assumption was true.

This is especially important when actor language drifts toward intent or responsibility. A verified user may still be deceived. Malware may drive an authenticated session. An application may trigger an operation under authority legitimately granted by an administrator. Technical provenance can establish bounded relationships such as authenticated, delegated, approved, triggered, or executed; those relationships do not automatically prove mental intent, organizational responsibility, or legal causation.

Actor-related provenance should therefore be evaluated as evidence with defined semantics and trust dependencies, not as an infallible chain from log entry to human responsibility.

More Provenance Is Not Automatically Better

The obvious response is to carry more identity and event context everywhere. That creates its own problems.

Longer delegation histories, cross-service user identifiers, device context, approval records, transaction details, and persistent correlation identifiers can improve investigation and accountability. They can also increase privacy exposure, retention burden, cross-domain correlation, disclosure risk, implementation complexity, and the damage caused by compromise of the audit system.

Existing standards already show sensitivity to that tradeoff. WebAuthn credentials are scoped to relying parties, and local biometric information used for user verification is not simply disclosed as raw biometric identity data to the relying party (W3C, 2026b). The fact that more information could be collected does not make indefinite collection a sound governance principle.

Automation creates an operational version of the same problem. Requiring fresh human interaction for every machine action would destroy legitimate unattended workflows. The target is therefore proposition-specific provenance: enough trustworthy context for the accountability consequence being governed, without collecting every possible identity relationship.

Actor Governance Is a Bounded Question

Once the easy thesis is removed, actor governance becomes less dramatic and more useful. In this paper, the term means asking which actor or authority relationship a later decision actually needs to establish and whether the available provenance supports that claim.

Sometimes the executing principal is enough. Sometimes the system must preserve a delegated user. Sometimes a privileged action needs an approval record and a fresh authentication event. Sometimes a transaction needs explicit confirmation. Sometimes the accountable subject is not a person at all but a workload, application owner, or organizational authority.

The boundary matters because accurate technical records can become misleading when later readers silently add semantics that were never established. A service-principal record should not become a claim about a human operator without evidence connecting the two. A WebAuthn user-verification result should not become proof of intent. A delegation history should not become current authorization merely because it is available.

The system should first identify the accountability proposition. Only then can it determine which actor, delegation, initiation, approval, or execution evidence must be preserved.

Conclusion - Attribute Only What the Provenance Supports

The original accountability concern survives, but not in its easiest form. Modern governance does not simply confuse credentials with humans. Identity systems distinguish human and non-human principals. Authentication can include user verification. Delegation can preserve subject and actor context. Privileged systems can require fresh activation, and transaction-specific confirmation can bind stronger evidence to consequential actions.

Those capabilities make the remaining boundary clearer. An audit record may be entirely correct about the principal that acted and still be insufficient for a later question about who initiated, delegated, approved, triggered, or explicitly confirmed the action. The mistake begins when a downstream decision treats evidence for one relationship as proof of another.

> **The conclusion of this paper is a synthesis rather than a quoted requirement from any one standard: governance should attribute only what its provenance can support.**

When accountability requires more than the immediate principal, the architecture should preserve the additional actor, delegation, initiation, approval, or execution evidence needed for that specific claim. It should not infer the missing relationship merely because a deeper answer would be convenient.

References

Databricks. (2026, May 4). Service principals. https://docs.databricks.com/aws/en/admin/users-groups/service-principals

Jones, M., Nadalin, A., Campbell, B., Bradley, J., & Mortimore, C. (2020). OAuth 2.0 token exchange (RFC 8693). Internet Engineering Task Force. https://doi.org/10.17487/RFC8693

Kent, K., & Souppaya, M. (2006). Guide to computer security log management (NIST Special Publication 800-92). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-92

Microsoft. (2026a). Authentication flow support in the Microsoft Authentication Library (MSAL): On-behalf-of (OBO). https://learn.microsoft.com/en-us/entra/msal/msal-authentication-flows

Microsoft. (2026b). Activate Microsoft Entra roles in Privileged Identity Management. https://learn.microsoft.com/en-us/entra/id-governance/privileged-identity-management/pim-how-to-activate-role

Temoshok, D., Fenton, J. L., Choong, Y.-Y., Lefkovitz, N. B., Regenscheid, A. R., Galluzzo, R., & Richer, J. P. (2025). Digital identity guidelines: Authentication and authenticator management (NIST Special Publication 800-63B-4). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-63B-4

World Wide Web Consortium. (2026a, June 4). Secure Payment Confirmation. https://www.w3.org/TR/2026/CRD-secure-payment-confirmation-20260604/

World Wide Web Consortium. (2026b). Web Authentication: An API for accessing Public Key Credentials - Level 3. https://www.w3.org/TR/webauthn-3/