Proposed Governance Norms and Minimum Organizational Competencies for Autonomous Systems and Agentic AI

Discussion Draft for the Autonomous Systems Governance Working Group

Purpose

To propose a technology-neutral baseline of governance outcomes and organizational competencies for autonomous and agentic systems, aligned with the human-control, accountability, transparency, and traceability objectives of ABA Resolution 604.

Prepared by John M. Willis

Founder & Chief Systems Architect, AGCP.ai

An initiative of Sustainable Future Tech, Inc.

July 2026

Status: Discussion draft. This document does not represent an official position of the American Bar Association, the ASG-WG, or any ABA section or committee.

Proposed Governance Norms and Minimum Organizational Competencies for Autonomous Systems and Agentic AI

Discussion Draft

Executive Summary

Autonomous and AI-enabled systems can make or materially influence decisions and, when integrated with tools, workflows, and operational systems, can initiate actions that produce legal, financial, operational, safety, privacy, or cybersecurity consequences. The operational-capability premise is addressed in runtime-governance architecture research (Willis, 2026a), while ABA Resolution 604, the OECD AI Recommendation, and the NIST AI Risk Management Framework establish the corresponding concerns of human control, accountability, transparency, traceability, and risk management (American Bar Association, 2023; OECD, 2024; Tabassi, 2023).

This discussion draft proposes ten governance norms and a corresponding set of minimum organizational competencies in the ASG-WG focus areas of data management, risk management, and operations management. The norms state the outcomes and conditions that governed systems and responsible organizations should achieve. The competencies identify the demonstrable organizational capabilities needed to establish, operate, coordinate, assess, and improve governance in those three functional domains.

The framework is intentionally technology-neutral and risk-proportionate. It does not require any one attorney, technologist, risk professional, executive, or other participant to possess every competency, prescribe a single architecture, or resolve questions of liability under particular facts. It instead proposes a common organizational baseline that legal, technical, risk, assurance, data, and operational functions can collectively implement and demonstrate. This approach is consistent with the voluntary, use-case-agnostic orientation of the NIST AI Risk Management Framework and the management-system orientation of ISO/IEC 42001 (International Organization for Standardization, 2023; Tabassi, 2023).

1. Purpose, Scope, and Status

1.1 Purpose

The purpose of this document is to support ASG-WG discussion concerning the minimum governance norms and organizational competencies needed to govern autonomous systems and agentic AI in commercial, professional, public-sector, and other consequential environments.

The framework advances the three objectives expressed in ABA Resolution 604: human authority, oversight, and control; accountability of responsible individuals and organizations for consequences; and transparency and traceability through documentation of key decisions concerning AI systems and their outcomes (American Bar Association, 2023).

1.2 Scope

The proposed norms apply to systems that can recommend, propose, coordinate, authorize, initiate, or execute actions affecting people, organizations, assets, infrastructure, data, services, legal interests, or authoritative records. Examples include AI agents, autonomous workflows, orchestration systems, enterprise copilots with tool access, robotic systems, programmatic decision systems, and multi-agent environments.

Application should be proportionate to the nature of the system, the extent of delegated authority, the reversibility of actions, the severity and scale of foreseeable harm, the degree of human reliance, and the legal or operational significance of the affected interests (European Parliament and Council of the European Union, 2024; International Organization for Standardization, 2018; Tabassi, 2023).

1.3 Status and Intended Effect

This is a discussion draft. It is not an official ABA or ASG-WG position, a statement of law, a professional standard, or a presumption of negligence or liability. The proposed norms may inform future guidance, professional education, governance frameworks, contracting practices, assurance criteria, and sector-specific standards.

2. Framework and Key Distinctions

2.1 Norms, Competencies, Controls, and Evidence

For purposes of this document:

- A norm is an expected governance outcome, behavior, or condition.
- A competency is a demonstrable organizational capability to achieve, sustain, assess, or improve a governance outcome within a defined functional domain.
- A control is a legal, organizational, procedural, or technical mechanism used to implement or support a competency and its associated norm.

Evidence is information demonstrating whether accountability is assigned, the competency and its controls are implemented and coordinated, and the intended governance outcome is being achieved in practice. These distinctions matter because an organization may endorse a norm without possessing the competency, controls, or evidence necessary to realize it. Conversely, a sophisticated control does not establish competent governance unless it contributes to the intended accountable outcome.

2.2 Minimum Organizational Competency

For purposes of this framework, minimum governance competency means a demonstrable organizational capability to govern autonomous systems responsibly, lawfully, and effectively within a defined functional domain. A competency ordinarily requires assigned accountability, appropriate multidisciplinary expertise, documented processes, operational controls, coordination across functions, and evidence that the capability operates as intended. Responsibilities may be allocated among legal counsel, compliance, risk management, data governance, cybersecurity, architecture, engineering, audit, assurance, product management, operations, executive leadership, and affected-domain specialists. No single participant is

expected to perform every function, but the organization should be able to demonstrate that each applicable competency is owned, implemented, coordinated, and effective.

2.3 Proposal, Authorization, Enforcement, and Execution

A central premise of this framework is that a model output, recommendation, plan, workflow step, tool request, or proposed state transition is not itself execution. Governance should distinguish proposal formation, evaluation, authorization, enforcement, operational commitment, and the resulting consequence. These stages may occur close together, but they serve different legal and technical functions. Established computer-security and distributed-systems literature supports the underlying concerns of complete mediation, event ordering, and time-of-check/time-of-use risk (Bishop & Dilger, 1996; Lamport, 1978; Saltzer & Schroeder, 1975). Runtime-governance research applies those foundations to the distinction among proposal, authorization, binding, and execution (Willis, 2026a, 2026b).

3. Proposed Governance Norms

The following norms are proposed as a minimum, technology-neutral baseline. Each should be implemented in a manner proportionate to the system, use case, authority, affected interests, and foreseeable consequences.

Norm 1. Human Authority and Accountable Legal Persons

Autonomous systems should remain subject to authority established by human beings or legally responsible organizations. Delegation to a system should not transfer or extinguish the responsibility of the persons or legal entities that design, deploy, direct, operate, integrate, or use it.

The organization should identify the responsible legal person or organizational authority for approving the use case, defining permitted authority, supervising operation, responding to harmful outcomes, and modifying, suspending, or terminating the system's authority (American Bar Association, 2023; European Parliament and Council of the European Union, 2024; OECD, 2024; Tabassi, 2023).

Norm 2. Explicit and Bounded Delegation

Authority delegated to an autonomous system should be expressly defined, attributable, limited, and reviewable. The delegation should identify the actions the system may perform or propose; the persons, systems, data, assets, and domains it may affect; applicable financial, temporal, geographic, and operational limits; and conditions requiring additional approval, escalation, suspension, or termination.

Authority should not expand merely because the system invokes another agent, service, model, workflow, tool, or provider. Recursive and cross-domain delegation should remain bounded by the authority originally granted unless additional authority is expressly established. This norm draws on longstanding principles of least privilege, role-bounded authorization, and the rejection of implicit trust (Rose et al., 2020; Sandhu et al., 1996; Saltzer & Schroeder, 1975). A sector-specific legal analogue appears in the HIPAA Privacy Rule's minimum-necessary standard, which generally requires covered entities to limit certain uses, disclosures, and requests for protected health information to what is needed for the intended

purpose, subject to specified exceptions (U.S. Department of Health and Human Services, 2003). The analogy is purpose-bounded access and delegation; it is not a claim that HIPAA applies to autonomous systems generally.

Norm 3. Separation of Proposal from Consequential Action

Governance should distinguish what a system recommends or proposes from what is authorized, enforced, committed, and executed. A proposal may be validly formed, procedurally approved, or initially admissible without remaining legitimate when operational consequence becomes possible.

Organizations should identify the point at which a proposed action becomes capable of changing authoritative state, transferring value, affecting legal rights, controlling equipment, disclosing information, or otherwise producing consequential effects. Distributed-systems literature establishes that ordering affects the meaning and correctness of operations (Lamport, 1978). The application of that distinction to proposal, binding, and operational commitment is developed in runtime-governance work (Willis, 2026a).

Norm 4. Time-of-Action Authority and Current-Condition Validation

A consequential action should proceed only when relevant authority and governance conditions remain valid at the time the action is about to occur. Earlier approval should not be treated as permanently sufficient when authority, evidence, policy, targets, dependencies, operational state, or environmental conditions may have changed.

Where circumstances can materially change between evaluation and execution, the system should revalidate the conditions necessary for the action to remain legitimate. Material mismatch, expiration, revocation, substitution, or state drift should result in refusal, re-evaluation, or escalation rather than automatic continuation. The general concern is reflected in time-of-check/time-of-use analysis and zero-trust guidance, while its application to commit-bound governance is developed in runtime-governance work (Bishop & Dilger, 1996; Rose et al., 2020; Willis, 2026a).

Norm 5. Risk-Proportionate and Effective Human Oversight

Human oversight should be designed according to the nature, scale, reversibility, uncertainty, and potential consequences of the autonomous activity. Relevant mechanisms may include prior authorization, bounded delegation, separation of duties, quorum approval, exception review, escalation, supervisory review, pause or termination authority, and post-event examination.

Human-in-the-loop approval is one oversight pattern, not the sole measure of human control. For high-risk AI systems within its scope, Article 14 of the EU AI Act requires oversight measures proportionate to risk, autonomy, and context and enables assigned persons to understand system capabilities and limitations, monitor operation, interpret outputs, override or reverse outputs, intervene, or safely stop the system (European Parliament and Council of the European Union, 2024, art. 14). Human-automation research likewise supports treating the type and level of human involvement as a design variable rather than a binary condition (Parasuraman et al., 2000).

Norm 6. Effective Enforcement, Non-Bypassability, and Safe Refusal

Governance requirements should be capable of affecting whether a consequential action occurs. Warning, scoring, monitoring, explanation, logging, or after-the-fact detection alone should not be treated as execution control.

Where mandatory conditions are not satisfied, the system should prevent, refuse, defer, constrain, or escalate the action as appropriate. Materially consequential actions should not have an alternative execution path that bypasses required governance. Refusal should fail safely, preserve evidence, and avoid converting an invalid proposal into operational consequence. Complete mediation and fail-safe defaults are longstanding security-design principles (Saltzer & Schroeder, 1975). CISA secure-by-design guidance emphasizes manufacturer responsibility for security outcomes and secure defaults (Cybersecurity and Infrastructure Security Agency, 2023). NIST's control catalog, assessment procedures, and Cybersecurity Framework support implementation and assessment of enforceable cybersecurity controls without prescribing a single technical implementation (Joint Task Force, 2020, 2022; National Institute of Standards and Technology, 2024).

Norm 7. Data, Context, Evidence, and Provenance Integrity

Governance decisions should rely on information sufficiently reliable for the intended decision. Organizations should identify authoritative sources, provenance, quality limitations, uncertainty, temporal validity, applicable policy versions, and the identity and authority of material contributors.

Model-generated assertions, cached information, runtime observations, third-party data, and unverified contextual claims should not automatically be treated as authoritative. Missing, stale, conflicting, incomplete, or unverifiable evidence should trigger proportionate refusal, qualification, re-evaluation, or escalation. The EU AI Act and NIST AI RMF support risk-based attention to data quality, relevance, limitations, and trustworthiness in consequential AI uses (European Parliament and Council of the European Union, 2024; Tabassi, 2023). PROV-DM provides a domain-agnostic model for representing the entities, activities, agents, derivations, and timing that constitute provenance (Moreau & Missier, 2013).

Norm 8. Transparency, Traceability, and Evidentiary Sufficiency

Organizations should preserve information sufficient to understand, challenge, assess, and reconstruct consequential autonomous actions. The record should be proportionate to risk and should identify the initiating actor and purpose, proposed action and target, applicable authority and policy, material data and evidence, approvals or interventions, the governance outcome, the action actually executed, and the resulting operational effect.

Traceability should support more than a chronology of events. It should permit an authorized reviewer to determine why an action was permitted, refused, deferred, or escalated under the conditions that existed at the relevant time. Accountability and transparency authorities support documentation sufficient to understand and challenge consequential AI decisions and outcomes (American Bar Association, 2023; European Parliament and Council of the European Union, 2024; OECD, 2024). PROV-DM supports representing the origin, derivation, activities, responsible agents, and timing associated with material information (Moreau & Missier, 2013). The NIST AI RMF supports risk-based consideration of information quality and

trustworthiness, subject to lawful confidentiality, privacy, security, privilege, and intellectual-property protections (Tabassi, 2023).

Norm 9. Governance Continuity Across Systems and Organizational Boundaries

Governance obligations should remain effective when actions cross agents, services, vendors, technical platforms, organizational units, contractual boundaries, or jurisdictions. A receiving system should not assume that an instruction remains legitimate solely because it originated from a trusted source.

Cross-system handoffs should preserve, convey, or independently re-establish relevant identity, purpose, authority, limitations, policy obligations, evidence, provenance, and accountability. Zero-trust principles reject implicit trust across system boundaries, while interoperable provenance models support attribution and verification across heterogeneous environments (Moreau & Missier, 2013; Rose et al., 2020). Contracting with a provider should not eliminate the deploying organization's responsibility to understand and govern consequential uses of the provider's system. This proposed allocation principle is consistent with the EU AI Act's differentiated provider and deployer duties and the OECD accountability principle, while the precise allocation of legal responsibility remains dependent on applicable law and contract (European Parliament and Council of the European Union, 2024; OECD, 2024).

Norm 10. Lifecycle Review, Incident Response, Redress, and Governance Evolution

Governance should continue after deployment. Organizations should review whether delegations remain appropriate, controls remain effective, use cases have expanded, new capabilities create additional risk, exceptions are becoming routine, or operational evidence shows that governance is too permissive, too restrictive, or misaligned with its intended purpose.

Organizations should maintain processes for incident containment, investigation, evidence preservation, correction, recovery, remediation, and legally appropriate challenge or redress. Material changes to capability, authority, intended use, operating context, data, policy, or control architecture should trigger renewed governance review. Continual improvement and post-market monitoring are reflected in ISO/IEC 42001, the EU AI Act, and the NIST AI RMF (European Parliament and Council of the European Union, 2024; International Organization for Standardization, 2023; Tabassi, 2023). Cybersecurity governance and integrated incident response are addressed directly by NIST CSF 2.0 and NIST incident-response guidance (National Institute of Standards and Technology, 2024; Nelson et al., 2025).

4. Proposed Minimum Organizational Competencies

The competencies below describe minimum organizational capabilities rather than individual credentialing or training requirements. An organization may distribute responsibility across multiple functions, but it should be able to demonstrate that each applicable capability has an accountable owner, suitable expertise, documented processes, effective controls, and evidence of operation. Individual professional knowledge and skill remain important inputs to organizational competency, but they are not substitutes for a functioning organizational capability.

4.1 Cross-Domain Foundational Competencies

Across data management, risk management, and operations management, the organization should be able to demonstrate the following foundational capabilities:

- Maintain a shared understanding of the distinctions among reasoning, recommendation, proposal formation, admissibility evaluation, authorization, enforcement, operational commitment, execution, and post-event audit.
- Assign and document the human beings, legal entities, and organizational functions responsible for defining, delegating, supervising, enforcing, modifying, and terminating autonomous authority.
- Maintain an inventory of the actions a system can perform or initiate and the people, systems, assets, data, legal interests, and operational domains those actions may affect.
- Distinguish identity and permission from authority that remains valid under current conditions at the time of action.
- Identify and document where governance decisions are made, where they are enforced, where bypass paths may exist, and where legal or operational consequence becomes possible.
- Determine when human oversight is required and whether the oversight is timely, informed, authorized, competent, independent where necessary, and technically effective.
- Identify and retain the evidence necessary to demonstrate accountable governance, explain or challenge a decision, support review, and reconstruct a consequential action.
- Recognize and manage limitations arising from uncertainty, stale information, incomplete context, conflicting authority, changing state, and cross-system handoffs.
- Distinguish governance objectives from the specific legal, organizational, procedural, and technical controls used to implement them.
- Communicate governance requirements across legal, technical, risk, data, assurance, and operational disciplines using a shared vocabulary and clearly allocated responsibilities.

4.2 Data Management Competencies

These are minimum organizational capabilities, not requirements that any single lawyer, technologist, or data professional possess them all. They may be allocated among legal, privacy, data governance, records management, cybersecurity, architecture, engineering, assurance, and operations personnel. The organization should nevertheless be able to demonstrate that the capabilities are assigned, coordinated, and operating effectively. The EU AI Act and NIST AI RMF provide foundations for data governance, quality, recordkeeping, and risk-based qualification of information used by AI systems; PROV-DM provides the corresponding model for provenance representation (European Parliament and Council of the European Union, 2024; Moreau & Missier, 2013; Tabassi, 2023):

- Maintain an inventory of the data, contextual information, evidence, records, and authoritative sources used to form, evaluate, authorize, or execute autonomous actions.
- Establish and apply criteria for assessing data quality, relevance, completeness, timeliness, provenance, uncertainty, representativeness, privacy, security, and permissible use in relation to the specific decision.
- Distinguish and appropriately label authoritative records, predictions, inferences, generated content, cached information, telemetry, and unverified assertions.

- Preserve sufficient decision context to understand why an action was permitted, refused, deferred, or escalated, including applicable versions of data, policies, models, rules, and evidence.
- Establish retention, access, disclosure, correction, deletion, minimization, confidentiality, and privilege requirements for governance evidence.
- Maintain data and evidence lineage when information moves among agents, platforms, vendors, organizational units, or jurisdictions.
- Define qualification criteria for information used in consequential decisions and determine when uncertainty, inconsistency, or conflict requires qualification, escalation, or refusal.
- Detect and respond to stale, missing, conflicting, corrupted, substituted, or unverifiable data and evidence before consequential action occurs.
- Retrieve and substantiate, or where appropriate independently verify, the data and evidence basis of a prior governance outcome.
- Translate applicable data-governance, records-management, confidentiality, and privacy obligations into operational requirements for autonomous action.

4.3 Risk Management Competencies

These are minimum organizational capabilities, not requirements that one lawyer, risk professional, or model specialist perform every task. They may be distributed among legal, enterprise risk, model risk, compliance, product, security, audit, assurance, and affected-domain specialists, but the organization should be able to demonstrate coordinated risk governance and accountable decision-making. Article 9 of the EU AI Act provides a direct example of an iterative risk-management requirement for high-risk AI systems (European Parliament and Council of the European Union, 2024, art. 9). The 2026 interagency model-risk guidance provides a useful banking-sector analogue for risk-based governance, validation, effective challenge, ongoing monitoring, clear role allocation, documentation, and third-party oversight. It is supervisory guidance rather than a regulation, does not establish enforceable standards, and expressly excludes generative and agentic AI models from its scope; it is cited here only as an analogous governance practice rather than directly applicable agentic-AI guidance (Board of Governors of the Federal Reserve System et al., 2026). The competency set also aligns with ISO 31000, ISO/IEC 42001, and the NIST AI RMF (International Organization for Standardization, 2018, 2023; Tabassi, 2023):

- Maintain an inventory of consequential autonomous use cases, action classes, execution targets, affected parties, dependencies, and delegated authorities.
- Evaluate foreseeable legal, financial, safety, privacy, cybersecurity, civil-rights, operational, reputational, and systemic harms.
- Assess the scale, velocity, reversibility, uncertainty, persistence, cumulative effect, and cross-domain propagation of potential consequences.
- Classify risk using criteria appropriate to the organization and sector, and define the governance rigor required for each risk tier.
- Map risks and legal obligations to accountable persons, authority limits, governance controls, evidence requirements, oversight mechanisms, and escalation paths.
- Evaluate risks created by recursive delegation, tool use, multi-agent coordination, third-party services, compositional workflows, and cross-system execution.
- Determine when separation of duties, human approval, independent review, quorum, external authorization, enhanced monitoring, or prohibition is required.

- Analyze reasonably foreseeable failure scenarios, including stale approval, revoked authority, data corruption, model error, control bypass, conflicting actions, unauthorized self-modification, evidence loss, and inability to reconstruct events.
- Evaluate residual risk and determine whether deployment, continued operation, expanded authority, or a proposed material change remains acceptable.
- Use incidents, refusals, exceptions, overrides, operational metrics, and assurance findings to reassess risk and improve governance over time.

4.4 Operations Management Competencies

These are minimum organizational capabilities, not requirements that any single attorney, architect, engineer, or operator possess them all. They may be allocated among legal, security, architecture, engineering, platform operations, incident response, audit, assurance, procurement, and vendor-management functions. The organization should nevertheless be able to demonstrate that operational governance is assigned, coordinated, testable, and effective. Complete mediation and secure-by-design engineering provide the underlying enforcement orientation (Cybersecurity and Infrastructure Security Agency, 2023; Saltzer & Schroeder, 1975). NIST guidance provides complementary foundations for cybersecurity governance, control implementation and assessment, and incident response (Joint Task Force, 2020, 2022; National Institute of Standards and Technology, 2024; Nelson et al., 2025):

- Maintain a documented understanding of how consequential autonomous actions move from initiation and proposal through evaluation, authorization, execution, evidence generation, and closure.
- Identify the points at which an autonomous action can produce legal or operational consequence and ensure that required governance review and control occur before those points.
- Ensure that the action ultimately presented for execution corresponds to the action, target, purpose, parameters, limitations, evidence, and authority that were reviewed and authorized.
- Revalidate authority, policy, evidence, dependencies, target conditions, and other material factors when circumstances may have changed before consequence occurs.
- Maintain effective means to refuse, defer, expire, revoke, escalate, suspend, reverse where feasible, or safely stop an autonomous action, including timely human intervention when required.
- Identify, prevent, or otherwise effectively control alternate or privileged pathways that could bypass required governance, including direct tool access, alternate interfaces, maintenance paths, and cross-system workarounds.
- Preserve reliable records of proposals, evaluations, approvals, refusals, escalations, execution attempts, execution outcomes, and resulting state changes.
- Govern material changes to policy, configuration, authority, software, models, data, and infrastructure so that changes do not create unreviewed authority or execution paths.
- Test whether governance controls are appropriately designed and operate effectively under normal, adverse, boundary, replay, failure, and bypass scenarios.
- Monitor indicators of operational and governance performance, including abnormal refusal rates, repeated overrides, approval bottlenecks, evidence failures, execution mismatches, and governance delay.

- Investigate incidents and anomalous behavior, preserve evidence, identify root causes, contain harm, restore safe operation, and document corrective action and lessons learned.
- Coordinate governance responsibilities across internal teams, vendors, service providers, regulators, auditors, affected parties, and other external stakeholders.

5. Organizational Application and Demonstration of Competency

5.1 Role Allocation

Organizations should assign ownership for each norm and organizational competency. The assignment should identify the accountable function, decision authority, escalation path, supporting roles, required independence, and evidence of performance. Responsibility may be distributed, but it should not be ambiguous (International Organization for Standardization, 2023; OECD, 2024; Tabassi, 2023).

5.2 Risk-Proportionate Evidence of Competency

Evidence of organizational competency should be proportionate to risk. Evidence may include action and authority inventories, governance charters, delegation instruments, data and evidence maps, risk assessments, architecture diagrams, control descriptions, operating procedures, test results, incident records, governance receipts, refusal records, audit evidence, training records, and documented review decisions. NIST SP 800-53A provides customizable assessment procedures for determining whether controls are implemented correctly, operating as intended, and producing the desired outcome (Joint Task Force, 2022).

5.3 Assessment Questions

A practical assessment of organizational competency should be able to answer at least the following questions:

- What consequential actions can the system propose or perform?
- Who or what grants authority, and what are the limits of that authority?
- What conditions must be satisfied when consequence becomes possible?
- Where is the decision enforced, and can the enforcement path be bypassed?
- What happens when required conditions are not satisfied?
- What information establishes authoritative state, context, and evidence?
- How is governance preserved across systems, agents, vendors, and organizational boundaries?
- Can the organization reconstruct why a specific action was permitted or refused?
- Who is accountable for consequences and remediation?
- What evidence demonstrates that governance operates as designed in practice?

5.4 Relationship to Existing Law and Standards

These norms and competencies are intended to complement, not replace, applicable law, professional obligations, contractual duties, sector regulation, privacy and data-governance requirements, risk-management frameworks, AI management systems, and established assurance practices (European Parliament and Council of the European Union, 2024; International Organization for Standardization, 2023; OECD, 2024; Tabassi, 2023). They likewise complement cybersecurity controls and secure-by-design guidance (Cybersecurity

and Infrastructure Security Agency, 2023; National Institute of Standards and Technology, 2024). The execution-focused lens used in this document is developed in the runtime-governance synthesis literature (Willis, 2026c).

Several existing regimes provide useful analogues without establishing universal rules for autonomous systems. Article 9 of the EU AI Act uses an iterative lifecycle risk-management model for high-risk AI systems. The 2026 interagency model-risk guidance illustrates tailored governance, effective challenge, ongoing monitoring, documentation, clear accountability, and third-party oversight in banking. The HIPAA minimum-necessary standard illustrates purpose-bounded access and disclosure. These analogues can help legal and governance professionals interpret the proposed norms, but each use case still requires analysis of the law and professional obligations actually applicable to it (Board of Governors of the Federal Reserve System et al., 2026; European Parliament and Council of the European Union, 2024, art. 9; U.S. Department of Health and Human Services, 2003).

6. Intended Uses

The proposed norms and organizational competencies may support:

- ASG-WG discussion, consensus development, and future work products;
- sector-specific legal and governance guidance;
- professional education concerning the legal, technical, risk, data, assurance, and operational functions that contribute to organizational competency;
- organizational policy, delegation, and oversight design;
- procurement, contracting, due diligence, and third-party governance;
- system architecture and operational-control review;
- risk, audit, assurance, and control-assessment criteria;
- incident investigation, accountability, and redress processes; and
- future standards, best practices, or model governance provisions.

7. Conclusion

Responsible governance of autonomous systems requires more than the existence of policy, a human-approval step, or a post-event audit trail. It requires a demonstrable connection among human authority, bounded delegation, qualified information, risk-proportionate oversight, effective enforcement, operational evidence, and accountable consequences.

The proposed norms describe the minimum outcomes that should exist. The proposed organizational competencies describe what responsible organizations should be demonstrably capable of doing to make those outcomes real. Together, they provide a practical starting point for developing legal and governance frameworks that preserve human responsibility while enabling responsible innovation in autonomous and agentic systems.

8. Questions for ASG-WG Consideration

The following questions may help structure review of the proposed baseline:

- Do the proposed norms capture the minimum outcomes needed to give practical effect to human control, accountability, transparency, and traceability?

- Which norms should apply universally, and which should be triggered by defined levels of risk, delegated authority, or consequence?
- Which minimum organizational competencies are essential in each of the three functional domains, and what evidence should demonstrate that the organization possesses and sustains them?
- Where should sector-specific duties, professional obligations, and insurability considerations modify or extend the baseline?
- How should responsibility be allocated among AI providers, vendors, deployers, integrators, operators, and professional users, and which governance duties should remain non-delegable notwithstanding contractual allocation?

References

Independent Authorities and Literature

American Bar Association. (2023). Resolution 604 and report. Adopted by the House of Delegates, February 6, 2023. https://www.americanbar.org/content/dam/aba/directories/policy/midyear-2023/604-midyear-2023.pdf

Bishop, M., & Dilger, M. (1996). *Checking for race conditions in file accesses*. Computing Systems, 9(2), 131-152. https://www.usenix.org/legacy/publications/compsystems/1996/spr_bishop.pdf

Board of Governors of the Federal Reserve System, Federal Deposit Insurance Corporation, & Office of the Comptroller of the Currency. (2026, April 17). Supervisory guidance on model risk management (attachment to Federal Reserve SR Letter 26-2; corresponding OCC Bulletin 2026-13 and FDIC FIL-15-2026). https://www.federalreserve.gov/supervisionreg/srletters/SR2602a1.pdf

Cybersecurity and Infrastructure Security Agency. (2023). Shifting the balance of cybersecurity risk: Principles and approaches for secure by design software. https://www.cisa.gov/sites/default/files/2023-10/SecureByDesign_508c.pdf

European Parliament and Council of the European Union. (2024). Regulation (EU) 2024/1689 of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). Official Journal of the European Union, L 2024/1689. https://data.europa.eu/eli/reg/2024/1689/oj

International Organization for Standardization. (2018). Risk management - Guidelines (ISO 31000:2018). https://www.iso.org/standard/65694.html

International Organization for Standardization. (2023). Information technology - Artificial intelligence - Management system (ISO/IEC 42001:2023). https://www.iso.org/standard/42001.html

Joint Task Force. (2020). *Security and privacy controls for information systems and organizations (NIST Special Publication 800-53, Revision 5)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53r5

Joint Task Force. (2022). *Assessing security and privacy controls in information systems and organizations (NIST Special Publication 800-53A, Revision 5)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53Ar5

Lamport, L. (1978). *Time, clocks, and the ordering of events in a distributed system*. Communications of the ACM, 21(7), 558-565. https://doi.org/10.1145/359545.359563

Moreau, L., & Missier, P. (Eds.). (2013, April 30). PROV-DM: The PROV data model. W3C Recommendation. https://www.w3.org/TR/prov-dm/

National Institute of Standards and Technology. (2024). *The NIST Cybersecurity Framework (CSF) 2.0 (NIST CSWP 29)*. https://doi.org/10.6028/NIST.CSWP.29

Nelson, A., Rekhi, S., Souppaya, M., & Scarfone, K. (2025). *Incident response recommendations and considerations for cybersecurity risk management: A CSF 2.0 community profile (NIST Special Publication 800-61, Revision 3)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r3

OECD. (2024). *Recommendation of the Council on Artificial Intelligence (OECD/LEGAL/0449, revised May 3, 2024)*. https://legalinstruments.oecd.org/en/instruments/OECD-LEGAL-0449

Parasuraman, R., Sheridan, T. B., & Wickens, C. D. (2000). *A model for types and levels of human interaction with automation*. IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans, 30(3), 286-297. https://doi.org/10.1109/3468.844354

Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture (NIST Special Publication 800-207)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-207

Saltzer, J. H., & Schroeder, M. D. (1975). *The protection of information in computer systems*. Proceedings of the IEEE, 63(9), 1278-1308. https://doi.org/10.1109/PROC.1975.9939

Sandhu, R. S., Coyne, E. J., Feinstein, H. L., & Youman, C. E. (1996). *Role-based access control models*. Computer, 29(2), 38-47. https://doi.org/10.1109/2.485845

Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0) (NIST AI 100-1)*. National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

U.S. Department of Health and Human Services, Office for Civil Rights. (2003, April 3). Minimum necessary requirement. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/minimum-necessary-requirement/index.html

Related Runtime Governance Works by the Author

Willis, J. M. (2026a). AGCP operating model and semantic architecture: Proposal schemas, governance context, canonical state, and commit-bound execution semantics for autonomous systems. Sustainable Future Tech, Inc. https://doi.org/10.5281/zenodo.20330062

Willis, J. M. (2026b). *Formal execution semantics and safety invariants for governance control planes in autonomous systems*. Sustainable Future Tech, Inc. https://doi.org/10.5281/zenodo.19241732

Willis, J. M. (2026c). *Runtime execution governance for AI systems: A cross-platform synthesis and architectural framework*. Sustainable Future Tech, Inc. https://doi.org/10.5281/zenodo.19341177

Prepared for discussion by John M. Willis | AGCP.ai | Sustainable Future Tech, Inc.