

AI Governance aka Artificial Intelligence Governance

TO: Lorna Scheel,
Administrator, U.S. Technical Advisory Group to ISO/TC 258, Project Management
Institute

FROM: Basil C. Puglisi, M.P.A.

CC: ISO Team, American National Standards Institute

DATE: September 15, 2026

SUBJECT: AI Governance: The Definition, Its Scope, and the Record Behind It
(ISO/DIS 21520, Clause 3.2)

Purpose. This memo supports my request to file a public technical comment on ISO/DIS 21520, Clause 3.2. It sets out the proposed definition, what that definition covers, how it has developed, and the public record that places it among the first operational definitions of AI Governance.

Current Clause 3.2 Definition

ISO/DIS 21520 defines artificial intelligence governance as an “organizational framework of policies, roles, responsibilities and processes to ensure the accountable, secure, ethical and effective use of artificial intelligence systems.” Unlike the neighboring entries in Clause 3, which adopt definitions from ISO/IEC 22989:2022 or ISO/IEC 2382:2015, Clause 3.2 cites no source vocabulary.

Technical Deficiency

The definition establishes organizational structure but does not establish the decision-authority threshold that distinguishes governance from oversight, management, Responsible AI controls, or human participation. A March 2026 review of ISO, NIST, the OECD, the EU AI Act, and UNESCO found that each addressed governance, yet none defined AI Governance as a standalone term. Standards bodies define the management systems they certify, and risk frameworks define the functions they measure. Intergovernmental bodies define the principles they negotiate, and regulators define the obligations they enforce. That gap let an organization claim AI Governance through policies, teams, or certification while no named person held the authority to stop an output or answered for it.

Proposed Definition

Applied: AI Governance exists when a qualified human holds binding authority at specific checkpoints, with personal accountability for the outputs that pass through.

Full scope: AI Governance is the system of decision authority, accountability structures, and oversight mechanisms through which named humans hold binding power to accept, modify, or reject AI outputs at defined checkpoints, with authority that cannot be delegated to the AI system being governed and accountability that survives audit through moral, professional, civil, and criminal channels.

AI Governance aka Artificial Intelligence Governance

Operational Scope

Element	Meaning
Named human	A person, not a team or a department, whose identity attaches to the decision
Binding authority	The power to accept, modify, or reject an AI output at a defined checkpoint
Accountability	Moral, professional, civil, and criminal channels that reach that person and survive audit
Human arbitration only	AI cannot satisfy, close, or validate a checkpoint for another AI, and no platform count, confidence score, or convergence level substitutes for human arbitration
Qualified	Capability, not credentials alone: the governance competence to direct, challenge, verify, and own the work done with AI. The human acts as the generalist who owns direction, purpose, and context without needing to out-know the specialist machine, and checkpoint practice builds that capability over time.
High-speed systems	The checkpoint moves upstream, where the named human signs the policy that authorizes automated action

Relationship to Responsible AI

Ethical AI supplies the values, and ethics frameworks govern AI behavior. Responsible AI builds those values into systems through controls, monitoring, and validation. It runs as machine validating machine, agents running the pipeline, or a human in the loop, and it is often the right choice when stakes allow process controls to suffice and outputs are reversible. Its ceiling, absent an assigned human authority, is individual decision accountability, because none of those modes, by itself, requires a named person to hold binding authority over an individual output. Automation with guardrails remains automation, and a human in the loop without binding authority is presence without accountability. AI Governance begins when a named human holds binding authority at a defined checkpoint and answers for what passes through. Checkpoint-Based Governance is the mechanism that converts that presence into authority.

Regulatory Convergence

Regulation has moved toward the same threshold on its own terms. New York's Part 161, effective June 1, 2026, governs the use of AI in court papers while requiring the filer to review the paper and independently verify its contents. Regulation (EU) 2026/1744 deferred the EU AI Act's high-risk obligations, including Article 14 human oversight, to December 2, 2027 for Annex III systems, without removing them.

AI Governance aka Artificial Intelligence Governance

Development Record

The applied definition stands as published, while the full scope now uses accept, modify, or reject as the decision outcomes. It also names the professional channel in place of employment and carries the constitutional prohibition that AI cannot approve AI.

- **September 23, 2025.** Checkpoint-Based Governance published mandatory human arbitration at defined decision points, with every decision logged to an identifiable reviewer.
- **November 20, 2025.** The Checkpoint-Based Governance constitution established that human arbitration retains final decision authority and that no AI may approve another AI's decision.
- **November 2025.** *Governing AI: When Capability Exceeds Control* (ISBN 9798349677687) opened with Geoffrey Hinton's warnings, examined each risk domain he named, and answered them with Checkpoint-Based Governance.
- **March 14, 2026.** The formal definition and its accountability channels appeared at basilpuglisi.com, following the institutional review described above.
- **March 18, 2026.** Owen Ambur, co-chair of the StratML Committee that developed StratML Part 1, later published as ISO 17469-1, independently mapped the definition into StratML (<https://stratml.us/docs/AIGDAF.xml>).

Editorial Notes on Clause 3

- Clause 3.1 refers to “AI systems (3.2),” but the AI system definition appears at 3.4.
- Clause 3.4 defines an “engineering system,” while its Note 1 refers to “the engineered system.”

Request

I ask the U.S. TAG to consider revising Clause 3.2 so that the definition of artificial intelligence governance expressly identifies binding human decision authority and attributable human accountability as elements that distinguish governance from organizational process, oversight, and Responsible AI controls. I offer the definition above as proposed language and would welcome the opportunity to support consideration of this issue as a technical expert.

Contact: Phone: +1.631.220.2754 | e-mail: me@basilpuglisi.com | basilpuglisi.com