

UNITY IMPACT PARTNERS

STRATEGIC PERSPECTIVES™ | NO. 1 | OCTOBER 2026

THE TRUST IMPERATIVE

Governing AI and Preparing for the Post-Quantum Transition

A Leadership Perspective on Identity, Authority, Evidence, and Organizational
Stewardship



October 2026

Co-prepared by Unity Impact Partners and Qrypto Federal

Executive Abstract

Artificial intelligence is transforming how organizations access knowledge, make decisions, and translate intelligence into action. At the same time, advances in quantum computing are driving the transition to new cryptographic standards for protecting information, securing communications, and establishing digital trust. Together, these developments present a defining opportunity to advance innovation, reinforce accountability, and build organizational resilience.

As AI becomes increasingly integrated into organizational operations, including agentic workflows, leaders must establish confidence in the information informing decisions, the identities and authority of those acting, and the integrity of the systems connecting them. This requires an integrated approach to governance, security, and organizational readiness that aligns technological capabilities with human judgment, clearly defined responsibilities, and measurable outcomes.

The Trust Imperative examines the convergence of AI governance and the post-quantum transition, providing insight into the evolving policy and standards landscape, shared trust requirements, and practical considerations for organizational readiness. Through illustrative applications and five leadership questions, this perspective invites leaders to assess their current capabilities, recognize emerging opportunities, and determine their next steps. The opportunity is to make trust a defining organizational capability, one that supports confident innovation, responsible decisions, and the intentional stewardship of the people, information, resources, and future possibilities entrusted to our care.

1. Why This Matters Now

AI is increasingly integrated into operations, analysis, acquisition, services, and decisions. Public authorities are developing different approaches to transparency, accountability, safety, procurement, and risk management. In parallel, NIST has finalized foundational post-quantum cryptography standards, and migration planning is underway. The leadership challenge is to coordinate these changes across strategy, people, information, technology, and governance. [1–5]

Policy priorities and terminology vary by administration and jurisdiction. U.S. executive-branch memoranda express direction for covered federal agencies; statutory obligations, contract requirements, and voluntary technical frameworks carry different legal effects. A common governance baseline can connect these requirements to the decisions and workflows they affect. For example, an AI-supported acquisition process may depend on verified employee and supplier identities, digitally signed submissions, encrypted exchanges, and third-party platforms. The reliability of the recommendation depends on the integrity of each

connection, including the cryptographic mechanisms protecting identity, records, and communications. This intersection makes AI oversight and post-quantum planning complementary leadership responsibilities. [1, 2, 9]

2. Trust as an Organizational Capability

AI governance and post-quantum migration share essential requirements: clear identity, appropriate access, verified sources, explicit authority, protected information, traceable action, and evidence of outcomes. These elements form a trust architecture connecting people, processes, technology, and accountable decisions.

An information system may retrieve, generate, or connect information. Organizational trust requires an additional determination: whether that information is reliable, relevant, sufficient, and authorized for its intended use.

Responsible AI involves evaluation, oversight, data governance, and accountable use. Post-quantum readiness involves cryptographic discovery, prioritization, interoperability, migration, and validation. Both benefit from coherent governance that connects technical decisions with organizational responsibilities.

3. Trust Has Boundaries

Trust depends on knowing the context in which information and authority may move. Boundaries may be defined by client, role, jurisdiction, system, data sensitivity, contractual commitment, or mission. An organization can establish a repeatable sequence: Know → Trust → Authorize → Act → Evidence.

- Know: establish identity, provenance, context, and the conditions relevant to the decision.
- Trust: evaluate the quality, sufficiency, and limitations of the available evidence.
- Authorize: establish decision rights, permissions, purpose, and operating boundaries.
- Act: carry out approved decisions and workflows within those boundaries.
- Evidence: preserve records of actions, outcomes, exceptions, and learning.

Consider a procurement team evaluating supplier submissions. An analytical system can organize records and identify a preferred supplier, but missing certifications, outdated pricing, or a recommendation based on restricted information could lead to an unsupported award, contractual exposure, or an unfair evaluation. Before approval, the responsible official establishes the completeness and provenance of the record, confirms the evaluation criteria and permissions, reviews material assumptions and exceptions, and documents the authorized

decision and its rationale. This example illustrates how trustworthy action depends on evidence, context, and accountable authority.

4. Federal Policy, Standards, and Executive Priorities

Federal AI policy combines innovation and adoption objectives with agency-specific governance and acquisition requirements. OMB Memoranda M-25-21 and M-25-22, both issued April 3, 2025, address federal AI use and acquisition for covered agencies. OMB M-26-04, issued December 11, 2025, provides additional direction on unbiased AI principles. These memoranda are executive-branch instruments with defined applicability, rather than universal requirements for every business. Federal contractors may encounter relevant expectations through solicitation terms, contract clauses, supplier assurances, or agency acquisition procedures, depending on the procurement and scope of work. The NIST AI Risk Management Framework and Generative AI Profile are voluntary resources unless made applicable by law, contract, or another controlling requirement. [1–4, 10]

Federal policy and procurement direction can evolve. Leaders can distinguish the current agency requirement, the governing contract, and the technical practice being adopted, while preserving evidence of how each applies. The June 24, 2026 OMB memorandum M-26-15 also directs covered federal agencies to execute a prioritized post-quantum migration and develop agency migration plans; it does not apply to national security systems. The current OMB memorandum index lists M-25-21 and M-25-22 alongside subsequent guidance; a final prepublication check remains appropriate. [1, 2, 9, 10]

5. U.S. State and International Approaches

U.S. states address AI through different combinations of consumer transparency, algorithmic discrimination protections, sector-specific requirements, and government-use controls. Definitions, covered actors, enforcement mechanisms, exceptions, and operative dates vary. Appendix A compares selected enacted measures in California, Colorado, and Texas to illustrate these differences. This selected analysis does not purport to be a complete inventory of the 50 states, District of Columbia, and Puerto Rico. [6, 11–13]

International approaches also differ in legal force and scope. The European Union AI Act creates binding, risk-based obligations for covered activities, with requirements applying in stages. The AI Omnibus entered into force July 27, 2026, extending the application of specified high-risk-system requirements to December 2, 2027 (Annex III) and August 2, 2028 (certain Annex I product systems); other obligations have their own dates. The United Kingdom has emphasized context-specific oversight through existing regulators and applicable sectoral law. Canada

combines existing legal requirements and public-sector instruments with a voluntary code for advanced generative AI. Singapore offers a voluntary Model AI Governance Framework for Agentic AI, updated in May 2026. For organizations working across borders, the practical objective is a consistent internal governance foundation mapped to the binding duties, guidance, and contractual commitments applicable to each operation. [7, 14–17]

6. The Post-Quantum Transition

NIST finalized FIPS 203 (ML-KEM), FIPS 204 (ML-DSA), and FIPS 205 (SLH-DSA) in August 2024. These standards address quantum-resistant key establishment and digital signatures. Post-quantum migration focuses especially on public-key mechanisms whose security could be compromised by a sufficiently capable quantum computer. An adversary may collect encrypted information today and attempt to decrypt it in the future; information requiring confidentiality for many years therefore warrants early attention. NIST migration resources and OMB M-26-15 emphasize cryptographic discovery, risk prioritization, interoperability, and testing. [5, 8, 9]

A practical approach begins with a cryptographic inventory: identify vulnerable public-key mechanisms supporting authentication, key establishment, digital signatures, protected communications, long-lived data, and third-party interfaces. Assess the duration of information sensitivity, system criticality, vendor dependencies, and upgrade pathways. Prioritize and test migration in phases, document configuration-specific results, and maintain cryptographic agility as standards and implementations evolve. A migration plan demonstrates preparedness; deployed changes require testing and operational acceptance; independent validation or certification, when required, is a separate determination. [5, 8, 9]

7. Governing Organizational Intelligence

As organizations connect data, knowledge, AI services, and authorized agents, governance can follow the information and decision lifecycle. Evidence, interpretation, assumptions, alternatives, recommendations, decisions, authorized actions, and outcomes are distinct records with different owners, purposes, and evidentiary requirements. Separating them makes it possible to distinguish what was observed, what was inferred, what was approved, and what occurred.

An illustrative operating sequence is Evidence → Interpretation and Assumptions → Alternatives and Recommendation → Decision Authority and Rationale → Authorized Action → Outcome Measurement → Learning. Decision rights remain with the

accountable person or body, while traceable records support review, continuity, and improvement across teams and systems.

8. Five Questions for Leadership

These questions are designed to invite discovery, informed choice, and a meaningful next step:

- What do we know about this use case, and what evidence gives us confidence in that understanding?
- What makes the information and evidence reliable and appropriate for the decision we are making?
- Who has the authority to decide and act, and how will we make those boundaries clear?
- How can we observe performance, recognize emerging risks, and respond when circumstances change?
- What will we measure and learn so that our next decision is better informed?

Together, the questions provide a technology-neutral starting point. Leaders can apply them to a specific AI use case, an agentic workflow, a cryptographic dependency, or an enterprise transformation initiative.

9. Translating Insight into Action

Trust can be developed through a focused sequence applied to one consequential workflow. Leaders can establish its intended value and accountable owner, understand the information and authority involved, examine the AI and cryptographic dependencies, choose proportionate improvements, and evaluate whether the changes produce better decisions and outcomes.

- Select a consequential workflow. Define the intended outcome, affected people, decision owner, and the consequence of an incorrect or delayed decision.
- Establish information and authority boundaries. Trace sources, sensitive data flows, identities, permissions, approvals, and the evidence needed to act.
- Assess technology dependencies. Identify AI-supported recommendations, external providers, cryptographic mechanisms, long-lived information, and potential points of failure.
- Prioritize coordinated improvements. Connect applicable legal, contractual, and technical requirements to proportionate controls, responsible owners, and achievable milestones.
- Measure outcomes and learn. Record decisions and exceptions; assess information quality, authorized actions, control performance, and realized

value; use the evidence to guide the next improvement. Measures may include decision traceability, resolution of material exceptions, and verified completion of priority controls.

10. Trust in Action: From Insight to Stewardship

The convergence of AI adoption and post-quantum migration creates an opportunity to advance how organizations establish trust, exercise authority, and demonstrate accountability. Leaders can begin with one consequential decision, identify the evidence and responsibilities that govern it, and translate that understanding into a measurable improvement.

The next opportunity begins with a decision that matters. Select one workflow where information quality, authority, or technology dependencies influence outcomes. Establish who is accountable, what evidence is sufficient, and how a sound decision will be recognized. Then choose an improvement that can be implemented, observed, and evaluated.

This disciplined practice creates organizational value beyond an individual technology deployment. It improves decision continuity, clarifies responsibilities across teams and suppliers, supports more deliberate investment, and provides a stronger basis for adapting as AI capabilities, security standards, and legal requirements evolve.

The opportunity before us is to make trust an intentional part of how organizations create value and shape the future. Every decision carries influence beyond its immediate outcome. Stewardship invites us to recognize that influence, exercise it responsibly, and consider how the choices we make today advance human potential, organizational progress, and the possibilities available to future generations.

Appendix A. Selected U.S. AI Legislation: Comparative Approaches

Scope. This dated comparative analysis illustrates different enacted regulatory approaches through selected measures in California, Colorado, and Texas. It is not a comprehensive inventory of AI legislation across all U.S. jurisdictions.

Regulatory approaches and leadership implications

Transparency and content provenance

Illustrative enacted measure: California SB 942 (2024). Specified generative-AI providers and covered content; disclosure and provenance requirements, subject to statutory definitions and exceptions.

Leadership implication: Identify which provider and content obligations apply; connect provenance controls to publication and distribution workflows.

High-risk decisions and discrimination

Illustrative enacted measure: Colorado SB 24-205, amended by SB 25B-004. Developer and deployer duties concerning specified high-risk AI systems and algorithmic discrimination; substantive commencement adjusted to June 30, 2026 by the 2025 amendment.

Leadership implication: Map consequential decisions, covered roles, required notices, assessments, and human accountability; reconfirm operative text before reliance.

Prohibited uses and disclosures

Illustrative enacted measure: Texas HB 149 (2025). Responsible Artificial Intelligence Governance Act; specified prohibited practices, transparency and public-sector provisions, with applicability depending on actor and use. Effective January 1, 2026.

Leadership implication: Determine whether activities fall within prohibited-use or disclosure provisions and document the basis for compliance.

What the comparison reveals

These measures illustrate three recurring governance dimensions: the origin and labeling of AI-generated content; the fairness and accountability of consequential automated decisions; and restrictions or disclosures tied to particular uses. A common internal governance approach can clarify identity, evidence, authority, records, and oversight while allowing jurisdiction-specific requirements to be mapped to the relevant workflow. Statutory scope, amendments, and operative dates require review against the official enacted record before any compliance conclusion.

Official legislative records

California SB 942 (2024):

https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB942

Colorado SB 24-205: <https://leg.colorado.gov/bills/sb24-205>

Colorado SB 25B-004: <https://leg.colorado.gov/bills/sb25b-004>

Texas HB 149 (89th Regular Session):

<https://capitol.texas.gov/BillLookup/History.aspx?LegSess=89R&Bill=HB149>

Appendix B. Sources and Citation Map

- [1] U.S. OMB, M-25-21, Accelerating Federal Use of AI through Innovation, Governance, and Public Trust (April 3, 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-21-Accelerating-Federal-Use-of-AI-through-Innovation-Governance-and-Public-Trust.pdf>
- [2] U.S. OMB, M-25-22, Driving Efficient Acquisition of Artificial Intelligence in Government (April 3, 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/02/M-25-22-Driving-Efficient-Acquisition-of-Artificial-Intelligence-in-Government.pdf>
- [3] NIST, AI Risk Management Framework (AI RMF 1.0, 2023). <https://www.nist.gov/itl/ai-risk-management-framework>
- [4] NIST, Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile (AI 600-1, 2024). <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf>
- [5] NIST, Post-Quantum Cryptography standards and program. <https://www.nist.gov/pqc>
- [6] National Conference of State Legislatures, Artificial Intelligence Legislation Database (legislative discovery; verify enacted law in official state records). <https://www.ncsl.org/financial-services/artificial-intelligence-legislation-database>
- [7] European Commission, AI Act Regulatory Framework and implementation timetable. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [8] NIST NCCoE, Migration to Post-Quantum Cryptography. <https://pages.nist.gov/nccoe-migration-post-quantum-cryptography/>
- [9] U.S. OMB, M-26-15, Execution of the Migration to Post-Quantum Cryptography (June 24, 2026). <https://www.whitehouse.gov/wp-content/uploads/2026/06/M-26-15-Execution-of-the-Migration-to-Post-Quantum-Cryptography.pdf>
- [10] U.S. OMB, M-26-04, Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles (December 11, 2025). <https://www.whitehouse.gov/wp-content/uploads/2025/12/M-26-04-Increasing-Public-Trust-in-Artificial-Intelligence-Through-Unbiased-AI-Principles-1.pdf>
- [11] California Legislature, SB 942 (2024), enacted legislative record. https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202320240SB942

[12] Colorado General Assembly, SB 24-205 (2024) and SB 25B-004 (2025), official legislative records. <https://leg.colorado.gov/bills/sb24-205> ; <https://leg.colorado.gov/bills/sb25b-004>

[13] Texas Legislature, HB 149 (2025), official legislative history. <https://capitol.texas.gov/BillLookup/History.aspx?LegSess=89R&Bill=HB149>

[14] European Commission, AI Omnibus enters into force (July 27, 2026), revised high-risk implementation dates. <https://digital-strategy.ec.europa.eu/en/news/ai-omnibus-enters-force>

[15] UK Department for Science, Innovation and Technology, AI regulation: a pro-innovation approach (2023; government response 2024). <https://www.gov.uk/government/publications/ai-regulation-a-pro-innovation-approach>

[16] Innovation, Science and Economic Development Canada, Voluntary Code of Conduct for Advanced Generative AI Systems. <https://ised-isde.canada.ca/site/ised/en/archived-artificial-intelligence-and-data-act>

[17] Singapore IMDA, Updated Model AI Governance Framework for Agentic AI (May 20, 2026). <https://www.imda.gov.sg/resources/press-releases-factsheets-and-speeches/factsheets/2026/updated-model-ai-governance-framework-for-agentic-ai>