THE ARCHITECTURE OF AUTHORITY

How Humanity Removed the Friction Between Knowing and Acting

Joshua Johosky
Founder & Chief Executive Officer | Principal Architect
Oak & Sparrow Systems Enterprise LLC
Manteca, California

Research basis: v1.5 Final Canonical Revision | 11 September 2026
Full-Manuscript External Edition v1.5-E | 14 September 2026

EXTERNAL RESEARCH EDITION

Ability is not authority.

Authorship and published analysis are attributable. Protected source identities and affiliations are not. This edition preserves the full research structure while limiting implementation disclosure.

Publication and Source-Handling Notice

Authorship of this paper is attributable to Joshua Johosky. The author's commercial affiliation is disclosed; it is not evidence for the research findings. The paper may be read, cited with attribution, and shared.

Where a source interaction was expressly conducted under the Chatham House Rule, information may be used only without revealing the identity or affiliation of the speaker or any other participant. Identifying contextual details must also be excluded. This treatment applies to the protected interaction, not to the authorship of this published paper or to its attributable public sources. Use of the Rule does not imply sponsorship, review, or endorsement by Chatham House. [SH1]

This source-handling convention does not create legal privilege or authorize disclosure prohibited by a confidentiality agreement, an NDA, a trade-secret obligation, a security restriction, or another applicable duty. Stricter obligations continue to govern. Material that cannot be separated from protected information is excluded rather than made publishable by removing a name. Non-public material is not used as the sole support for public factual allegations.

Public claims concerning governments, companies, institutions, technologies, and individuals are grounded in attributable public sources. Allegations are identified as allegations. Disputed findings remain disputed. Association is not treated as evidence of participation or wrongdoing.

Private customer and partner environments, non-public commercial relationships and pilots, credentials, implementation-specific architecture, unpublished test results, roadmaps, and protected intellectual property are outside this edition's evidentiary record. Its conceptual analysis is not a product specification, security certification, implementation guarantee, or disclosure of any customer deployment.

The paper distinguishes documented fact, corroborated fact, attributable testimony or allegation, reported allegation, and author inference. A stronger narrative interpretation does not outrank the evidence supporting it.

> **The governing rule is simple: the claim may not exceed the record.**

Claim Classification

Grade	Meaning	Use
A	Primary documentary fact	Statute, court opinion, procurement record, agency document, contemporaneous official record.
B	Independently corroborated fact	May support a central proposition when corroboration is genuinely independent.
C	Attributable testimony or allegation	Mapped and tested; does not prove a central proposition by itself.
D	Reported allegation	Lead only unless independently established.
E	Author inference or theoretical proposition	Must be identified as inference and tested against contrary cases.

Abstract

This paper reconstructs a recurring architectural problem in modern institutions: technical systems have progressively reduced the physical, administrative, and cognitive friction separating information from action. The investigation begins with intelligence and law-enforcement information systems in the United States, then tests the emerging pattern against contractor dependence, commercial cloud infrastructure, local surveillance technology, digital public infrastructure, comparative international systems, and contemporary government AI. The design is conceptual, historical-comparative, and systems-architectural. It is theory-generating rather than a statistical or causal validation study. The paper does not claim a single covert lineage, common political objective, or controlling organization. The continuity identified here is functional.

The paper proposes administrative compression as a middle-range explanatory construct: institutions face persistent incentives to reduce the labor, latency, handoffs, infrastructure, and organizational scale required to resolve information and produce action. Fear and crisis can accelerate that process; competition and commercial reuse can scale it; path dependence can preserve it. The central governance problem appears when systems that remove useless computational friction also remove friction that had functioned as a practical constraint on authority. The paper therefore distinguishes computational friction from governance friction and tests three analytical outcomes: impunity, in which legitimate action fails because knowledge and authority do not connect; abuse, in which capability outruns enforceable authorization; and governed execution, in which the correct information reaches the correct authority, an authorized action executes, and the chain remains reconstructable afterward.

The historical synthesis generates seven prospective hypotheses and an engineering prescription; it does not claim that the cases estimate their causal effects or establish the proposed protocol as optimal. The distinctive contribution is narrower: a cross-era architecture centered on the consequential state transition and on the symmetric problem of wrongful execution and wrongful non-execution - how standing, delegation, purpose, jurisdiction, temporal validity, authorization, obligation, refusal, revocation, evidence, and accountable disposition survive to the point of commitment in increasingly agentic and distributed systems. The resulting proposition is simple: ability is not authority. Architecture can preserve and evidence externally legitimate authority; it cannot manufacture substantive legitimacy, justice, or lawful power.

CONTENTS

PART I The Problem: From Knowing To Acting

PART II Historical Reconstruction

PART III Comparative And Global Tests

PART IV The General Model And Its Falsifiers

PART V From Administrative Friction To Explicit Authority

PART VI Control Cases, Measurement, And Due Process

PART VII The Machine-Speed State

PART VIII The Authority Protocol

PART IX Engineering Convergence And Claim Ceiling

APPENDICES Capability-Authority Ledger; Falsification; Source Register; Research Boundary; Authority Questions; Adversarial Scenarios; Measurement Dictionary

PART I

THE PROBLEM: FROM KNOWING TO ACTING

1. Introduction: The Apparent Problem

The apparent problem is artificial intelligence.

It's not.

Artificial intelligence is just the newest layer in a much older shift in how institutions work.

For most of institutional history, knowing something and acting on it were separated by practical limits. Records had to be found. Identities had to be confirmed. Agencies had to communicate. Jurisdiction had to be sorted out. Evidence had to be interpreted. One person had to make a decision. Often, someone else had to approve it. Only then could anything happen.

Those limits were inefficient. But they also mattered.

They created a gap between what institutions could do and what they could actually carry out.

Modern information systems have spent decades closing that gap.

Paper became databases. Databases became networks. Networks became interoperable systems. Interoperability made large-scale resolution possible. Cloud infrastructure made those systems flexible and scalable. Sensors expanded what institutions could observe. Artificial intelligence is now speeding up interpretation. Autonomous agents are starting to speed up execution.

None of these changes automatically means wrongdoing. Most were built to solve real problems, and many have created real public value.

The question is more specific.

What happens when the technical limits that once kept an institution from acting faster than its authority could be established start to disappear?

This paper looks at that question across historical and contemporary information systems, public administration, intelligence, policing, digital public infrastructure, commercial technology, and artificial intelligence.

The claim is not that these systems share a secret lineage, a common political goal, or a single controlling organization.

They don't need to.

The continuity is functional.

Across different technologies, political systems, threats, institutions, and markets, the same pattern keeps showing up: a limit on what an institution can know becomes visible; technology reduces that limit; the resulting capability becomes infrastructure; the infrastructure becomes something people depend on; and governance has to deal with a boundary that used to exist partly because the underlying capability was technically hard or impossible.

The result is a shift from accidental limits to explicit governance.

The central point of this paper is simple:

> **Ability is not authority.**

A system's ability to observe, resolve, infer, recommend, or act does not make the resulting action legitimate.

For much of history, that distinction was partly enforced by the limits of the technology itself.

We can't assume that anymore.

1.1 The research question

The research question is not whether information technology makes institutions more powerful. That proposition is too broad to be useful. The question is whether a repeatable structure appears when institutions use technology to remove a bottleneck between what they can know and what they can do, and whether that structure explains both successful governance and failure better than narrower stories about surveillance, secrecy, privatization, or artificial intelligence alone.

The investigation therefore treats the historical record as a falsification problem. A theory that explains only abuse is incomplete. A theory that explains only institutional failure is incomplete. A theory that treats integration itself as harmful fails as soon as integrated systems prevent a bombing, deliver a benefit to the correct person, identify fraud accurately, or allow two public institutions to exchange a record under valid authority. The useful model has to explain why the same broad class of information capability can produce public value, abuse, or paralysis depending on where authority, scope, jurisdiction, and accountability sit.

> **Collection is not knowledge. Knowledge is not resolution. Resolution is not authority. Authority is not execution.**

2. Method: Follow the Function, Not the Organization

The first version of this investigation looked for continuity among institutions: Watergate, the Church Committee, NSA collection, INSLAW and PROMIS, post-9/11 information sharing, intelligence contractors, commercial cloud, automated license-plate readers, digital identity, and artificial intelligence. That approach repeatedly ran into a legitimate evidentiary wall. Direct technological or organizational lineage could not be established across the entire sequence. In several places the record contained explicit counterevidence.

That failure improved the method. The relevant continuity is not necessarily organizational. It is functional. Each generation confronts a bottleneck in institutional knowing or acting. A new system reduces that bottleneck. The capability becomes normal. The next generation inherits the new baseline and confronts the next bottleneck.

This paper therefore asks the same questions of each case: What information existed? Who could access it? What could the system resolve that previously required more labor or time? What authority governed access and action? Where did friction remain? What happened when the capability was used? Could the action be reconstructed afterward? What governance response followed failure or success?

Counterevidence is part of the record, not an inconvenience. The INSLAW section includes the Bua findings and later litigation that cut against the strongest misconduct narrative. The post-9/11 section includes successful information-sharing cases. The cloud section treats multi-cloud procurement as a deliberate attempt to reduce single-vendor lock-in. Estonia is used as a control against the proposition that interoperability necessarily requires centralized possession. These are not concessions around the thesis. They are the tests that determine what the thesis is allowed to become.

2.1 The information-action chain

> **OBSERVE → IDENTIFY → RESOLVE → CORRELATE → QUERY → SHARE → INTERPRET → AUTHORIZE → EXECUTE → AUDIT → ACCOUNT**

The chain is intentionally longer than a conventional data lifecycle. The point is to separate technical states that are often collapsed in practice. A record may exist without being discoverable. A person may be identified without the relevant records being resolved to that identity. Information may be resolved without being legally usable. A recommendation may be generated without authority to execute it. An action may execute without a durable record sufficient to establish who authorized it and why.

Those distinctions become increasingly important as latency falls. When every step required human effort measured in hours or days, the boundaries were partly enforced by time. When the same chain can execute in seconds or milliseconds, architecture has to carry distinctions that delay once carried accidentally.

2.2 Research design and inferential scope

The method is structured narrative comparison and theory construction, not a controlled experiment, statistical analysis, or systematic review. Cases were assembled to reconstruct successive changes in collection, retrieval, resolution, sharing, procurement, sensing, interoperability, AI interpretation, and agentic execution, and then deliberately challenged with counterexamples. That design can establish documentary facts, identify recurring mechanisms, generate hypotheses, and test whether a proposed explanation survives obvious disconfirming cases. It cannot estimate causal effect sizes, prevalence, or comparative performance across a defined population.

The seven hypotheses introduced later are therefore prospective hypotheses generated by the comparative synthesis. Terms such as increases, predicts, and outperforms identify propositions for subsequent empirical testing; they are not reported estimates from the present corpus. No observation-level dataset, sampling frame, regression model, confidence interval, or experimental treatment is claimed here.

2.3 Case selection and evidence protocol

The case corpus is purposive rather than exhaustive. Cases entered because they exposed a distinct boundary in the information-action chain, supplied a strong primary documentary record, or offered a meaningful control against an emerging proposition. The corpus is intentionally heterogeneous: national-security collection, prosecutorial information systems, post-9/11 sharing, contractor dependence, cloud procurement, automated license-plate readers, digital public infrastructure, and government AI do not share one legal regime or one political purpose. They are compared only at the functional level identified in this paper.

That choice creates a real limitation. A purposive corpus cannot establish how common the pattern is, and theoretical sampling can become confirmatory sampling if contrary cases are not actively sought. The paper therefore records counterevidence and negative results, uses successful integration and ex-ante governance as controls, and states independent falsifiers. A future empirical study should predefine the case universe, inclusion and exclusion rules, coding protocol, rejected-case register, and outcome measures before observing results.

2.4 Construct boundaries

Administrative compression is not a single scalar. In this paper it is a vector of changes in elapsed time, human handoffs, labor input, cost, systems queried, organizational reach, resolution capacity, and autonomous execution. Those dimensions can move independently. Any future composite index would require validation rather than assumption.

The paper also separates institution-side administrative compression from citizen-side administrative burden. Public-administration scholarship uses administrative burden to describe learning, compliance, and psychological costs experienced by people interacting with the state. Faster institutional processing can reduce those costs, leave them unchanged, or transfer them onto citizens through faster error, harder contestation, or more demanding verification. Compression is therefore not treated as synonymous with social benefit.

Finally, computational friction and governance friction are functional categories rather than labels attached permanently to a human step. The same manual review can be useless delay in one context and a legally significant opportunity for individualized judgment in another. The question is whether the safeguard's function survives, not whether the old form remains slow.

2.5 Positionality and commercial-interest boundary

The author is Founder & CEO and Principal Architect of Oak & Sparrow Systems Enterprise LLC (OASSE), which develops systems concerned with pre-execution governance. That commercial interest creates an obvious reason to keep the evidentiary boundary explicit. Company products, pilots, private partner discussions, and non-public implementation results are not independent evidence for the historical thesis and are not used to validate H1-H7. Any implementation derived from this framework should be evaluated against credible alternatives by independent or reproducible testing. Authorship of this paper is not a determination of patent inventorship, ownership of intellectual property, or the contributions of other engineering participants.

PART II

HISTORICAL RECONSTRUCTION

3. Before the Network: Observation, Collection, and the First Governance Shock

The modern story does not begin with artificial intelligence or even with databases. It begins with the industrialization of collection. The Church Committee's investigation documented intelligence activities that had remained unknown to the public, including NSA Projects SHAMROCK and MINARET. Senate records describe SHAMROCK as involving private telegraph companies turning international communications over to the government over decades. In the later years of the program, NSA analysts reviewed on the order of 150,000 telegrams per month. [S1-S2]

The point is not merely that surveillance occurred. Private communications infrastructure had become part of public intelligence capability. The state did not need to own every wire to gain access to communications moving across them. That public-private pattern will recur later under very different legal regimes and technologies.

Watergate and the intelligence investigations of the 1970s produced the first major governance inversion in this paper. The problem was not insufficient information sharing. It was unchecked collection, secrecy, misuse of authority, and weak oversight. The response included stronger statutory and institutional boundaries. The Privacy Act, congressional intelligence oversight, and the Foreign Intelligence Surveillance Act belong to that corrective era.

This matters because the later history is sometimes told as a simple march toward integration. It was not. Institutions first learned that unconstrained information capability could be dangerous. They deliberately added friction. The later post-9/11 architecture would then discover that some of the resulting barriers could themselves become dangerous when they prevented information from reaching the people who needed it.

> **The first lesson was that collection without governance could become abuse.**

3.1 Private infrastructure, public authority

SHAMROCK establishes an early version of a pattern that later technologies make easier to miss. The government did not need to own the communications network to gain intelligence value from it. Private carriers controlled the physical transport layer; public institutions supplied the intelligence purpose. The distinction matters because later public-private systems are often evaluated as if ownership alone determines power. It does not. Practical capability can be assembled across institutional boundaries.

The correction that followed the 1970s also demonstrates that governance can be added after a capability is already operational. Congressional investigation exposed collection practices that had developed under secrecy and broad executive claims. FISA then inserted a specialized judicial process into a domain where technical access had previously outrun external review. The point is not the 1978 statute solved every later surveillance problem. It is that the institutional response recognized a recurring architectural requirement: the existence of a collection path cannot itself answer whether a particular use of that path is authorized.

This is the first place where the paper's distinction between computational and governance friction becomes visible. Requiring an application, a legal predicate, minimization, or judicial review adds delay. But the delay is not a technical defect. It is part of the authority system. Later generations would repeatedly face the same design problem in different forms: how to remove the friction that prevents legitimate work without also removing the friction that distinguishes legitimate work from mere capability.

> **Access to infrastructure is a capability. The authority to use that access is a separate state.**

4. PROMIS and the Database Turn: From Records to Resolution

PROMIS belongs in this history because it represents a different technical step. The important move was not merely storing more records. It was making case information dynamically searchable and usable across prosecutorial work. The INSLAW dispute later became surrounded by extraordinary allegations, many of which remain disputed or unsupported. This paper does not require those allegations to be true.

The documented contract history is enough to expose the architectural transition. The Department of Justice sought a national implementation of prosecutorial case-management capability. The dispute over enhanced PROMIS then became a dispute over software ownership, technical control, contracting authority, and the practical relationship between a public institution and the private firm whose software embodied a significant part of the capability.

The evidentiary record is deliberately mixed. A bankruptcy court found serious misconduct by the Department; later proceedings altered the legal posture, and a Court of Federal Claims decision favored the government. The House Judiciary Committee majority found strong evidence of misconduct, while the Bua investigation rejected the broader criminal conspiracy allegations, and DOJ later stated that it found no credible evidence supporting those claims. [S3]

The strongest lesson therefore survives without resolving the disputed conspiracy narrative. Once an institution depends on specialized software to organize and retrieve operational knowledge, ownership of the code, technical expertise, data rights, procurement authority, and the ability to replace the supplier become governance variables. The database is no longer merely a tool sitting beside the institution. It begins to shape what the institution can know efficiently.

This is the first clear appearance of epistemic dependency in the paper: the possibility that the institution formally owns the mission while specialized actors possess disproportionate knowledge of the machinery through which the mission is carried out.

> **The database did not merely store institutional memory. It changed the cost of resolving it.**

4.1 What the INSLAW record can and cannot carry

The INSLAW record is unusually useful because it punishes overstatement. One branch of the record contains findings severe enough to support a story of deliberate government misconduct. Another contains later judicial and executive findings that substantially cut against that story. If the architecture depends on choosing only one side, it is not robust enough to use.

The narrower proposition survives both versions. DOJ needed a software system capable of making prosecutorial case information operational across offices. INSLAW asserted proprietary rights in enhancements that it said made PROMIS materially more capable. The government disputed the scope of those rights. The litigation then forced questions that recur in modern technology procurement: Who owns the technical layer? Who understands it well enough to evaluate it? What happens when the institution depends on a system while the supplier and the institution disagree about the rights necessary to operate, modify, or distribute it?

That is why this paper uses epistemic dependency rather than capture as the baseline description. Capture implies a stronger causal claim about who has bent public decision-making toward whose interests. Dependency requires less. It exists when the mission owner cannot easily reproduce, evaluate, replace, or operate an important technical capability without specialized outside knowledge. Dependency can exist with honest officials, lawful contracts, competitive procurement, and no corrupt intent at all.

The distinction is not semantic. It protects the analysis from turning every vendor relationship into evidence of capture. A government can depend on a private technical system and still govern it well. The risk appears when dependency combines with weak internal expertise, proprietary barriers, concentrated procurement authority, poor auditability, or an inability to switch providers without mission loss.

> **Dependency is observable before capture is provable.**

5. The Governance Inversion: When Friction Became the Threat

The September 11 attacks reversed the dominant information-governance problem. The 9/11 Commission concluded that the federal government possessed vast amounts of information but had a weak system for processing and using it. It called for a shift from a 'need to know' culture toward a 'need to share' model and described the desired architecture as a decentralized network crossing old institutional divides. [S4-S6]

The Commission's account of the pre-9/11 'wall' is especially important. Procedures governing intelligence and criminal information sharing were misunderstood and misapplied, producing substantially less coordination than the written rules allowed. The failure was not a lack of collection alone. It was a failure of institutional resolution and handoff. Information existed in parts of the system without becoming actionable knowledge at the correct boundary. [S4]

The response was deliberate integration. Congress and the executive branch built information-sharing procedures spanning federal, state, and local entities. The Information Sharing Environment and fusion-center model sought to make relevant information discoverable and exchangeable across organizational boundaries while retaining rules for access, dissemination, privacy, and audit.

The Commission itself recognized the danger in the correction. It recommended that expanded information sharing be accompanied by privacy safeguards, review, and oversight. Its technical vision even anticipated rule-based access and audit trails for queries across a distributed network. [S5-S6]

That is a critical finding for the present theory. The historical record does not support a simple opposition between integration and governance. The intended architecture was integration under rules. The failure condition arises when the capability to connect information scales faster than the mechanisms that determine who may access it, for what purpose, and with what accountability.

> **The second lesson was that governance friction can become failure when it blocks legitimate information flow.**

5.1 The wall as an information-action failure

The pre-9/11 wall is a particularly strong case because it demonstrates that the existence of information is not the same thing as institutional knowledge. The 9/11 Commission did not describe a government that simply knew nothing. It described agencies, databases, investigators, analysts, and legal rules that held different pieces of a problem without reliably producing the shared operational picture necessary for action. [S4-S6]

That distinction can be stated more precisely. Possessed knowledge means that some component of an institution has the information. Resolved knowledge means that the information has been connected to the correct identity, event, case, or context. Actionable institutional knowledge means that the authorized decision-maker receives it in a legally and procedurally usable form. Those states are not interchangeable.

The post-9/11 reform project therefore attacked a real bottleneck. The Information Sharing Environment was not merely a larger database. It was an attempt to connect policies, procedures, technologies, people, systems, and information across federal, state, local, tribal, and private-sector boundaries. The objective was to make handoff reliable enough that a piece of information did not become operationally irrelevant merely because it sat on the wrong side of an organizational boundary.

But the same reform created a new governance surface. Every new path for legitimate sharing is also a new path that must be scoped, authenticated, logged, and limited. The architecture cannot solve the previous failure by pretending the new capability has no cost. It has to move the control point.

COLLECTION != KNOWLEDGE != RESOLUTION != AUTHORITY != ACTION

6. The Public-Private Information State

The next transition is organizational as much as technical. As information systems became more complex, government increasingly relied on private contractors for engineering, program management, intelligence support, acquisition support, and other specialized functions. GAO repeatedly warned that contractor work approaching inherently governmental functions creates risks to government control and accountability. In 2008, GAO quoted the Director of National Intelligence as saying that the Intelligence Community had been left with little choice but to use contractors for work that might be 'borderline inherently governmental.' [S7-S9]

The problem is not that contractors are inherently improper. GAO also recognizes that contractors can improve economy, efficiency, and effectiveness. The problem is dependency at the boundary where private expertise becomes necessary to understand, maintain, specify, or challenge systems carrying public authority.

That dependency has at least three forms. Procurement dependency appears when an incumbent possesses data, integration knowledge, or infrastructure that makes competition difficult. Epistemic dependency appears when the public institution cannot independently understand or evaluate the capability it is buying. Infrastructural dependency appears when removing or changing the private system would materially impair the public mission.

These conditions do not prove administrative capture. They are capture conditions: structural circumstances under which formal public authority can coexist with practical dependence on actors outside the institution.

The distinction matters. A government can retain legal sovereignty while losing technical autonomy. It can possess the right to decide while lacking enough internal capacity to specify, test, replace, or operate the machinery through which the decision is implemented.

> **Formal authority can remain public while operational knowledge becomes private.**

6.1 When expertise becomes a control variable

Federal acquisition records make the dependency problem concrete. GAO has repeatedly distinguished ordinary contracting from situations in which contractor work approaches inherently governmental or critical functions. The concern is not private participation by itself. It is whether government retains enough qualified personnel, technical knowledge, and decision capacity to remain the actual principal rather than merely the legal owner of the mission. [S7-S9]

GAO's 2010 review of competition in federal contracting provides a particularly clean path-dependence example. For 27 of 47 noncompetitive DOD contracts in its sample, the government could not compete the requirement because it lacked access to proprietary technical data. Officials described some programs as effectively stuck with the incumbent, and GAO reported cases in which reliance on contractor expertise had become so deep that government officials found it difficult even to make decisions or set requirements. [S23]

That is administrative memory expressed through procurement. An earlier decision not to purchase or maintain technical data can become a later constraint on competition. The immediate choice may have saved money or simplified acquisition. Years later, the institution inherits a narrower set of possible actions. Path dependence is therefore not an abstract property of technology markets. It can be documented in the contract file.

The public-private information state emerges when these dependencies stack. Private firms may supply communications infrastructure, software, integration, analysts, cloud, identity services, sensors, or models. Public institutions retain legal authority, but the practical ability to exercise that authority increasingly depends on systems that are designed, operated, or deeply understood outside the state.

> **Sovereignty is not only the right to decide. It also requires the practical capacity to understand and control the machinery that carries the decision.**

7. The Hyperscale State: When Sovereign Capability Runs on Commercial Infrastructure

Commercial cloud moves the dependency problem deeper into the stack. The CIA's C2S procurement placed commercial cloud infrastructure inside the intelligence mission. GAO's 2013 decision on IBM's protest of the AWS award sustained parts of the protest because the CIA's price evaluation lacked a common basis and because a software-certification requirement had been materially relaxed only for the awardee. The decision is evidence of a procurement correction, not proof of capture. [S10]

The larger structural fact is more important. Classified computing increasingly moved onto commercial hyperscale infrastructure. Later multi-cloud programs broadened the supplier base. The Department of Defense's Joint Warfighting Cloud Capability selected AWS, Microsoft, Google, and Oracle to provide enterprise cloud capability spanning classification levels and tactical use. That architecture is partly an anti-lock-in response: instead of placing the entire mission on one provider, the department deliberately created a multi-vendor environment.

That counterexample sharpens the dependency claim. Multi-cloud can reduce vendor dependency while leaving sector dependency intact. The state may be able to switch among providers while remaining dependent on a small commercial class capable of supplying classified hyperscale compute, networking, storage, edge services, and increasingly AI.

The question therefore changes from 'Does one company control the state?' to a narrower and more defensible question: 'Can the state perform the mission at comparable scale without the commercial capability class?' If the answer is no, the relevant dependency is infrastructural even when competition remains healthy.

This is not unique to national security. The European Union's Interoperable Europe Act explicitly identifies technological lock-in and dependence on non-EU providers as strategic concerns while simultaneously pushing public-sector interoperability. The same policy can therefore pursue more integration and more technological sovereignty at the same time. [S11]

> **Vendor competition can reduce lock-in without eliminating dependency on the capability class.**

7.1 Vendor dependency, platform dependency, sector dependency

Cloud infrastructure requires a more precise dependency vocabulary. Vendor dependency exists when leaving one supplier is prohibitively difficult. Platform dependency exists when an institution becomes bound to an interconnected ecosystem of identity, productivity, security, data, and application services. Sector dependency exists when the institution can move among vendors but cannot perform the mission at comparable scale without the small class of providers capable of supplying the required infrastructure.

JWCC is important because it contains both the dependency and the mitigation. DOD explicitly sought a multi-cloud capability spanning unclassified, secret, and top secret environments from the continental United States to the tactical edge. DOD officials also stated that nothing inside the department met that requirement at the time. [S24-S25] The government therefore acknowledged a capability gap while designing a procurement structure intended to avoid a single-provider answer.

That is not evidence that the government surrendered control. It is evidence that modern sovereign capability can depend on a commercial capability class even when procurement remains competitive. The distinction matters because a theory of capture would predict that competition itself is compromised. A theory of structural dependency predicts something different: competition can remain real while the mission becomes difficult to perform outside the sector.

The architecture becomes still more consequential when the same commercial substrate supports data analytics and artificial intelligence. Compute is not merely storage for a later decision. It is increasingly part of the machinery that produces the recommendation, inference, or operational plan. The private layer moves upward from hosting information toward participating in its interpretation.

> **Multi-cloud can reduce dependence on one company while preserving dependence on the hyperscale capability class.**

8. Capability Compression: National Power at Local Scale

The history changes again when capabilities that once required national-scale institutions become cheap enough for smaller organizations. Automated license-plate-reader networks are a clean contemporary example. Flock's network has expanded to roughly 120,000 AI-enabled cameras across 49 states. The technical system allows local officers to search movement information across a network whose scale would have been difficult to reproduce locally only a generation earlier. [S12]

This paper calls that capability compression: technology lowers the organizational size required to exercise a capability. The relevant unit of power is no longer the institution that owns every sensor. It is the actor who can query the network.

The governance record is equally important. Recent reporting documented numerous allegations and cases of law-enforcement personnel misusing license-plate-reader systems for personal surveillance. In response, Flock announced stronger controls including mandatory criminal case numbers and automated review for abnormal search activity. Some departments also introduced or strengthened audits. [S13-S14]

That sequence is almost a laboratory demonstration of governance lag: deploy, normalize, expand access, discover misuse, add controls. The vendor did not create the legal authority of police officers, but product architecture materially affected how easily that authority could be exceeded and how easily misuse could later be reconstructed.

The lesson is not that automated license-plate readers are inherently illegitimate. The same capability can support legitimate investigations. The architectural question is whether the query boundary carries purpose, authority, scope, and audit at the moment of use rather than relying entirely on policy language outside the system.

> **When capability compresses, governance has to travel with it.**

8.1 Query power is a form of compressed organizational power

The significance of a networked ALPR system is easier to see if the unit of analysis is the query rather than the camera. A local department does not need to deploy a national sensor network itself if it can obtain searchable access to observations generated across a much larger network. The organizational burden of surveillance falls while the resolution capability rises.

This is why audit architecture matters. If a query requires no durable purpose field, no case association, no supervisor review, and no anomaly detection, then the practical cost of misuse may be only the effort required to type the search. A policy manual can prohibit misuse, but the technical system has still made the prohibited action cheap. When the interface later requires a case identifier or the system automatically detects abnormal search behavior, governance moves closer to the capability boundary.

That movement should not be overstated. A mandatory field can be falsified. An audit can be ignored. A case number does not establish constitutional reasonableness. Engineered controls are therefore not substitutes for law, training, supervision, courts, or institutional accountability. They are mechanisms for making those external rules harder to bypass invisibly.

Capability compression will recur beyond policing. Cloud APIs, generative models, autonomous tools, biometric services, and commercial data brokers can make capabilities once reserved for large institutions available to small teams or individual operators. The governance problem scales downward with the capability.

> **The smaller the organization required to exercise a capability, the less we can rely on organizational size itself as a control.**

PART III

COMPARATIVE AND GLOBAL TESTS

9. Zooming Out: The Global Resolution Ratchet

The American history initially suggested that fear might be the root cause. The global comparison weakens that conclusion. Different political systems are independently building digital identity, interoperable public services, secure data exchange, cloud infrastructure, and AI-enabled administration for reasons that extend far beyond national security.

The World Bank's Global Digital Public Infrastructure program supports more than 80 countries in building or scaling digital identity, payments, and secure data-sharing systems. The stated goals are access, inclusion, service delivery, and economic opportunity. [S15] The European Union is pursuing cross-border interoperability to reduce administrative burdens and allow public services to function across national boundaries. [S11] China has formally directed construction of nationally integrated government data systems and cross-level, cross-department data sharing. [S16]

Estonia provides the most useful control case. Its X-tee data-exchange layer enables secure exchange among independent information systems using authentication, multi-level authorization, encryption, signatures, and logging. The architecture demonstrates that high interoperability does not require one central database possessing every record. [S17]

That counterexample forces an important distinction: resolution does not require possession. An institution can become highly capable of resolving authoritative information across distributed systems while leaving the source data under separate control. The state can ask who a person is, what authoritative record exists, and whether the requesting service is allowed to retrieve it without collapsing every dataset into a single repository.

The convergence therefore survives regime type. Democracies, authoritarian states, small digital states, developing states, and supranational institutions all gain from making administration more machine-readable and interoperable. Their legal safeguards, political objectives, and institutional constraints differ. The functional pressure toward resolution remains.

> **The global convergence is toward resolution, not necessarily centralized possession.**

9.1 A comparative test across political systems

The global comparison is not intended to erase the legal and moral differences among political systems. Estonia, the European Union, India, Singapore, China, and World Bank-supported DPI programs do not share one theory of rights, one security model, or one relationship between citizen and state. Treating them as equivalent would destroy the comparison rather than strengthen it.

What they do share is an administrative incentive to resolve identity, authoritative records, payments, eligibility, transactions, and interagency data more efficiently. The World Bank now supports more than 80 countries in building or scaling digital identity, payment, and secure data-sharing foundations. [S15] The OECD reports that interoperable data systems exist in 30 of 36 OECD countries, even though adoption across public institutions remains incomplete. [S18]

This is the point at which fear fails as the sole explanation. A benefits portal, tax exchange, digital identity credential, or cross-border diploma verification service may have little to do with terrorism or policing. The common pressure is administrative: reduce repeated verification, reduce manual handoff, reduce time, and allow authoritative information to be reused across institutional boundaries.

Estonia then supplies the architectural falsifier. If the theory required centralized possession, X-tee would break it. Its distributed exchange model shows that an institution can gain resolution capability without turning the exchange layer into a universal data owner. The design choice shifts the governance question from 'Who owns all the data?' toward 'Who may ask which authoritative source for what, under which identity, and with what log?'

> **Resolution is a capability of connection. It does not require universal possession.**

10. The Invisible Hand: Administrative Compression

The global comparison points to a deeper driver than fear. Institutions face more people, more transactions, more data, more law, more mobility, more services, and more interdependence than unaided human administration can process efficiently. Complexity exceeds human administrative bandwidth.

Technology answers by compressing administrative work. A record lookup becomes a database query. Identity verification becomes a digital credential. Interagency correspondence becomes an API call. Infrastructure procurement becomes elastic cloud capacity. Manual pattern review becomes machine-assisted analysis. A sequence that once required many people and days of work can collapse into one machine-mediated transaction.

The selection pressure is straightforward: do more with less. That pressure exists in ordinary government, not only during emergency. It explains why digital public infrastructure spreads in countries that are not responding to a common security crisis and why the OECD treats data, interoperability, digital identity, and AI as foundational government capabilities. The OECD's 2026 Digital Government Outlook reports AI use in almost every OECD government while warning that governance, procurement capacity, risk assessment, and impact evaluation lag adoption. [S18]

Fear remains important, but as an accelerator. A crisis raises the perceived cost of under-capability and shortens the time available for governance design. Competition adds another accelerator: an institution or state that believes peers or adversaries are gaining an information advantage faces pressure to match it. Commercial reuse lowers the cost of adoption further.

The resulting ratchet can be stated without assuming malicious intent. Complexity creates pressure for compression. Compression creates capability. Capability becomes infrastructure. Infrastructure creates dependency and a new baseline. The next problem is solved on top of that baseline.

> **COMPLEXITY → COMPRESSION → CAPABILITY → INFRASTRUCTURE → DEPENDENCY → NORMALIZATION**

10.1 Administration as an information-processing system

Administrative compression is an information-processing response to scale. Every institution has a finite capacity to remember, retrieve, compare, coordinate, decide, and execute. Population growth, legal complexity, commerce, mobility, communications, and data all increase the number of states the institution may need to distinguish. Human labor can absorb some of that growth, but not indefinitely.

The historical sequence is therefore cumulative. Paper externalized memory. Indexes made memory searchable. Databases compressed retrieval. Relational systems compressed association. Networks compressed distance. Digital identity compressed repeated identity proofing. Interoperability compressed institutional handoff. Cloud compressed infrastructure provisioning. Sensors compressed observation. AI compresses interpretation. Agents begin to compress execution.

The sequence is not technologically deterministic. Institutions choose what to digitize, what to connect, and what to automate. Courts, budgets, politics, procurement, social resistance, and technical failure all alter the path. The claim is that each successful compression changes the feasible set for the next institutional decision.

That is why normalization matters. Once a capability becomes ordinary infrastructure, future policymakers do not experience its adoption as a fresh expansion of power. They experience its absence as a loss of capability. The baseline moves.

Yesterday's exceptional capability becomes tomorrow's minimum operating requirement.

11. The Fear-Capability Ratchet

Fear is not the root of the system, but it changes the rate at which the system moves. The Cold War, Watergate, terrorism, cyber conflict, violent crime, pandemics, geopolitical competition, and artificial intelligence have each created periods in which the cost of being unable to see or act appears unusually high.

That changes institutional risk. Removing or declining a capability can create a visible future failure for which a specific official may be blamed. Retaining a capability can create privacy, dependency, or abuse risks whose costs are diffuse, delayed, or distributed across many actors. Under uncertainty, institutions therefore face an asymmetric incentive to retain capability.

Once the capability has been purchased, integrated, staffed, and normalized, the next crisis does not begin from zero. It inherits existing contracts, databases, APIs, sensors, cloud environments, legal authorities, clearances, and institutional habits. Capability is cumulative even when the justification changes.

This is the fear-capability ratchet: crisis accelerates capability; capability becomes infrastructure; infrastructure survives the crisis; the next crisis builds on the inherited baseline.

The theory does not require every capability to persist forever. Programs are terminated, statutes sunset, contracts end, and courts constrain government. The claim is narrower: technical and institutional capability has a tendency to become path-dependent because removal has direct costs while retention preserves optionality.

Fear supplies urgency. Path dependence supplies memory.

11.1 Why crisis changes the adoption curve

Crisis compresses decision time before it compresses technology. An institution that believes failure is imminent will tolerate higher acquisition cost, greater uncertainty, temporary authority, or imperfect governance in exchange for speed. That does not make the resulting decision irrational. It changes which error appears more dangerous.

The asymmetry can be described without pretending to calculate it precisely. The official who declines a capability may later face a visible failure attributed to what the institution could not see or do. The costs of retaining the capability - privacy loss, dependency, mission creep, or future misuse - may be delayed, diffuse, and harder to assign to one decision-maker. The first risk is often legible before the second.

This produces a retention bias. A capability introduced during emergency conditions may acquire ordinary users, workflows, budgets, contracts, integrations, and constituencies. Even when the original emergency recedes, removal now has an operational cost. The institution is no longer deciding whether to adopt a hypothetical capability. It is deciding whether to take away something people have learned to rely on.

The ratchet is therefore institutional rather than conspiratorial. No central actor has to coordinate it. Local actors can make individually defensible decisions that cumulatively move the baseline toward greater capability.

> **Urgency rewards capability now. Accountability often arrives later.**

12. The Disappearing Human Bottleneck

Administrative friction has always been expensive. It is also easy to misunderstand because not all friction serves the same purpose.

Computational friction is the delay created because the system cannot efficiently perform a legitimate operation: the file cannot be found, the databases cannot communicate, the identity cannot be resolved, the analyst cannot process the volume, the infrastructure cannot scale. Removing that friction generally improves the institution.

Governance friction is different. A warrant is friction. Jurisdiction is friction. A second approval is friction. An appeal is friction. Separation of duties is friction. Procurement review is friction. Human review of a consequential edge case is friction. These mechanisms deliberately sacrifice speed because speed is not the only institutional value.

For much of history, computational and governance friction were entangled. The same practical slowness that made legitimate administration inefficient also limited the scale of misuse. A person had to pull the file. A person had to make the call. A person had to notice the anomaly. Those human bottlenecks were not designed as constitutional safeguards, but they sometimes functioned as accidental constraints on power.

Digitization systematically removes those constraints. The correct response is not to rebuild useless inefficiency. It is to distinguish which friction was accidental and which boundary still has to exist when the rest of the system becomes fast.

> **The question is not how much friction to preserve. It is where friction belongs.**

12.1 Selective friction

The phrase selective friction is intended to avoid a false choice. The alternative to a machine-speed state is not a return to paper. Institutions should eliminate delays that serve no legitimate governance function. A person should not have to carry a folder across a building merely to preserve the appearance of control.

The question is which boundaries should remain consequential even when the underlying computation becomes nearly free. A warrant should still mean something when a query takes milliseconds. Jurisdiction should still mean something when databases are interoperable. A second signature should still mean something when the second signer is an automated service. An appeal should still be possible when the original decision was generated instantly.

This reframes human review as one control among several rather than a universal answer. A human can rubber-stamp an automated recommendation. A human can lack enough time or information to challenge it. In some cases a deterministic policy check, separation of duties, cryptographic authorization, rate limit, purpose binding, or mandatory hold may preserve the boundary more reliably than ceremonial human presence.

The objective is not human speed. It is authority integrity.

> **Remove computational friction. Preserve governance friction where legitimacy depends on it.**

13. Two Symmetric Failure Modes

Once the information-action chain is separated into distinct states, two apparently opposite institutional failures become structurally related.

Impunity occurs when relevant information exists somewhere in the system but fragmentation of knowledge, jurisdiction, responsibility, evidence, or authority prevents proportional legitimate action. The failure is not necessarily corruption. It can emerge from compartmentalization, referral, evidentiary thresholds, unclear ownership, incompatible systems, resource constraints, or diffusion of responsibility.

Abuse occurs when capability outruns the controls that should constrain it. The actor can query, correlate, infer, or execute more easily than the institution can verify purpose, authority, proportionality, or accountability.

These failures point in opposite directions. Impunity contains too much disabling friction at the wrong boundary. Abuse contains too little governing friction at the consequential boundary.

A third state is therefore necessary: governed execution. The correct information reaches the correct actor; jurisdiction is valid; authorization is explicit; the action is proportional to the evidence and purpose; execution remains inside the granted capability; and the chain can be reconstructed afterward.

State	Information	Authority boundary	Outcome
Impunity	Exists but fails to resolve/handoff	Fragmented or unusable	Required legitimate action fails
Abuse	Available and actionable	Weak, bypassed, or unaudited	Unauthorized/excessive action executes
Governed execution	Resolved to correct actor	Explicit and reconstructable	Authorized action executes

> **The target is not maximum action or minimum action. It is authorized action.**

13.1 A three-state operating model

The symmetry between impunity and abuse is the point at which the paper stops being a surveillance history and becomes a governance model. Both failures can occur in the same institution. A system may be unable to connect information needed for one legitimate action while simultaneously making another class of action too easy to execute without adequate authority.

The model therefore rejects both maximal integration and maximal separation. More sharing is not always safer. More compartmentalization is not always safer. The relevant question is whether the architecture allows information to cross the boundaries required for legitimate action while preserving the boundaries required for legitimate authority.

Governed execution is the middle state, but it is not compromise for its own sake. It has specific properties: the correct identity is resolved; the information has a lawful purpose; the receiving actor has jurisdiction; the handoff preserves material context; the action is explicitly authorized; the consequence is proportional; and the record survives long enough to reconstruct the decision.

Those properties are testable. They can be examined in an intelligence system, a benefits system, a police query, a cloud control plane, or an AI agent. That portability is what makes the model potentially useful beyond the cases that generated it.

> **The system should be fast everywhere except where legitimacy requires a boundary.**

13.2 False authorization, false inhibition, and the objective function of governance

The symmetry becomes sharper when the two failure states are treated as separate observables. False authorization occurs when a consequential transition commits without a valid authority basis at the moment of commitment. False inhibition occurs when an action that was validly authorized and institutionally due for execution or authoritative disposition is blocked, abandoned, lost in handoff, or left unresolved without a legitimate reason. The first is wrongful action. The second is wrongful non-action.

This distinction moves the paper beyond denial-centric access control. A security system can appear safe by refusing nearly everything. A public institution cannot. If a benefit payment, emergency response, legally required disclosure, authorized investigation, licensing action, or corrective remedy disappears because the authority path is fragmented, the absence of execution is itself a governance outcome. The relevant denominator therefore includes consequential proposals and obligations that should have reached an accountable disposition, not only requests that happened to execute.

The two error classes also have different consequence structures. A false authorization can disclose a secret, move money, alter status, or create an irreversible physical effect. A false inhibition can withhold a benefit, miss a deadline, prevent a required intervention, or strand a lawful request until the right disappears in practice. No single denial-rate metric can represent both.

This paper therefore treats governed execution as a multi-objective problem. The architecture should reduce wrongful execution without converting safety into institutional paralysis, and it should preserve legitimate action without treating urgency as permission. Which error is more costly is a domain-specific question that must be specified externally to the mechanism being tested.

> **Governance must account for both the action that should not have happened and the action that should have happened but did not.**

14. Case Study: The Action Gap

The Epstein federal-investigation record is useful here only if handled more conservatively than the strongest public narratives around it. The documented 2006-2008 federal investigation, the non-prosecution agreement, later criticism of victim treatment, and subsequent institutional reviews establish a complicated record of investigative capability, prosecutorial discretion, referral, and accountability. DOJ's Office of Professional Responsibility ultimately found poor judgment in the resolution but did not find professional misconduct by Department attorneys. [S19]

That finding is important counterevidence against any theory that treats a broad protection conspiracy as established fact. The case should therefore not carry the thesis.

What it can illustrate is the action gap. Institutions can possess substantial information without that information becoming the same actionable institutional knowledge for every prosecutor, investigator, jurisdiction, or later decision-maker. Legal rules do not automatically impute knowledge across every government component. Evidence can be collected without being resolved to the correct case. A referral can transfer formal responsibility while weakening practical ownership. A broad agreement can close future pathways even when later observers believe additional investigation was warranted.

The architectural lesson does not require assigning hidden intent. It is enough to observe that information, authority, jurisdiction, and responsibility are separable. When those states fail to align, the existence of information does not guarantee action.

That is why the original Architecture of Impunity work belongs inside this broader paper as a failure mode rather than as the universal explanation. Institutional impunity is what the system produces when the path from knowledge to legitimate action breaks.

> **Possession by the government is not the same thing as actionable knowledge by the responsible authority.**

14.1 Why the Epstein record is a test, not a foundation

The Epstein record is particularly vulnerable to retrospective over-coherence. Once later facts are known, earlier fragments can look as though they should have produced one obvious institutional conclusion. That is exactly why the case must be handled as an action-gap test rather than proof of a hidden unified actor.

DOJ's Office of Professional Responsibility found no professional misconduct and no evidence that the prosecutors' decisions were based on corrupt or impermissible considerations such as wealth, status, or associations. It nevertheless concluded that former U.S. Attorney Alexander Acosta exercised poor judgment, including by resolving the federal investigation before significant investigative steps were completed and by relying on actions of state officials over whom he had no authority. [S19]

Those findings support a narrower architectural proposition. Formal authority can be divided among actors who do not control one another. A decision can depend on downstream action by another institution. Closing one pathway can alter the practical value of information that remains elsewhere. None of that requires a secret coordinating mechanism.

The case therefore supplies a warning about the phrase 'the government knew.' Government is not one memory, one prosecutor, one database, or one authority. The relevant question is always: which component knew what, when, in what evidentiary form, under whose jurisdiction, with which available action?

> **The government can possess information without the responsible authority possessing actionable institutional knowledge.**

15. Falsification: When Integration Works

A theory of information power that explains only failure is not useful. The same architecture has produced successful outcomes.

Fusion centers provide one control. GAO's 2014 assessment found an average capability score of about 92 out of 100 across the national network, indicating that centers generally had policies and procedures in place for key information-sharing activities, although the score did not itself prove specific homeland-security outcomes. Earlier GAO recommendations led DHS to expand performance measurement beyond baseline capability. [S20-S21]

Section 702 provides another control. Publicly released government examples describe cases in which foreign-intelligence collection produced a lead, the lead was handed to the FBI, domestic investigative authority was used, and a legitimate law-enforcement action followed. Whatever broader controversies surround Section 702, such cases falsify the proposition that integration is inherently abusive.

Estonia provides an architectural control. High interoperability can coexist with distributed data ownership, authenticated exchange, authorization, encryption, signatures, and logging. The existence of the exchange layer does not require the exchange layer to become the owner of every record. [S17]

JWCC provides a procurement control. The government can recognize lock-in risk and deliberately choose a multi-provider architecture. The mitigation may be incomplete, but it demonstrates that dependency is not an inevitable binary state.

The success cases identify the variables that matter: scope clarity, resolution accuracy, jurisdiction alignment, reliable handoff, explicit authorization, proportionality, retained internal capacity, replaceability, auditability, and clear accountability.

> **Integration is valuable when authority, scope, handoff, and accountability remain aligned.**

15.1 What would disprove the model

The success cases are not decorative balance. They define the falsifier. If highly integrated systems routinely produce legitimate outcomes without clear authority, audit, or accountability, then the execution-boundary thesis is too strong. If fragmented systems routinely avoid impunity despite unreliable handoff, then the action-gap model is too strong. If vendor dependence does not reduce reversibility or internal control, then the path-dependence claim is too strong.

The fusion-center record is useful because it shows institutional learning. GAO reported that the 2013 annual assessment produced an average capability score of about 92 out of 100 and that the score generally indicated policies and procedures were in place for key information-sharing activities. GAO also cautioned that the score did not establish specific homeland-security impacts. [S20] That distinction is exactly the kind of claim ceiling this paper requires.

Likewise, the fact that a Section 702-derived lead can move through an interagency handoff into a lawful domestic investigation demonstrates that integrated collection and sharing can produce legitimate value. It does not settle every constitutional or policy controversy around Section 702. The case proves only what it proves: a high-capability information chain can work when the handoff reaches a valid authority and the subsequent action occurs under a lawful process.

A useful theory has to survive both sides of the record. The same variables should explain why one chain works and another fails.

> **A theory that explains only its favorite failures is not a theory of the system.**

16. From Legibility to Interpretability

The historical sequence has so far industrialized memory, retrieval, identity, exchange, and correlation. Artificial intelligence changes the next bottleneck: interpretation.

Traditional record systems assume that a human already knows what to search for. Relational systems allow the human to search attributes and relationships. Networked surveillance allows selectors to be resolved across communications or movement. Machine-learning systems increasingly allow the human to state a pattern, description, risk, anomaly, or objective and ask the system to identify what deserves attention.

That is a qualitative shift because the system begins to participate in target formation. It is no longer only retrieving a known record. It is helping determine which records, people, transactions, or events become salient.

The global adoption data show that this transition is already administrative rather than speculative. The OECD reports AI use in almost every OECD government while also reporting that evaluation, risk assessment, procurement support, and other governance mechanisms remain uneven. [S18]

International governance is beginning to respond. The Council of Europe's Framework Convention on Artificial Intelligence applies principles of accountability, transparency, oversight, risk and impact management to public authorities and, in relevant ways, private actors acting on their behalf. [S22] The architecture is again following the capability.

> **The machine is moving from finding the record to helping decide what deserves to become a record of concern.**

16.1 AI changes the bottleneck, not the principle

AI does not create the distinction between ability and authority. It changes how often the distinction must be enforced and how quickly a failure can propagate. The historical systems in this paper generally required humans to formulate queries, interpret outputs, and carry recommendations into separate execution systems. AI compresses the interpretation step and can increasingly generate the next action itself.

The OECD's 2026 evidence captures the asymmetry. AI is used in at least one area of government in 97 percent of OECD countries, yet only 28 percent report financial or non-financial assessment of government AI use cases, and central procurement support remains incomplete. [S18] Adoption is therefore not waiting for governance maturity.

The Council of Europe Framework Convention represents a different response: build accountability, transparency, oversight, risk and impact management, and procedural safeguards around the lifecycle of AI systems. The European Union ratified the Convention in May 2026. [S22] The significance for this paper is not that the treaty endorses one technical architecture. It is that international law is attempting to make governance obligations explicit as machine capability expands.

The next step is agentic execution. Once a system can invoke tools, the output is no longer merely speech. It can become a state transition in another system. At that point, the authority boundary has to survive across APIs, credentials, delegation chains, and machine identities.

> **AI accelerates interpretation. Agents make authorization an execution problem.**

17. The Execution Boundary

Interpretation is not the endpoint. Agents move the problem one step further.

An analytical system can produce a recommendation that a human must review. An agentic system can potentially acquire credentials, call tools, modify records, send communications, move funds, provision infrastructure, change permissions, or invoke other systems. The bottleneck after interpretation is execution.

At that boundary the old assumption becomes dangerous: if the system can do it, the system can do it. Technical capability silently becomes operational permission.

That equivalence has never been legitimate. It was merely less visible when the machinery required a person to carry the recommendation across the final boundary.

The architecture of authority therefore has to make the boundary explicit. Identity is not permission. Authentication is not authorization. A model recommendation is not authority. Access to a tool is not authority. A queued action is not authority. Execution is the point at which the system must be able to establish not merely that an action is possible, but that this actor, under this authority, for this purpose, within this scope, at this time, may cause this state transition.

ABILITY ≠ AUTHORITY

AUTHORIZED ACTION = EXECUTED ACTION

17.1 What an execution boundary has to establish

An execution boundary is not one approval screen. It is the set of conditions that must be true before a consequential state transition can occur. At minimum, the system must be able to establish the actor, the delegated authority, the target resource, the permitted action, the purpose or policy basis where required, the relevant time or context, and the evidence necessary to reconstruct the decision.

This becomes more difficult in agentic systems because delegation can be recursive. A human authorizes an application. The application invokes an agent. The agent invokes a model. The model selects a tool. The tool calls an external service. The external service changes state. Authentication at every hop can be perfectly valid while the final action still exceeds the authority originally granted.

The architecture therefore has to distinguish identity continuity from authority continuity. A valid credential proves that a principal or service is recognized. It does not prove that every action reachable through that credential remains inside the original grant. Delegation needs scope, constraints, and evidence of provenance.

A queue can contain a technically valid request. A tool can accept it. A credential can authenticate it. None of those facts answer the authorization question. The distinction concerns the authority for the action, not the successful movement of the request through an operational system.

> **Authentication answers who or what is present. Authorization answers what may happen next.**

RELATED WORK

AUTOMATION, DUE PROCESS, ACCOUNTABILITY, AND ADMINISTRATIVE BURDEN

Related Work: Where This Paper Sits

The broad proposition that information technology changes administrative discretion is not new. Bovens and Zouridis described a shift from street-level to system-level bureaucracies in which system designers become consequential administrative actors and constitutional questions of due process and fairness move into technical systems. [S45] Citron's technological-due-process account similarly showed that automated administrative systems can weaken notice, hearings, transparency, and participatory safeguards when legal rules are translated into code. [S46] Those works are foundational antecedents to the disappearing-human-bottleneck and policy-compiler problems developed here.

The accountability literature also precedes this paper's claim that technical architecture can help enforce institutional rules. Kroll and colleagues developed technical mechanisms for verifying procedural regularity and compliance with legal standards in automated decision systems. [S47] Yeung's account of algorithmic regulation distinguishes systems that recommend from systems that automatically execute changes, placing legitimacy at the center of algorithmic governance. [S48] Wieringa's systematic review demonstrates that algorithmic accountability is already a mature field with established questions about actors, forums, criteria, relationships, and consequences. [S50]

More recent sociotechnical and lifecycle work narrows the novelty further. Selbst and colleagues warn that technically elegant interventions fail when their abstraction boundary excludes the surrounding social and institutional system. [S49] Raji and colleagues propose end-to-end internal auditing and lifecycle documentation so organizational decisions and harms can be reconstructed across development and deployment. [S51] Busuioc identifies public-sector expertise asymmetries and the weakness of nominal human oversight when reviewers lack the information or power required to challenge automated systems. [S52]

The administrative-burden literature supplies an additional boundary. Moynihan, Herd, and Harvey define burden from the citizen side through learning, compliance, and psychological costs. [S53] This paper uses administrative compression from the institution side: the labor, latency, handoffs, infrastructure, and processing capacity required to move from information toward action. The two can move in opposite directions. Institutional compression can reduce citizen burden, or it can transfer burden by making an erroneous or opaque decision faster and harder to contest.

Technical authorization prior art: bounded comparator

The authorization-engineering primitives are not blank-slate inventions. Zanzibar demonstrates a globally distributed relationship-based authorization service separated from application logic. [S54] WAVE demonstrates decentralized authorization with transitive delegation and cryptographically enforced attenuation. [S55] NIST Zero Trust Architecture separates policy decision and enforcement functions and rejects implicit trust based on network position; SP 800-207A extends that model toward granular application-level enforcement in cloud-native and multi-cloud systems. [S56][S57] OAuth Rich Authorization Requests carries structured, fine-grained authorization information beyond coarse scopes. [S58] Cedar provides analyzable policy-as-code outside application logic, and its verification-guided development work shows how authorization semantics can be tested with mechanical proof, differential testing, and property-based methods. [S59][S60]

The 2026 standards landscape narrows the technical novelty further. OpenID AuthZEN Authorization API 1.0 is a final specification for interoperable communication between policy decision points and policy enforcement points. [S65] The IETF OAuth Transaction Tokens working-group draft carries user or workload identity together with authorization and request context through a multi-service call chain in short-lived, transaction-bound tokens. [S66] An August 2026 Internet-Draft on delegated refresh tokens for OAuth Token Exchange preserves subject and actor chains, confines resources and scopes, re-evaluates current policy, bounds delegation lifetime, and provides task-scoped revocation for long-running delegated workflows. [S67]

Agent-specific work reaches the same execution boundary from another direction. NIST treats software- and AI-agent identification, authorization, auditing, and non-repudiation as an active standards problem. [S42] AWS Dogwood explicitly places a runtime control layer at the agent tool-call boundary because tool calls are where agent behavior begins to affect the external world. [S68] AIP, Overlaying Governance, and Mandato propose closely related delegation, attenuation, provenance, and mandate-enforcement mechanisms; these are 2026 preprints and are treated here as emerging technical comparators rather than settled standards. [S62][S63][S64] The broader agent-security literature also makes clear that authority enforcement does not replace prompt-injection, model, tool, supply-chain, or ordinary software-security defenses. [S61]

These systems occupy important technical territory, but they do not by themselves resolve the institutional predicates at the center of this paper. A valid policy decision, delegated capability, or transaction-bound token does not independently establish lawful standing, jurisdiction, evidentiary sufficiency, proportionality, an institutional duty to act, a viable appeal, or the legitimacy of the rule being enforced. The relevant contribution here is therefore not the invention of authorization. It is the attempt to keep those institutional predicates attached to consequential action - and consequential non-action - as execution becomes faster, recursive, distributed, and autonomous.

Novelty claim after social-science and authorization prior art

The contribution claimed here is therefore not that automation changes discretion, that code can affect due process, that algorithms require accountability, that human review can become ceremonial, or that authorization can be externalized into software. Nor are transitive delegation, policy-as-code, fine-grained authorization, permission attenuation, transaction binding, or runtime tool-call gating independently new. Those propositions and mechanisms have substantial prior literature and deployed antecedents. The narrower contribution is a cross-era institutional synthesis focused on the authorization claim attached to the consequential commit and on the symmetric treatment of wrongful execution and wrongful non-execution: how standing, delegation, purpose, jurisdiction, temporal validity, evidentiary sufficiency, obligation, refusal, revocation, appeal, remedy, and reconstructability survive recursive and distributed machine execution.

That claim ceiling also limits the engineering prescription. Mature authorization systems can decide and enforce increasingly rich permissions, and contemporary standards can preserve authorization context through distributed call and delegation chains. The remaining institutional question is whether the authority represented by those mechanisms was legitimate for this actor, purpose, jurisdiction, time, evidence state, and consequence - and whether an authorized or required action silently failed rather than reaching an accountable disposition. An authority protocol can preserve and evidence authority supplied by law, policy, consent, institutional role, or another legitimate source. It cannot determine that the underlying law is just, make an unlawful power lawful, resolve every contested standard into code, or replace courts, politics, supervision, and affected-person remedies.

PART IV

THE GENERAL MODEL AND ITS FALSIFIERS

18. A General Model of Information-Age Institutional Authority

The comparative record supports a model with five forces and one recurring governance problem.

Complexity creates the need. Administrative compression supplies the incentive to reduce human effort. Fear and crisis accelerate adoption. Competition and commercial markets scale the available capability. Infrastructure and path dependence preserve each gain as the baseline for the next generation.

The recurring governance problem is synchronization. Capability, authority, jurisdiction, and accountability do not automatically scale together.

If capability grows faster than governed authority, abuse risk increases. If the action required of an institution exceeds its ability to resolve information and coordinate authority, impunity and failure risk increase. The target operating region is the alignment of capability, authority, and accountable execution.

This is not a numerical equality. It is an architectural requirement: every consequential action should have enough capability to execute what legitimate authority requires, but no consequential action should execute merely because the technical system can perform it.

18.1 The seven control boundaries

Scope

The information and capability are available for a defined legitimate purpose rather than unlimited future use.

Identity and resolution

The system has resolved the correct person, entity, record, or target with an accuracy appropriate to the consequence.

Jurisdiction

The actor or institution receiving the information has authority over the relevant matter.

Handoff

The information reaches the correct decision-maker without losing material context or evidentiary status.

Authorization

The action is legally and procedurally permitted for the actor, purpose, scope, and time.

Proportionality

The consequence matches the evidence, risk, and authority rather than merely the maximum capability available.

Auditability

The institution can reconstruct what was known, what was queried, which authority applied, what executed, and who was responsible.

18.2 The causal model and its limits

The model can now be stated in full. Complexity creates pressure for administrative compression. Compression produces new capability. Fear and competition can accelerate adoption. Markets make the capability reusable. Infrastructure and procurement preserve it as a baseline. Governance then determines whether the new capability produces governed execution, abuse, or impunity.

The model does not predict outcomes from technology alone. The same technology can sit in different legal and organizational systems and produce different results. The modifiers are the control boundaries: scope, identity and resolution, jurisdiction, handoff, authorization, proportionality, auditability, internal capacity, replaceability, and accountability.

This makes the framework falsifiable at the case level. For any consequential system, identify the capability, identify the authority, locate the friction, identify who can remove it, observe what executed, and determine whether the decision can be reconstructed. If the framework cannot distinguish a legitimate success from an abuse or a failure to act, it has not earned generality.

The strongest claim this paper can presently make is therefore architectural rather than predictive: as institutional information-processing capacity approaches machine-speed, legitimacy increasingly depends on preserving explicit authorization, accountability, and audit boundaries that were previously supplied in part by human and organizational friction.

> **Capability scales automatically with technology. Legitimacy does not.**

19. Adversarial Falsification: Trying to Break the Model

A framework this broad can become unfalsifiable. If every capability increase counts as compression, every commercial relationship as dependency, and every later control as governance lag, the theory can explain anything after the fact. That makes it rhetorically useful and analytically weak. The analysis therefore treats the strongest alternatives as competing models and asks what evidence would force the architecture of authority to narrow or fail.

19.1 Alternative explanation one: this is ordinary modernization, not a new authority problem

The simplest counterargument is that the entire history can be explained as ordinary modernization. Governments replaced paper with databases, databases with networks, and local servers with cloud infrastructure because newer tools are faster and cheaper. On this view, there is no distinct architecture-of-authority problem. There is only technology adoption plus familiar public law.

That explanation accounts for much of the adoption pressure. It does not fully account for the repeated need to redesign control surfaces after the capability changes. SHAMROCK did not merely process telegrams faster; the scale and secrecy of collection became an oversight problem. Post-9/11 reforms did not merely buy better databases; they changed who could share what with whom. Flock's later requirement for case identifiers did not increase recognition accuracy; it inserted purpose and accountability into the query path. The Council of Europe AI convention and OECD AI-governance work similarly focus on accountability, risk, transparency, and oversight rather than computational performance alone.

The modernization explanation therefore remains part of the model, but it is incomplete. Administrative compression explains why institutions adopt capability. It does not eliminate the separate question of how authority survives the change in capability.

19.2 Alternative explanation two: law is enough; software does not need to carry authority

A stronger legal objection is that authorization already exists in constitutions, statutes, regulations, warrants, contracts, and agency policy. Software should execute what officials tell it to execute. Requiring the technical system itself to participate in authorization risks turning legal judgment into code.

That objection is partly correct. Architecture cannot manufacture lawful authority. A policy engine cannot decide that an agency possesses a power the law never granted. The claim in this paper is therefore not that software replaces law. It is that technical systems increasingly determine whether legal limits are enforceable at operational speed. Logging, role separation, purpose binding, approval requirements, retention rules, query constraints, and credential scope are technical properties through which an existing legal boundary can become easier or harder to violate.

The falsifier here is straightforward. If high-speed consequential systems consistently remain within valid authority without technical enforcement, durable audit, or reconstructable authorization, then the execution-boundary claim is overstated. The present record does not establish that proposition. It instead shows repeated institutional investment in exactly those control mechanisms.

19.3 Alternative explanation three: human review is the missing safeguard

It is tempting to reduce the solution to a human-in-the-loop rule. That would be too easy. Humans can rubber-stamp machine recommendations, misunderstand system behavior, operate outside their jurisdiction, misuse legitimate credentials, or become the least auditable component in the chain. A slow human decision is not automatically a legitimate one.

The relevant variable is not whether a human clicked approve. It is whether the approving actor possessed the required standing and authority, whether the decision was made against sufficient evidence, whether the action stayed inside scope, and whether the chain can later be reconstructed. Human review is one possible governance mechanism. It is not the definition of governance friction.

19.4 Alternative explanation four: vendor dependency is just specialization

Private specialization is normal. Governments do not manufacture every vehicle, router, database, or processor they use. A dependency theory that treats every supplier relationship as capture would be useless.

Recent GAO work sharpens the distinction. In 2026, officials at 11 of 24 selected federal agencies reported new technical considerations around multi-vendor cloud interoperability, while 10 reported resource constraints affecting cloud workforce acquisition. GAO also recorded the Department of Defense's objection to treating vendor lock-in as inherently negative: provider-specific capability can be strategically correct when dependencies are understood and mission-critical workloads have documented exit strategies. [S26] That is important counterevidence to any blanket anti-dependency claim.

The narrower variable is reversibility. Dependency becomes governance-relevant when the institution cannot credibly understand, replace, exit, or constrain the capability without unacceptable mission loss. GAO has separately documented cloud lock-in risks arising from provider-specific services, government skill gaps, restrictive licensing, and even contractual arrangements requiring an agency to pay to regain ownership of data at contract end. [S29-S31] The distinction is therefore not public versus private. It is specialization with credible exit versus specialization without practical autonomy.

19.5 Alternative explanation five: crisis capability does not always ratchet

The fear-capability ratchet is a tendency, not a conservation law. Programs terminate. Authorities sunset. Courts narrow powers. Contracts are recompeted. Technologies become obsolete. A theory that predicts permanent retention of every emergency capability is false.

The narrower hypothesis concerns inherited baseline. Even when a particular program disappears, institutions often retain some combination of technical expertise, legal concepts, procurement vehicles, data standards, infrastructure, or expectations created during the expansion. A valid disconfirming case would show a major crisis capability removed together with its enabling infrastructure and institutional assumptions, leaving no meaningful increase in the next baseline. Such cases should be actively sought rather than explained away.

19.6 Alternative explanation six: high interoperability can be governed without an execution gate

Estonia is the strongest control against an overly centralized reading of this paper. X-tee demonstrates that high administrative resolution can be built around distributed ownership, authenticated exchange, authorization, encryption, signatures, and logging. It does not require a universal central database. The case also warns against making the execution boundary do work that belongs earlier in the chain.

Some systems are primarily information-exchange systems. Their consequential boundary may be disclosure rather than downstream physical action. In those systems, the relevant authorization check belongs at the data exchange itself. The general principle survives only if 'execution' is defined as the consequential state transition under analysis, not merely as an AI agent calling a tool.

19.7 Alternative explanation seven: abuse can occur in low-technology systems

Nothing in this paper implies that advanced technology is necessary for abuse. Low-capability institutions can abuse authority through coercion, discrimination, corruption, unlawful detention, falsified records, or simple human misconduct. Technology changes scale, latency, reach, reproducibility, and the number of actions one actor can initiate. It does not create the underlying possibility of illegitimate power.

This is a useful limit on the theory. The architecture of authority is a model of information-mediated institutional action, not a universal theory of political abuse.

19.8 Alternative explanation eight: successful high-capability systems may precede governance rather than follow it

Governance lag is not universal. Some institutions conduct impact assessments, procurement review, security authorization, privacy analysis, testing, or staged pilots before broad deployment. OECD data show that such practices exist even though they are uneven: 39 percent of OECD countries require pre-deployment AI risk assessments, 33 percent report internal review committees, and 31 percent conduct post-deployment audits. [S30] These are not failures of the model. They are the control group the model predicts should perform better if the relevant safeguards are substantive rather than ceremonial.

The empirical burden is therefore higher than showing that governance exists. Future work has to test whether pre-deployment controls actually correlate with fewer unauthorized outcomes, faster detection, better reversibility, or clearer accountability. The present paper establishes the architectural variables. It does not yet establish their causal effect size.

19.9 What survives the attack

After these objections, several earlier formulations have to be narrowed. Administrative compression is not the exclusive cause of digital government. Fear does not explain every deployment. Vendor dependency is not capture. Human review is not synonymous with authority. Integration is not inherently dangerous. Governance does not always arrive late. Advanced technology is not necessary for abuse.

What survives is smaller and stronger. First, technical systems can materially change the cost, speed, and scale of institutional action. Second, legal authority and technical capability are separate variables. Third, removing technical friction can remove practical constraints that previously limited both legitimate action and misuse. Fourth, successful systems require the relevant authority boundary to remain enforceable and reconstructable at the speed and scale of the capability. Fifth, the correct location of that boundary depends on the consequential state transition being governed.

> **The theory survives only in its narrower form: capability changes the action surface; authority still has to govern the transition.**

19.10 Prospective hypotheses generated by the comparative analysis

The propositions below are hypotheses for future testing. The historical and comparative material motivates them and supplies plausibility and falsification cases; it does not estimate their causal effects or validate their comparative performance.

H1 - Compression hypothesis. Holding other factors constant, technologies that reduce the labor, latency, or organizational scale required for a consequential institutional action increase the number or speed of actions the institution can practically initiate.

H2 - Misalignment hypothesis. When practical capability expands faster than enforceable authorization and audit, the opportunity set for unauthorized or excessive action expands.

H3 - Fragmentation hypothesis. When information, jurisdiction, evidentiary usability, and responsibility remain fragmented beyond the institution's coordination capacity, the probability of required legitimate action failing increases.

H4 - Selective-friction hypothesis. Systems that remove computational friction while preserving explicit control at the consequential authority boundary should outperform systems that either preserve generalized fragmentation or remove both forms of friction.

H5 - Reversibility hypothesis. Private specialization becomes a stronger governance risk as switching cost, proprietary dependence, skill asymmetry, and absence of credible exit increase.

H6 - Governance-timing hypothesis. Governance lag should be most pronounced where urgency, uncertainty, and visible cost of under-capability are high; it should be weaker where deployment is staged and ex ante controls are operational rather than ceremonial.

H7 - Distributed-resolution hypothesis. High institutional resolution can coexist with lower central-possession risk when exchange is federated, authority is verified at the boundary, and source ownership remains distributed.

These hypotheses can fail independently. That is intentional. A framework becomes more useful when future evidence can remove one mechanism without forcing every other mechanism to survive by definition.

PART V

FROM ADMINISTRATIVE FRICTION TO EXPLICIT AUTHORITY

20. The Constitutional Function of Friction

The word friction can sound like an argument for inefficiency. That is not the claim. Modern administration contains enormous amounts of delay that serve no legitimate constitutional or institutional purpose. A lost file is not due process. An incompatible database is not separation of powers. A six-week wait caused by a broken intake system is not a safeguard. The point is that institutions historically experienced several kinds of delay at the same time, and digitization often removes them together. Once that happens, the architecture has to distinguish the delays we wanted to eliminate from the boundaries we still need.

A constitutional system is full of deliberate impedance. Warrants make some searches slower because an independent authority must evaluate a justification. Appeals make finality slower because an affected person receives another decision point. Separation of duties makes execution slower because one actor cannot initiate, approve, and conceal the same consequential act. Notice-and-comment procedures slow rulemaking because the state is expected to expose a proposed exercise of power before it becomes binding. Procurement review slows acquisition because public money and public capability are not supposed to move on the same terms as a private purchase. None of these mechanisms exists because slowness is intrinsically virtuous. Slowness is the visible cost of a deeper requirement: authority has to be established before some state transitions are allowed to become final.

That distinction becomes easier to see in software. A database can return an answer in milliseconds. The existence of the answer says nothing about whether the requesting actor may use it for the proposed purpose. An identity provider can establish who is making a request. Authentication still does not establish what the authenticated actor may do. A model can infer that a transaction resembles fraud. The inference does not create legal authority to freeze an account, deny a benefit, open an investigation, or publish an accusation. A workflow engine can route an action to the correct endpoint. Routing is not authorization.

The machine therefore exposes a distinction that paper administration could obscure. In a slow institution, permission is often embedded in the social process around the action: the person knows which supervisor signed the form, which office owns the matter, which exception applies, and which consequences require another review. When that process is converted into APIs and autonomous workflows, the informal context can disappear while the technical ability remains. The system becomes excellent at moving the object and weak at proving why the object was allowed to move.

This is why the relevant design objective is not frictionlessness. It is selective friction. Computational friction should continue to fall where it serves no legitimate purpose. Governance friction should become more precise, more machine-readable, and in many cases faster than its paper predecessor. A good authorization boundary does not recreate a filing cabinet. It replaces an accidental delay with an explicit decision. The ideal system can be both faster and more constrained because it spends less time discovering information and more certainty establishing the authority attached to the action.

The historical arc of this paper can be restated in those terms. The Church-era correction added deliberate boundaries after collection capability outran oversight. The post-9/11 correction removed disabling barriers that prevented legitimate handoff while attempting to preserve privacy and audit. Modern digital public infrastructure demonstrates that interoperability can be designed around authenticated exchange rather than indiscriminate possession. Contemporary AI governance increasingly places risk assessment and approval before deployment and monitoring after deployment. These are different institutions solving different problems, but the recurring move is the same: relocate friction from technical incapacity to explicit governance.

That relocation matters because technical incapacity is a poor safeguard. It is arbitrary, uneven, and temporary. It protects rights only until someone builds a better system. A legitimate boundary has to survive technological improvement. If a right exists only because the database is slow, the right disappears when the database becomes fast. If jurisdiction exists only because agencies cannot exchange data, jurisdiction disappears when interoperability arrives. If human review exists only because the model cannot execute, review disappears when the model receives tools. Durable governance cannot depend on the persistence of technical weakness.

The constitutional function of friction is therefore not to make the state slow. It is to keep consequential power conditional. The condition may be a warrant, statutory authority, delegated role, consent, evidentiary threshold, budget approval, safety rule, or other legitimate source of permission. The engineering problem is to make that condition visible at the point where capability becomes consequence.

21. Authority as a State-Transition Property

Most information-governance frameworks are organized around data: who may collect it, where it may be stored, how long it may be retained, who may access it, and when it must be deleted. Those questions remain essential. They are not sufficient for agentic systems because an agent can use perfectly lawful access to cause a consequential state transition. The governance object is no longer only the information. It is the action.

A state transition is any operation that changes the world represented by the system or causes a downstream system to do so. Sending a message changes a communications state. Changing a permission alters an access-control state. Issuing a payment changes a financial state. Creating a case changes an institutional state. Provisioning infrastructure changes an operational state. Filing a report, suspending an account, denying a claim, scheduling an inspection, or opening a physical door are all transitions. Some are reversible. Some are not. Some are trivial. Some alter rights, money, liberty, reputation, or safety.

Traditional access control answers a narrower question: may this identity access this resource or operation? That remains necessary, but it does not always encode purpose, current evidence, delegated authority, temporal validity, conflict among policies, or the difference between proposing and committing an action. A credential may prove that an employee belongs to an organization while saying nothing about whether this particular employee may execute this particular action in this particular case.

The execution boundary is where those variables have to converge. At minimum, a consequential transition should be able to answer: who or what is acting; on whose authority; against which object; for what declared purpose; under which rule; using which evidence or facts; within what scope; for how long; with which required approvals; and with what durable record. The precise implementation will differ by domain. The invariant is that capability alone cannot fill in the missing answers.

This also clarifies the role of models. A model can be extremely useful in the evidence path without occupying the authority path. It can classify, summarize, detect anomalies, generate options, estimate risk, or recommend a next step. Those outputs can become facts considered by a policy engine or human authority. The model's confidence does not become jurisdiction. Its fluency does not become standing. Its access token does not become delegation.

The distinction is especially important in multi-agent systems. One agent may discover a resource, another may plan, another may retrieve data, another may call an external service, and another may commit the result. If authorization is inferred from the existence of the chain, delegation becomes self-expanding. Every successful handoff can silently widen the effective authority of the system. A governed architecture instead treats delegation as an object with scope, provenance, expiration, and constraints that survive each handoff.

The same principle applies to queues. A queued action is evidence that something requested execution. It is not evidence that execution was authorized. In conventional systems, the distinction is easy to miss because queue consumers often inherit broad service credentials. In an agentic environment, that design turns the queue into an authority laundering mechanism: once a proposal enters the trusted execution path, downstream components assume the upstream component already resolved permission.

A mature execution architecture therefore separates proposal from commitment. The proposal can be probabilistic, creative, or model-generated. Commitment should be governed by a deterministic or otherwise independently enforceable authority mechanism appropriate to the consequence. This does not mean every rule must be simple or every policy static. It means the component proposing the action cannot be the sole source of permission for its own consequence.

The resulting model is more precise than human-in-the-loop. A human can be one source of authority, but a human click is not automatically legitimate either. A person can approve outside jurisdiction, approve without reading, approve under the wrong role, or approve an action prohibited by law. The goal is authority-in-the-loop, with human review used where human judgment is actually part of the valid authority chain.

22. Government AI as a Live Test of the Theory

The transition described in this paper is no longer historical reconstruction. Government AI provides a live comparative test. The OECD's 2026 Digital Government Outlook reports that AI is used in at least one area of government in 35 of 36 OECD countries. Adoption is strongest in internal processes and public services, while policymaking and oversight remain more limited. The pattern is consistent with the administrative-compression hypothesis: adoption moves fastest where tasks are standardized, data is available, and the expected efficiency gain is legible. [S30]

The same dataset provides a direct test of governance lag. Nearly all OECD countries have an AI-in-government strategy, and 30 of 36 have at least one institution responsible for governing public-sector AI. But concrete controls are less common. Fourteen of 36 require pre-deployment risk assessments, 12 have internal review committees, and 11 conduct post-deployment audits. Only six report an open algorithm register. The strategy layer is therefore substantially more mature than the operational-control layer. [S30]

That finding should not be overstated. It does not show that 61 percent of OECD governments deploy high-risk AI without any controls. National legal systems contain procurement law, privacy law, administrative law, sector rules, security requirements, and other constraints not captured by a single AI-specific indicator. It does show something narrower and more relevant: explicit AI governance is still being translated from principles and institutions into routine operational checkpoints.

The OECD itself makes the distinction. Its 2026 analysis says guardrails are widespread but concrete, enforceable controls remain uneven. It notes that ex-ante checks can be embedded into approval and procurement workflows as checkpoints and sign-offs before deployment, while warning that ex-ante review alone is insufficient without monitoring and post-deployment audit. That is almost exactly the lifecycle problem identified by the execution-boundary model: authorization is not a one-time certification of the technology. It is a continuing relationship among purpose, deployment, action, monitoring, and accountability. [S30]

The emergence of agentic AI makes the distinction more urgent. OECD's turning-point analysis describes agentic systems as capable of placing and executing sequences of actions across government tools and workflows. It specifically identifies the need for oversight, traceable decision trails, accountability, user control, and the ability to pause, reverse, or challenge automated actions. [S32] The governance problem has therefore moved beyond whether a model produces a correct answer. It now includes whether the system may carry the answer into the world.

This is a useful falsifier because governments are not all adopting the same architecture. Some jurisdictions require mandatory ex-ante assessments. Others rely more heavily on guidance. Some use review committees. Some emphasize transparency registers. Some are adapting existing administrative-law controls rather than creating new AI-specific institutions. If the theory is correct, the successful systems should converge functionally even when their mechanisms differ: consequential capability should encounter explicit authority, scope, review, and audit before or at the point of effect.

Canada's mandatory Algorithmic Impact Assessment is one example cited by OECD. It is completed before federal deployment of an automated decision system and is used to verify compliance with binding requirements. Australia has piloted an assurance framework for ex-ante assessment. Other governments are building post-deployment audit or central review mechanisms. [S30] These are not proof that an execution boundary must be implemented in software. They are evidence that the institutional problem is increasingly being expressed as a pre-deployment and lifecycle authorization problem rather than a purely ethical aspiration.

The gap between adoption and measurement is equally significant. OECD reports that only 10 of 36 countries conduct any financial or non-financial impact measurement of AI use cases in government. A system can therefore be authorized to exist without the institution having mature evidence that it produces the intended result. [S30] That separates two questions that governance often conflates: may we deploy this capability, and is this capability actually achieving its purpose? Authority without evaluation can preserve a system long after its justification weakens.

Government AI is therefore a particularly clean test of the paper's central proposition. The capability is arriving. The governance response is arriving too, but unevenly and through multiple institutional forms. The next decade will reveal whether explicit authority and audit become native properties of machine-mediated administration or remain external documents attached to systems whose operational permissions are broader than their legal ones.

23. Reversibility, Exit, and the Problem of Infrastructure Memory

Path dependence is one of the easiest claims in this paper to overstate. Infrastructure does not make change impossible. Governments migrate systems, cancel programs, change vendors, repeal statutes, close databases, and redesign workflows. The stronger claim is unnecessary. What matters is that each capability changes the cost of reversal.

A paper process can be abandoned with relatively little technical migration. A database accumulates schemas, integrations, records, reports, training, and institutional habits. A cloud environment accumulates identities, network patterns, security controls, deployment pipelines, data gravity, observability, and specialized staff. An AI system can add model-specific evaluation, prompt or policy layers, vector stores, agent tools, workflow dependencies, and new forms of institutional expertise. The longer the capability remains operational, the more surrounding processes assume its existence.

This is infrastructure memory. The system remembers not because the technology has agency, but because organizations build around it. Budgets are allocated against it. Staff are hired for it. Policy is written assuming it. Downstream systems consume its outputs. Users learn its interface. Vendors optimize for its APIs. Removal then requires not only deleting the technology but reconstructing the organizational functions that migrated into it.

That memory explains why exit strategy belongs inside authority architecture. A government can lawfully authorize a system today and still create a future governance problem if the authorization cannot be withdrawn without unacceptable mission failure. Revocability is therefore not merely a contract term. It is part of the practical meaning of public control.

The same applies at smaller scales. An officer's access should be revocable without destroying the investigative system. An agent's delegation should expire without disabling the entire workflow. A model should be replaceable without losing the policy layer that determines what actions are permitted. An archival provider should preserve evidence without acquiring authority to alter the evidence. The more tightly capability and authority are fused, the harder it becomes to replace the capability without destabilizing governance.

This is where vendor dependency and authority dependency diverge. A vendor can be difficult to replace while the government still retains clear legal and operational control over what the system may do. Conversely, a system can be technically portable while its authorization logic remains embedded in undocumented human practice. The relevant question is not simply whether a supplier can be changed. It is whether the institution can preserve the authority boundary while changing the implementation beneath it.

A well-governed system therefore needs an exit path at several layers: data portability, policy portability, evidence portability, identity portability, and operational continuity. If a model changes, the institution should not have to rediscover what it is legally allowed to do. If a cloud provider changes, the audit history should not become unintelligible. If a contractor leaves, the government should retain enough knowledge to specify the mission and evaluate the replacement.

Infrastructure memory is not inherently bad. It is one reason complex institutions can accumulate capability rather than relearn the same task every generation. The governance problem appears when memory becomes irreversibility. A capability that cannot be safely withdrawn begins to constrain the future choices of the authority that adopted it.

That yields another testable implication. Mature authority architectures should make revocation and substitution cheaper than architectures that treat governance as an emergent property of the implementation. The better the separation between policy, evidence, identity, authorization, and execution, the more of the system can change without changing who is allowed to do what.

24. The Public-Private Boundary Is an Interface, Not a Line

The public-private distinction is often described spatially: government is on one side, industry on the other, and contracting moves a function across the line. Modern information systems make that model increasingly inaccurate. The boundary behaves more like an interface.

A commercial provider may own the physical infrastructure while the government controls the workload. A contractor may write code while a public official retains approval authority. A private communications provider may be legally compelled to assist government collection. A vendor may operate a sensor network queried by public officers. A public digital-identity system may be consumed by private services. A government may use a commercial model inside a classified environment while retaining its own data, policy, and authorization layers. The relevant question is not where the server sits. It is which actor controls which state transition.

This interface model helps avoid two opposite errors. The first is assuming that private infrastructure privatizes sovereign authority. It does not necessarily. Governments routinely use private capacity while retaining legal authority. The second is assuming that formal public authority guarantees practical control. It does not necessarily. If the public institution cannot understand, specify, audit, replace, or constrain the machinery, formal authority may be operationally thin.

The architecture should therefore be decomposed by function. Who defines policy? Who establishes identity? Who controls credentials? Who can change the enforcement mechanism? Who owns the evidence? Who can alter retention? Who can approve an exception? Who can execute the action? Who can see the audit trail? Who can revoke access? Who can replace the underlying capability? A single label such as public, private, contractor, or vendor cannot answer those questions.

The SHAMROCK example shows the early form of this interface: private communications infrastructure and public intelligence purpose. Section 702 later formalized a very different legal relationship in which electronic communication service providers can be compelled to assist authorized acquisition. Commercial cloud creates another interface: private hyperscale infrastructure carries public workloads under government security and procurement requirements. Flock creates another: private sensing and query infrastructure becomes part of local public investigative capability.

The interface perspective also clarifies accountability. When something goes wrong, institutions often discover that responsibility was distributed while capability was integrated. The vendor says the customer configured the system. The agency says the vendor designed the interface. The contractor says the government approved the requirement. The official says the system generated the recommendation. Each statement can be partly true while the overall action lacks a clear accountable owner.

Authority architecture should make that diffusion harder. Every consequential transition should have an accountable institutional principal even when many private and public components participate. The provider may supply compute. The model may supply inference. The contractor may supply integration. The officer or agency may supply purpose. The legal regime may supply authority. The system should still be able to reconstruct which authority permitted the final effect.

This is why the public-private information state is not best understood as privatization alone. It is a coupled system. Public authority increasingly operates through private capability, while private capability becomes shaped by public legal, security, and procurement requirements. Governance has to operate at the interfaces where those systems meet.

25. Evidence, Receipts, and the Difference Between Logging and Accountability

Auditability appears throughout this paper, but logging and accountability are not the same thing. A system can produce enormous volumes of logs and still fail to answer the only questions that matter after a consequential action: what was proposed, what was known, what rule applied, who or what had authority, what decision was made, what executed, and whether the executed result matched the authorized one.

Conventional observability is optimized for operational diagnosis. It tells engineers that a request failed, a service slowed, a token expired, or an exception occurred. Security logging often records authentication, access, network events, and administrative changes. Those records are valuable, but they can leave the authority chain implicit. A successful API call proves that credentials worked. It does not prove that the call was permitted for the purpose that caused it.

An authority receipt is a different evidentiary object. It preserves the relationship between a proposed action, the authority decision that governed it, and the consequence that actually followed. Its institutional value is that a later reviewer can examine the basis of the decision rather than infer permission from a successful technical operation. This is an evidentiary requirement, not a prescribed record format.

The exact data retained should be proportionate. A receipt should not become an excuse to duplicate sensitive payloads indefinitely. In many systems the evidence can be represented through identifiers, hashes, references, policy versions, scoped facts, and other minimally sufficient artifacts. The objective is reconstructability, not indiscriminate retention.

This distinction becomes critical when AI participates in the chain. Model outputs are often non-deterministic, model versions change, retrieval context changes, and external tools can change state between proposal and execution. If the institution records only the final action, it may be impossible to reconstruct why the action was considered permissible. If it records only the model transcript, it may still lack the independent authority decision.

Evidence should therefore be generated at the boundary where permission is resolved, not reconstructed later from whatever telemetry happens to remain. This is the same reason financial systems produce transaction records rather than asking investigators to infer transactions from server logs. Consequential authority deserves its own evidence object.

The record also has to survive organizational change. If a vendor leaves, an employee departs, a model is upgraded, or a policy is revised, the institution should still be able to explain a prior action according to the authority that existed at the time. Otherwise accountability becomes dependent on the continued availability of the system that is being audited.

This is one place where immutability has real value, but the claim should remain precise. An immutable record does not prove that the original decision was legitimate. It proves that the preserved record has not silently changed after the fact, assuming the integrity mechanism itself is sound. Governance still requires valid inputs, valid authority, and a correct relationship between decision and execution.

The stronger architecture therefore combines prospective and retrospective control. Prospective authorization determines whether the action may proceed. Durable evidence makes the decision reconstructable afterward. Either half alone is incomplete. Blocking without evidence creates opaque power. Logging without blocking creates excellent documentation of preventable harm.

26. Machine-Speed Administration and the Last Bottleneck

The trajectory examined in this paper can now be seen as a sequence of bottlenecks. The first bottleneck was memory: institutions could not retain or organize everything they knew. Records externalized memory. The next was retrieval: the information existed but was expensive to find. Databases compressed retrieval. The next was resolution: records could be found but relationships among people, cases, transactions, and entities remained difficult to connect. Relational systems and identifiers compressed resolution.

Networks then compressed distance. Interoperability compressed institutional boundaries. Cloud compressed infrastructure provisioning. Sensors compressed observation. Machine learning compresses the work required to classify and correlate. Generative systems compress the production and interpretation of language. Agentic systems compress the sequence between interpretation and effect.

At each stage, the previous bottleneck becomes less visible because a generation grows up treating the solved capability as infrastructure. Searchable records feel ordinary until one reconstructs the world in which a prosecutor had to know which office held the file. Cloud capacity feels ordinary until one reconstructs the procurement and deployment cycle of physical infrastructure. Instant identity resolution feels ordinary until one reconstructs the institutional labor once required to prove who someone was across systems.

Authorization is becoming the last bottleneck because it is the point the system cannot legitimately compress away. It can compress the work required to establish authority. It can make policy machine-readable. It can resolve identity instantly. It can retrieve the governing rule. It can collect required approvals electronically. It can evaluate conditions at machine-speed. What it cannot do legitimately is remove the distinction between having the capability and having permission.

This does not mean authorization must remain slow. Quite the opposite. If machine-speed administration is inevitable in some domains, authority has to become machine-speed too. A system that can evaluate a million transactions per second cannot depend on a human manually approving each ordinary, low-risk event. The policy has to be encoded so that routine authorized actions can proceed automatically while exceptions, ambiguity, conflicts, or high-consequence cases are held for additional authority.

That is selective friction in operational form. Low-risk, well-defined, clearly authorized transitions can move with almost no perceptible delay. A transition that exceeds scope, lacks required evidence, conflicts with policy, or crosses a consequential threshold should encounter friction immediately. The system becomes fast by default only inside the space that authority has already defined.

This architecture is different from post-hoc monitoring. Monitoring assumes execution and then asks whether the result was acceptable. That remains necessary for drift, fraud, unexpected behavior, and system health. But it cannot substitute for pre-execution control where the consequence is irreversible or where the cost of reversal is high. A payment sent to the wrong destination, a secret disclosed to the wrong principal, or a physical command executed in the wrong environment may not be meaningfully repaired by a later alert.

The machine-speed state therefore does not eliminate governance. It changes its location. Governance moves out of the latency of the bureaucracy and into the semantics of the transaction. The critical question becomes whether the system can establish authority before commitment and preserve evidence after commitment without forcing every legitimate action back through manual administration.

That is the final bottleneck this paper identifies. It is not a bottleneck to be removed. It is a boundary to be engineered.

27. Research Implications and a Testable Program

If the architecture-of-authority framework is useful, it should generate research questions that can be answered without accepting the framework in advance. The next step is therefore empirical rather than rhetorical.

First, researchers can measure the relationship between capability compression and misuse. A cross-jurisdictional study of networked investigative tools could compare agencies with similar access but different requirements for case numbers, supervisory approval, query logging, anomaly detection, retention, and audit frequency. The hypothesis is not that more controls always produce better outcomes. It is that explicit controls at the query or execution boundary should reduce unauthorized use without requiring the underlying capability to be dismantled.

Second, digital-government systems can be compared on the relationship between interoperability and possession. Federated exchange architectures, centralized data platforms, and hybrid systems can be evaluated for service speed, data duplication, breach impact, access accountability, and reversibility. Estonia suggests that resolution can be separated from centralized possession. That proposition should be tested against other national and subnational architectures rather than treated as universal.

Third, procurement research can distinguish vendor dependence from sector dependence. Switching costs, proprietary interfaces, data egress, staff specialization, accreditation, and mission continuity can be measured separately. A multi-cloud environment may reduce one dimension while increasing another. The relevant dependent variable is not simply the number of vendors but the institution's practical ability to preserve the mission and authority model while substituting capability.

Fourth, government AI provides a natural experiment in governance sequencing. OECD countries now differ materially in whether they require pre-deployment assessment, internal review, post-deployment audit, transparency registers, and impact measurement. [S30] Future work can ask whether jurisdictions that operationalize these controls earlier experience fewer severe incidents, faster remediation, greater public trust, or more successful scaling than jurisdictions that rely primarily on principles and voluntary guidance.

Fifth, agentic systems create a direct engineering test. Systems can be evaluated on whether delegation is explicit, scoped, expiring, and auditable; whether tool access is separated from action authority; whether proposals are bound to authorization decisions; whether authorization survives retries and asynchronous queues; whether policy conflicts fail closed or fail open; and whether the executed state can be cryptographically or otherwise reliably tied to the authorized state.

Sixth, the impunity side of the theory needs equal attention. Over-governed systems can fail by creating so many approvals, silos, referrals, and ambiguous ownership boundaries that legitimate action becomes impossible. The correct metric is therefore not the number of controls. It is the rate at which authorized necessary actions execute correctly while unauthorized actions do not. A system that prevents all abuse by preventing all action is not well governed.

That produces a measurable objective: minimize the distance between authorized action and executed action in both directions. False authorization occurs when an unauthorized action executes. False inhibition occurs when an authorized and required action fails to execute. Governance quality is the institution's ability to reduce both errors under real operating conditions.

Finally, the framework can be tested historically. If administrative compression is a genuine driver, comparable transitions should appear outside the cases selected here: tax administration, benefits delivery, immigration processing, financial regulation, healthcare, emergency management, military logistics, and courts. If the pattern disappears outside surveillance and national security, the theory should be narrowed accordingly.

This paper therefore ends its research claim before its engineering prescription. The record supports a recurring synchronization problem among capability, authority, and accountable execution. It does not establish one universal mechanism for solving it. That distinction is deliberate. The next stage is to compare mechanisms.

28. Private Infrastructure and Public Authority Across Three Eras

The public-private pattern becomes clearer when three eras are placed beside one another without claiming that one system technologically descended from the previous one. SHAMROCK, Section 702, and classified commercial cloud involve radically different statutes, technologies, oversight mechanisms, and purposes. Their common feature is narrower: sovereign information capability can operate through infrastructure the state does not exclusively own.

In SHAMROCK, private telegraph carriers provided government intelligence agencies access to international communications moving through commercial systems. The later Church investigations exposed both the scale of collection and the weakness of the surrounding oversight. The resulting governance correction was substantial precisely because private possession of infrastructure had not prevented public power from operating through it.

Section 702 represents a different legal architecture. Acquisition occurs under statutory authority, targeting and minimization procedures, judicial review, and other oversight mechanisms. Electronic communication service providers can be directed to assist lawful acquisition. The comparison to SHAMROCK is therefore architectural, not legal: private communications capability remains part of the practical path through which public surveillance authority operates. Treating the two regimes as legally equivalent would erase the governance changes that are central to the paper's argument.

Commercial cloud moves the private component deeper into the information stack. The provider is no longer merely carrying a communication that government seeks to acquire. It may provide compute, storage, identity integration, networking, managed services, AI infrastructure, and classified environments on which the public mission itself runs. The government's authority remains public, but the capability surface becomes jointly constituted by public requirements and private technical systems.

This progression is important because the governance question changes as the interface deepens. In a carrier-assistance model, the primary question may be whether government has lawful authority to compel or receive the information. In a cloud model, governance also includes architecture, accreditation, service dependencies, data rights, operational continuity, technical expertise, and exit. The state is not simply asking a private actor for information. It is building part of its institutional nervous system on commercial machinery.

That does not mean the provider acquires sovereign authority. The distinction is exactly the point. Capability can be supplied privately while authority remains public. The risk appears when the institution loses the practical ability to understand or constrain the supplied capability. A government that can state the law but cannot determine how its own system implements the law has retained formal authority while weakening operational control.

The same interface now appears in AI. A public agency may use a commercial foundation model, commercial cloud, private data tooling, and contractor integration while applying public policy and public legal authority. The model provider does not become the agency. But if the agency cannot independently determine which data reaches the model, which tools it can invoke, how outputs become actions, or how the system can be stopped, the distinction between public authority and private capability becomes difficult to enforce operationally.

The historical lesson is therefore not privatization equals abuse. It is that ownership of infrastructure and possession of authority are separate variables. Governance has to bind them together explicitly wherever public power travels through private capability.

29. Resolution Without Possession: Identity, Interoperability, and Data Minimization

One of the most important corrections produced by the comparative research is that administrative resolution does not require universal data centralization. This matters because a simplistic version of the thesis could imply that increasing state capability necessarily produces a single consolidated database. The evidence does not support that claim.

Estonia's X-tee architecture demonstrates the alternative. Independent systems remain responsible for their data while a secure exchange layer provides authenticated, authorized, logged communication among them. The architecture can answer a practical administrative question by reaching the authoritative source rather than copying every source into one central repository. [S17]

That distinction changes the privacy and security analysis. Centralization can simplify analytics and operations, but it can also enlarge breach impact, create secondary-use incentives, and make a single platform an unusually powerful point of access. Federated resolution preserves other risks - including misuse of legitimate queries and correlation across sources - but it allows possession, exchange, and authority to be governed separately.

Digital identity makes the same distinction visible. An identity system can establish that a person is the holder of a credential without exposing every fact the issuing institution possesses about that person. A service can receive an assertion that an eligibility condition is satisfied without necessarily receiving the entire underlying record. The technical details vary, but the governance principle is stable: reveal or resolve what the transaction requires rather than treating maximum data availability as the default.

This is selective friction applied to information itself. The system can be fast at proving the fact while remaining constrained about moving the underlying data. It replaces a slow human verification process with a narrow machine-verifiable exchange rather than a universal permission to copy records.

The same principle can be applied to AI systems. A model or agent often does not need unrestricted access to a source database. It may need a bounded fact, a policy-relevant attribute, a retrieval result, or a capability exposed through a governed tool. Giving the model the smallest sufficient fact surface reduces both accidental disclosure and the amount of authority that can be inferred from mere access.

This suggests a broader architecture of institutional knowledge. Possession should remain with the authoritative source where practical. Resolution should occur through authenticated exchange. Interpretation should receive the minimum sufficient context. Authorization should be evaluated independently of the component requesting the information. Execution should be limited to the authorized consequence. Evidence should record the decision without unnecessarily reproducing the sensitive payload.

The result is not a return to silos. A silo prevents legitimate exchange because systems cannot or will not communicate. Federated governance allows exchange while preserving boundaries. The difference is whether the boundary is an accidental incompatibility or an explicit rule.

That distinction is central to the paper's broader claim. The alternative to frictionless centralization is not administrative paralysis. It is high-resolution administration with intentionally preserved control surfaces.

30. Administrative Compression Beyond Security

National-security cases are analytically useful because authority, secrecy, and information sharing are unusually visible. They also create a selection problem. If administrative compression appears only in intelligence and policing, the theory may be describing surveillance rather than government.

Routine public administration provides the stronger test. Tax agencies seek pre-populated returns and automated verification because manually reconciling millions of filings is expensive. Benefits systems use digital identity and interoperable records to reduce repeated documentation and accelerate eligibility decisions. Customs systems exchange data to clear legitimate trade while identifying anomalies. Health systems use digital records and exchange standards to make information available across providers. Licensing systems move from office visits and paper certificates toward online identity, registries, and machine-verifiable status.

The political purposes differ, but the information problem is similar. The institution needs to establish a fact about a person, entity, asset, transaction, or entitlement; determine which rule applies; and produce a valid administrative result. The traditional process often requires the citizen or official to move information manually among agencies. Digital public infrastructure attempts to make the institutions perform that coordination themselves.

The European once-only principle expresses the direction clearly: people should not have to repeatedly provide information the public administration already possesses under lawful authority. The convenience gain is real. So is the architectural consequence. The institution becomes responsible for resolving data across organizational boundaries rather than forcing the citizen to serve as the integration layer.

That shift is administrative compression in its least conspiratorial form. No crisis is required. No intelligence objective is required. The incentive is service quality, cost, accuracy, and speed. A government that can resolve identity and authoritative records automatically can process more transactions with fewer manual handoffs.

The same process can reduce injustice. Manual administration can lose records, impose unequal burdens on people with less time or mobility, generate inconsistent decisions, and force citizens to understand institutional boundaries that have no relationship to their lived problem. Compression can therefore improve access and consistency.

But the authority problem remains. A faster eligibility system can deny benefits faster. A more integrated tax system can propagate an incorrect identity match across more processes. A proactive service can act on behalf of a person before the person understands what data was used. Efficiency changes the scale and latency of both correct and incorrect administration.

That is why rights-oriented digital government increasingly emphasizes user control, transparency, data governance, complaint mechanisms, and audit alongside interoperability. The objective is not to stop the state from becoming legible to itself. It is to ensure that greater legibility does not erase the individual's ability to understand, contest, or constrain the resulting action.

The broader theory therefore survives outside security only in a qualified form: administrative compression is valuable precisely because it reduces useless friction. Its legitimacy depends on whether the resulting action remains contestable, attributable, authorized, and proportionate.

31. Error Propagation: Why Faster Resolution Changes the Cost of Being Wrong

Compression does more than accelerate correct action. It accelerates error propagation. This is one reason governance requirements can increase even when the underlying decision rule has not changed.

In a fragmented paper system, an incorrect record may remain local. That is inefficient when the record is correct and another office needs it. It is protective when the record is wrong. Interoperability removes both effects. A mistaken identity, stale status, incorrect risk flag, or bad source record can move farther and influence more decisions before a human notices the problem.

AI adds another propagation layer because it can transform an upstream error into a downstream interpretation. A wrong record can become a wrong summary. A wrong summary can become a wrong risk score. A wrong score can become a recommendation. An agent can then turn the recommendation into an action. Each layer can preserve the original error while making its provenance harder to see.

This does not argue against integration or AI. It argues for provenance and contestability as first-class properties. A consequential decision should be able to identify which facts materially supported it and where those facts came from. When a fact changes, the institution should know which downstream decisions may require reconsideration.

The same issue appears in human institutions, but machine-speed changes the blast radius. A clerk can make one wrong decision at a time. A flawed rule or model can apply the same error across thousands of cases before a complaint reaches the correct office. Automation therefore changes the relationship between average accuracy and institutional harm. A small error rate can still produce a large absolute number of consequential errors at scale.

This is another reason human review cannot be the universal answer. Sampling, escalation, and appeal may be more effective than placing a ceremonial human click in every transaction. The control has to match the failure mode. High-volume deterministic actions may need automated policy enforcement and statistical monitoring. Ambiguous high-consequence cases may need independent human judgment. Identity conflicts may need hold-and-resolve behavior rather than automatic denial or automatic approval.

The architecture should also distinguish reversible from irreversible actions. A recommendation displayed to an analyst is easier to correct than a payment released, credential revoked, record published, or physical command executed. The closer the action is to irreversibility, the stronger the case for resolving ambiguity before commitment.

This produces a risk model based not only on probability of error but on propagation and reversibility. Governance should become stricter as an action becomes easier to scale, harder to reverse, more consequential to rights or safety, and more dependent on uncertain inference.

The machine-speed state therefore needs more than accurate models. It needs architectures that contain error. A system that can act everywhere should also be able to stop locally.

32. The Difference Between Permission, Delegation, and Standing

Authority is often treated as a single variable: authorized or unauthorized. Real institutions are more complicated. Permission, delegation, and standing answer different questions, and collapsing them creates precisely the kind of ambiguity that autonomous systems exploit accidentally.

Permission answers whether an action is allowed under a rule. Delegation answers whether this actor may exercise that permission on behalf of another principal. Standing answers whether the actor is the appropriate party to invoke the process at all. An employee may have permission to access a system but no standing to open a matter unrelated to assigned duties. A service account may have technical permission to modify a record but no delegation to make a policy decision. A model may be permitted to recommend an action while lacking standing to commit it.

Human organizations carry these distinctions through job roles, professional norms, signatures, supervisory chains, and institutional memory. Software tends to compress them into credentials and role bindings because those are easy to enforce. The result can be technically valid access with institutionally invalid action.

Agentic systems make delegation especially important. A user may authorize an agent to schedule a meeting, but that does not imply authority to disclose confidential documents to attendees. An agency may authorize an analytical agent to retrieve records for a case, but that does not imply authority to open a new investigation based on an inferred relationship. A planning agent may delegate a subtask to another agent, but the sub-agent should not inherit broader authority than the original delegation.

Delegation therefore needs scope and provenance. The system should know the principal, the delegating authority, the permitted action class, relevant resource boundaries, expiration, conditions, and whether further delegation is allowed. Without those constraints, delegation becomes capability multiplication.

Standing matters because even a fully authorized operation can be illegitimate when invoked by the wrong actor. Courts formalize standing doctrine in one domain, but the architectural analogue is broader. The system should be able to ask not only whether this operation is legal in the abstract but whether this requester has a legitimate relationship to the matter that justifies invoking it.

This distinction also helps with privacy. Many privacy failures are not caused by a user lacking general system access. They occur because a user accesses a legitimate system for an illegitimate purpose. Purpose limitation is therefore partly a standing problem: why is this actor entitled to make this query in this context?

A case number field, supervisory approval, or assignment check is a crude but useful attempt to encode standing. More mature systems can bind queries and actions to cases, tasks, warrants, consents, contracts, or other authority artifacts. The mechanism can vary; the principle is that access should carry context.

The architecture of authority is therefore not a bigger access-control list. It is a system for preserving the distinctions that institutions already rely on when they decide who may cause which consequence on whose behalf.

33. Designing for Disagreement, Conflict, and Hold

Most clean architecture diagrams assume that authority resolves to one answer. Real institutions contain conflicting rules. A privacy requirement can conflict with a disclosure obligation. A security policy can conflict with an availability requirement. A local procedure can conflict with a higher-order statute. Two agencies can possess overlapping jurisdiction. A delegated authority can expire while an asynchronous action remains in flight. An emergency exception can exist without being applicable to the current facts.

A machine-speed system has to represent those conflicts rather than hiding them behind a binary allow/deny decision. The most important outcome in some cases is neither execution nor rejection. It is hold: the system has enough information to know that it cannot legitimately decide with the authority currently available.

Hold is governance friction made explicit. In a paper institution, ambiguity often produces a phone call, supervisor review, legal consultation, or delay while someone finds the missing document. Those mechanisms are slow but conceptually important. A digital system that removes them needs an equivalent state in which execution pauses without converting uncertainty into permission.

This matters because fail-open and fail-closed are both incomplete as universal rules. Fail-open can turn missing governance infrastructure into unauthorized action. Fail-closed can turn a temporary technical failure into denial of legitimate service, potentially at enormous scale. The correct behavior depends on consequence, reversibility, statutory obligation, emergency conditions, and the nature of the missing authority.

A mature authority system therefore needs precedence. Higher-order rules should be distinguishable from lower-order preferences. Mandatory prohibitions should not be overridden by convenience. Emergency authority should be explicit, scoped, time-limited, and auditable rather than inferred from urgency. Conflicts should produce a deterministic escalation path rather than whichever component executes first.

Transformation is another legitimate outcome. An action can be authorized only after sensitive fields are removed, the destination is narrowed, the amount is reduced, the scope is limited, or another condition is applied. Governance does not always choose between doing the requested thing and doing nothing. It can change the proposed action into the action authority actually permits.

That distinction is useful in AI systems because model proposals are often intentionally broad. The model can optimize for usefulness while the authority layer narrows the proposal to institutional limits. Creativity remains in the proposal path. Permission remains in the commitment path.

Disagreement also has to remain visible in evidence. If two policies conflict and a higher-order rule resolves the conflict, the receipt should preserve enough information to reconstruct that resolution. Otherwise the final action may look arbitrary even when the system behaved correctly.

The presence of a hold state is therefore a marker of institutional maturity. It acknowledges that uncertainty exists and refuses to let uncertainty silently become authority. A system that cannot say "not yet" is forced to treat every unresolved case as either permission or prohibition. Both are dangerous simplifications.

34. The Architecture of Authority as Institutional Memory

The final implication of the research is that authority itself has to become durable institutional memory. Modern systems are very good at remembering data and increasingly good at remembering behavior. They are less consistent at remembering why an action was legitimate at the moment it occurred.

Human institutions traditionally preserve authority through statutes, regulations, orders, contracts, policies, minutes, signatures, case files, and professional roles. The problem is not a lack of authority artifacts. It is that operational systems often consume only a thin projection of them: a role, a token, a Boolean permission, or a configuration flag.

As execution becomes automated, that thin projection becomes risky. The system can continue executing long after the institutional context that justified the permission has changed. A contract expires. A user changes roles. A policy is superseded. A consent is withdrawn. A case closes. A warrant reaches its limit. A temporary emergency authority ends. If the execution layer sees only a persistent credential, the historical authority can outlive itself technically.

Authority therefore needs lifecycle semantics. It must be created by a valid source, scoped, versioned, delegated where permitted, amended, suspended, expired, revoked, and retained as evidence after it is no longer active. Those are familiar properties in legal and administrative systems. The machine layer has to preserve them rather than flattening authority into access.

This also changes how policy updates should work. A new rule should not merely change future behavior; the institution should know which version governed a prior action. If a later investigation asks why the system acted on Tuesday, the answer cannot depend on today's policy configuration. Reconstructability requires historical policy state.

The same principle applies to facts. Authority may depend on facts that were true at the time and later changed. A person was employed, a license was valid, a threshold was exceeded, a threat condition existed, or an approval was active. The receipt should establish the decision context without pretending that current state is identical to historical state.

This is why authority architecture is also evidence architecture. The institution is not merely preventing unauthorized actions. It is preserving the institutional memory necessary to explain authorized ones.

That memory becomes more important as humans leave the immediate execution path. A person can later testify to why a decision was made, although memory is imperfect. An autonomous workflow has no independent institutional memory beyond what the system records. If the architecture does not preserve the authority chain, the reason for the action disappears even if every technical log survives.

The machine-speed institution therefore needs to remember more than what happened. It needs to remember what made what happened permissible.

35. The Practical Standard: Govern the Commit, Not the Conversation

A large share of contemporary AI governance is organized around the conversation: what the model said, whether the output was biased, whether a prohibited topic appeared, whether a response contained unsafe content, and whether the user received a disclosure. Those questions matter. They are not the whole problem once the model can act.

An agent can produce a perfectly acceptable sentence and still perform an unauthorized operation. It can correctly summarize a policy and then call the wrong tool. It can identify the correct customer and then disclose the customer's data to the wrong destination. It can propose a legitimate payment and execute it from an account outside its delegated scope. Content safety and execution authority overlap, but neither subsumes the other.

The practical standard is therefore to govern the commit. The conversation, plan, model output, or workflow proposal can remain flexible up to the point where the system is about to create a consequential external state. At that point the architecture should resolve the action against authority that is independent of the component proposing it.

This is analogous to transaction processing. A database can perform complex work before commit, but commit is the boundary at which the change becomes durable. An authority architecture treats consequential external action similarly. The system may explore, retrieve, reason, simulate, and propose, but the protected transition does not become real until the authorization condition is satisfied.

That approach has an important operational advantage. It avoids trying to predict and prohibit every undesirable internal reasoning path. The institution can instead define the smaller set of consequential actions it cares about protecting: release funds, expose protected data, alter permissions, submit a filing, send an external communication, invoke a physical actuator, modify a production system, or create another durable institutional effect.

The policy surface then becomes tractable. The institution does not need to understand every token the model produced to know that only a finance principal with the correct delegation may release a payment above a threshold. It does not need to classify every possible plan to know that a production credential cannot be used outside an approved deployment. It does not need to trust the model to remember a retention rule if the storage boundary enforces the rule independently.

This does not eliminate upstream governance. Data quality, model evaluation, prompt security, privacy, red teaming, content controls, and monitoring remain necessary. The point is architectural separation. Upstream controls improve the quality and safety of proposals. The execution boundary controls whether a proposal becomes consequence.

The separation also improves failure analysis. If the model generated a bad proposal but the authority layer blocked it, the model failed while the governed system succeeded. If the model generated a good proposal but the authority layer incorrectly denied it, the authorization system failed. If both approved an action that later proved harmful because the governing policy itself was wrong, the institutional policy failed. Those are different failure classes and should not be collapsed into "the AI failed."

This is especially important for accountability. Without separation, every incident becomes a debate over model behavior. With separation, investigators can ask a more disciplined sequence: what was proposed; what authority was presented; what policy applied; what verdict was produced; what executed; and whether the execution matched the verdict. The chain localizes responsibility.

The same architecture can preserve speed. Routine actions whose authority is clear can be authorized automatically. Ambiguous actions can be transformed or held. Prohibited actions can be denied. High-consequence exceptions can be escalated. The system does not require a human to watch every action. It requires every consequential action to encounter a valid authority decision.

That is the practical meaning of selective friction. The friction is not placed around intelligence, language, or computation generally. It is placed around commitment when commitment changes something the institution is responsible for.

This also explains why monitoring alone is insufficient. Monitoring is retrospective by default. It observes that a state transition occurred and attempts to determine whether the transition was acceptable. For reversible, low-consequence activity, that may be enough. For irreversible or high-consequence actions, the institution needs prospective control as well. The audit trail should document a decision that existed before the effect, not merely an alarm generated after it.

The standard can therefore be stated without reference to any particular vendor or product: consequential state transitions should be governed at the commit boundary by authority that the proposing component cannot manufacture for itself, and the decision should leave evidence sufficient to reconstruct the relationship among proposal, authority, verdict, and execution.

That is not a complete theory of AI safety. It is a narrower theory of institutional authority. It says where one class of governance must live once software can move directly from inference to action.

36. The Authority Gradient: Not Every Action Deserves the Same Friction

A theory of explicit authority fails if it treats every state transition as equally consequential. Reading a public record, drafting an internal note, changing a benefit determination, releasing funds, altering production permissions, and activating a physical system do not create the same exposure. If the architecture imposes the same authorization burden on each, governance becomes indiscriminate overhead rather than institutional control.

The useful variable is consequence. Consequence can arise from irreversibility, scale, rights impact, financial value, safety effect, confidentiality, jurisdiction, or the number of downstream systems that inherit the result. The more consequential the transition, the stronger the institution's reason to require explicit evidence of authority before commitment. Low-consequence and readily reversible actions can often proceed under standing policy. High-consequence or difficult-to-reverse actions should encounter stronger checks, narrower delegation, or a hold state.

This creates an authority gradient rather than a universal gate. The architecture can remain fast because routine authority is pre-resolved. A valid role, current delegation, bounded purpose, and low consequence can produce automatic authorization. The system becomes slower only where the institution has decided that uncertainty is more expensive than delay.

The gradient also makes proportionality an engineering property. A policy engine should not ask only whether an actor has some permission related to an object. It should ask whether the strength and scope of the presented authority are sufficient for the consequence being proposed. The same principal may be allowed to read, recommend, draft, or stage an action while lacking authority to commit it.

This is how selective friction avoids becoming a disguised argument for bureaucracy. It does not maximize review. It aligns the cost of review with the cost of being wrong.

37. Delegation Chains and the Problem of Authority Laundering

Agentic systems make delegation chains operationally important. A human principal may authorize one service, that service may invoke an agent, the agent may call another agent, and a downstream tool may finally create the consequential state transition. Every component can be authenticated while the authority of the original principal becomes progressively less visible.

That creates a risk this paper calls authority laundering. The term does not imply criminal intent. It describes a structural failure in which an action appears authorized because every local hop was technically permitted, even though no component can demonstrate that the final consequence remained inside the purpose and scope of the originating authority.

Conventional service credentials make this easy. A downstream service account may possess broad standing permissions because the infrastructure was designed for trusted application-to-application calls. Once autonomous components can select tools and compose workflows dynamically, possession of that credential can become a substitute for delegated authority. The system knows that the service may call the API. It may not know whether this particular agent may cause this particular effect for this particular principal now.

A governed delegation chain therefore has to preserve more than identity. It needs provenance: who delegated, what was delegated, for which purpose, under which limits, until when, and whether redelegation was permitted. A downstream component should receive no more authority than the chain can validly carry. If an upstream delegation expires or is revoked, downstream authority should fail with it rather than survive as an orphaned credential.

This principle has a direct institutional analogue. Public authority is routinely delegated, but delegation does not erase the source of authority. Agencies act under statutes; officials act under offices; employees act under roles; contractors act under instruments that define scope. Machine delegation should preserve the same logic instead of converting a chain of limited grants into a final actor with ambient power.

The rule is simple: delegation may transmit authority; it may not manufacture it.

38. Temporal Authority: Permission Has a Clock

Authority is not only scoped by actor and object. It is temporal. Warrants expire. emergency authorities sunset. contracts terminate. employment changes. approvals are superseded. consent is withdrawn. incident-response privileges end when the incident closes. A system that treats yesterday's authority as a permanent capability silently converts temporary permission into durable power.

This temporal problem becomes more severe in autonomous systems because execution can be delayed, retried, queued, or resumed. An action may have been valid when proposed and invalid when the worker eventually commits it. A long-running workflow can cross a policy change, a revocation, or a jurisdictional boundary while retaining the same technical credentials.

The execution boundary should therefore evaluate authority at commitment time, not merely at planning time. Earlier approval can remain relevant evidence, but the system should be able to determine whether the approval is still valid for the current action and context. Admission to a queue is not itself a determination of continuing authority. A queue preserves intent to attempt an action; it should not freeze authority in time.

Temporal semantics also improve emergency governance. Emergency access is often justified precisely because normal constraints are too slow for an exceptional condition. The architecture should make the exception visible, narrow, and self-expiring. Otherwise the temporary path becomes a permanent shadow permission that survives after the reason for it disappears.

The broader historical ratchet in this paper has a machine-scale equivalent. Institutions tend to retain capabilities after crises because infrastructure has memory. Individual systems can reproduce the same ratchet in miniature when temporary permissions are technically easier to grant than to revoke. Explicit expiry and revocation are therefore not housekeeping. They are anti-ratchet controls.

39. The Minimum Sufficient Authority Principle

The familiar security principle of least privilege limits the resources or operations available to a principal. That remains necessary, but agentic execution requires a related principle: minimum sufficient authority. The system should carry enough authority to complete the legitimate action and no more authority than the action requires.

The distinction matters because privilege and authority are not identical. A service can have technical privilege to write a database while lacking institutional authority to change a specific person's status. Conversely, an institution can have legal authority to produce an outcome while a particular service lacks the technical privilege needed to implement it. Governed execution requires both conditions to align.

Minimum sufficient authority also limits blast radius. If a workflow needs to send one approved notice, it should not inherit unrestricted messaging authority. If an agent needs to retrieve a defined record set, it should not receive a credential capable of reading an entire tenant. If a payment is authorized up to a defined amount and recipient, the execution artifact should not become a general-purpose payment credential.

This is not merely a security optimization. It preserves the meaning of authorization. The more ambient capability a component possesses, the harder it becomes to infer whether an observed action followed from legitimate authority or merely from available privilege. Narrow grants make the relationship between permission and execution legible.

Minimum sufficient authority therefore connects cybersecurity, administrative law, and systems architecture without collapsing them into one discipline. Security limits what compromise can reach. Institutional authority limits what legitimate actors may cause. The execution boundary has to satisfy both.

40. Revocation, Reversal, and the Right to Stop the Machine

A mature authority architecture needs a theory of stopping, not only starting. OECD's 2026 analysis of agentic AI in government explicitly identifies the ability to pause, reverse, or challenge automated actions as part of responsible deployment. It also argues that governments need to govern the action space granted to an agent, the approval points in the workflow, and the mechanisms for oversight and redress. [S32]

Pause, revoke, and reverse are different operations. Pause prevents the next transition while preserving the workflow. Revocation removes an authority that would otherwise permit future transitions. Reversal attempts to restore a prior state after an authorized or unauthorized action has already committed. Some actions support all three. Others are physically, legally, or socially irreversible.

That difference should affect authorization before execution. An action that can be cheaply reversed may tolerate more automation and less pre-commit friction than an action that cannot be recalled. Sending an internal draft is different from publishing a legal notice. Staging infrastructure is different from deleting evidence. Preparing a payment is different from settling it. The architecture should know which side of the reversible boundary it is approaching.

Redress also matters because authority can be valid and still produce a wrong result. A lawful process can operate on bad facts. A policy can be unjust or outdated. A model can misclassify. An authorized human can make a mistake. Governance therefore cannot end with proof that permission existed. Institutions also need mechanisms through which affected people can challenge the factual or legal basis of consequential action.

The Council of Europe's AI Framework Convention makes that distinction explicit at the legal level. It calls for relevant information about AI systems and their use to be documented and made available sufficiently for affected people to challenge decisions or the use of the system, alongside complaints and procedural safeguards. [S22] The engineering analogue is that the evidence chain must be intelligible enough to support review, not merely tamper-evident enough to prove that a machine ran.

The right to stop the machine is therefore not a single kill switch. It is a set of lifecycle controls that preserve institutional authority before, during, and after automated execution.

41. Authority Debt: The Governance Equivalent of Technical Debt

Digital systems accumulate technical debt when expedient implementation choices create future maintenance costs. Institutions can accumulate authority debt in the same way. Authority debt appears when a system is allowed to operate because everyone currently understands an informal exception, a broad credential, a temporary workaround, or an undocumented dependency, but the architecture cannot later reconstruct why that operation was legitimate.

The debt often remains invisible while the original team is present. People supply missing context from memory. They know which account is shared, which service is trusted, which approval happened in a meeting, and which policy exception was intended to be temporary. Automation removes those people from the immediate path without automatically encoding the context they were carrying.

Scale then compounds the debt. A one-off exception becomes a reusable integration. The integration becomes a platform assumption. New agents inherit the platform assumption as if it were a formal grant. Eventually the institution cannot remove the permission without breaking workflows that were built around it. What began as a governance shortcut becomes infrastructure memory.

Authority debt can therefore be measured indirectly. Indicators include long-lived broad credentials, undocumented delegation, inability to identify the source of a permission, policy rules with no owner, emergency access that does not expire, actions whose authorizing evidence cannot be reconstructed, and workflows that cannot be disabled without unacceptable mission impact.

The concept also explains why governance retrofits are expensive. Adding a policy engine after the fact is not enough if the institution cannot identify where authority entered the system or which downstream dependencies rely on ambient access. Paying down authority debt requires recovering the missing model of who may cause what, under which conditions, and then moving those conditions into enforceable boundaries.

This is another reason governance should be designed before machine autonomy scales. The cheapest moment to encode authority is before an informal permission becomes an ecosystem dependency.

42. Measuring Governed Execution

If governed execution is to be more than a design philosophy, it needs observable performance measures. Traditional AI evaluation focuses on model quality: accuracy, hallucination, bias, robustness, latency, or task completion. An authority architecture requires a second evaluation surface concerned with whether consequences matched permission.

At minimum, an institution can measure authorization coverage: the percentage of consequential transitions that encounter an explicit authority decision. It can measure decision provenance: the percentage of executed actions for which the governing policy version, authority source, actor identity, declared purpose, and relevant evidence can be reconstructed. It can measure enforcement fidelity: whether execution matched the verdict rather than bypassing or mutating it downstream.

Other measures address failure. False authorization is the rate at which an action executes despite lacking valid authority. False inhibition is the rate at which an action that was validly authorized and institutionally due for execution or authoritative disposition is blocked, abandoned, or left unresolved without a legitimate basis. Unauthorized bypass measures consequential transitions that occur outside the governed path. Stale-authority execution measures actions committed after the relevant grant expired or was revoked. Mean time to revoke measures how quickly an institution can terminate authority across dependent systems after a change in role, policy, incident state, or consent.

The hold state creates another metric: unresolved-authority rate. A high rate may indicate that policies are ambiguous, delegation is poorly modeled, or the system is attempting actions outside its intended operating envelope. A zero rate is not automatically good; it may indicate that uncertainty is being silently collapsed into allow or deny.

Finally, governance has to be evaluated against mission performance. A system that blocks every action has perfect prevention and zero utility. The relevant objective is not maximal denial. It is high authorization fidelity with acceptable latency and low interference with legitimate work. That is why the target state remains authorized action equals executed action rather than execution approaches zero.

These metrics make the theory falsifiable in deployed systems. If explicit execution governance does not reduce unauthorized transitions, improve reconstructability, or preserve legitimate throughput compared with monitoring-only controls, the architectural claim should be revised.

42.1 A multi-objective loss function

A serious comparison among governance architectures needs an objective function that does not reward a system merely for blocking more actions. One research form is:

$$L_G = w_{FA} C_{FA} + w_{FI} C_{FI} + w_{LAT} C_{LAT} + w_{REC} C_{REC}$$

C_{FA} is the consequence-weighted cost of false authorization. C_{FI} is the consequence-weighted cost of false inhibition. C_{LAT} is legitimate-action latency attributable to governance. C_{REC} is the residual cost of unreconstructability, failed remediation, or inability to explain and correct the outcome. The weights are domain-specific and should be specified by the governing institution or independent study design before comparative testing; they should not be tuned after results are known to favor a preferred architecture.

The formulation deliberately does not collapse legitimacy into a single universal score. A benefits system, intelligence workflow, cloud control plane, medical device, and physical robot should not share the same consequence weights. Its narrower purpose is to force an evaluation to price both wrongful execution and wrongful inhibition while keeping latency and reconstructability visible.

H4 can then be tested more precisely. A selective-friction architecture should outperform alternatives only if it lowers the domain-defined governance loss relative to monitoring-only control, generalized approval friction, or another credible baseline. If it reduces false authorization only by producing unacceptable false inhibition or latency, it has not won the comparison.

> **A governance architecture does not succeed by preventing action. It succeeds by aligning action and non-action with valid authority at acceptable cost.**

43. A Research Agenda for the Machine-Speed State

The paper's historical argument generates a research program rather than a final answer. The first question is comparative: do institutions with explicit pre-execution authority controls experience fewer consequential misuse events than institutions with equivalent capability governed primarily through policy and retrospective audit? The answer requires matched cases, not anecdotes.

The second question concerns latency. How much governance friction is actually necessary? If machine-readable authority can resolve routine cases in milliseconds, then the traditional tradeoff between control and speed may be partly artificial. Research should distinguish the latency created by authority resolution from the latency created by poorly designed bureaucracy.

The third question concerns reversibility. Systems should be compared by whether the consequence can be paused, rolled back, compensated, or challenged. The hypothesis is that irreversible actions require stronger ex-ante authority while reversible actions can tolerate greater reliance on monitoring and ex-post correction.

The fourth question concerns delegation depth. As the number of autonomous hops between principal and executor increases, does authority provenance degrade? If so, at what depth do conventional identity and access-management controls stop providing an adequate account of purpose and delegated scope?

The fifth question concerns institutional capacity. OECD's 2026 findings show that AI strategy is widespread while operational governance and impact measurement remain less mature. [S30] A useful comparative program would test whether retained public-sector technical expertise predicts better procurement, stronger reversibility, and more effective oversight after controlling for spending and system scale.

The sixth question concerns public legitimacy. Does reconstructable authority improve the ability of affected people, auditors, courts, and oversight bodies to understand and challenge machine-mediated action? Evidence architecture matters only if the evidence can actually be used by the institutions responsible for accountability.

The final question concerns the boundary itself. Some systems may prove that pre-execution control is unnecessary for broad classes of low-consequence activity. Others may show that human approval adds little value. The theory should not defend a preferred mechanism against those results. Its durable claim is only that consequence must remain conditional on valid authority. The mechanism is an empirical design question.

A serious architecture of authority should therefore be willing to discover that some proposed controls are unnecessary, that others are insufficient, and that different consequences require different boundaries. The claim ceiling remains in force even when the subject is the framework itself.

PART VI

CONTROL CASES, MEASUREMENT, AND DUE PROCESS

44. Comparative Authority I: When High Capability Does Not Produce Abuse

A theory about disappearing friction is weak if it can explain only failure. The stronger test is a system that resolves information quickly, crosses institutional boundaries, and still preserves legitimate control. Such cases matter because they separate capability from outcome. If high capability necessarily produced abuse, the architecture would be deterministic. The evidence does not support that stronger proposition.

Estonia remains the clearest control case in this paper because its interoperability model does not require every participating institution to surrender its data to one central repository. The exchange layer makes authorized queries possible across independently controlled systems. The architectural point is not that Estonia has solved every governance problem. It is that high resolution can coexist with distributed possession, authenticated exchange, logging, and institutional ownership of source data.

That control narrows the thesis. Centralization is not a necessary consequence of interoperability. Resolution is not possession. A system can reduce search and coordination friction while retaining boundaries around who controls the underlying record and who may request it. The relevant governance variable is therefore not simply how much information can move. It is how authority is represented at the moment the information is requested, combined, interpreted, or used.

The same logic appears in contemporary government AI guidance. OECD now treats agentic AI as a shift from systems that answer toward systems that act. Its 2026 analysis says governments need to govern the action space granted to an agent, the approval points in a workflow, and the mechanisms for oversight and redress. It also emphasizes traceable decision trails and the ability to pause, reverse, or challenge automated actions. [S32] Those are not proof of this paper's model. They are independent evidence that the operational boundary is becoming visible to institutions deploying the capability.

44.1 The control-case requirement

A useful control case should satisfy at least four conditions. First, the system must possess substantial information-processing capability. Second, it must actually reduce administrative friction rather than merely describe a future design. Third, it must preserve identifiable authority or accountability boundaries. Fourth, the evidence must allow the researcher to distinguish architectural controls from favorable outcomes caused by unrelated factors.

This requirement prevents a common analytical error. A successful institution cannot be treated as proof that its architecture caused the success. The claim ceiling permits a narrower inference: if a high-capability system operates with explicit boundaries, then the existence of that system falsifies the proposition that capability and governance are inherently opposed. The stronger causal question requires comparative outcome data.

> **High capability is not the failure condition. Unresolved or unenforced authority is the variable under examination.**

45. Comparative Authority II: When Low Technology Still Produces Abuse

The inverse control matters just as much. Abuse predates databases, cloud computing, artificial intelligence, and machine-speed administration. Paper institutions can discriminate. Human officials can exceed jurisdiction. Manual processes can conceal decisions, destroy evidence, deny due process, and exercise arbitrary discretion. Technology is therefore not the origin of institutional abuse.

This is not a concession at the edge of the theory. It is a necessary boundary. The paper is about the relationship between capability and authority, not a claim that automation creates illegitimate power from nothing. A low-technology institution can possess broad legal power, weak accountability, and harmful incentives. A high-technology institution can possess narrow authority and strong controls. Technology changes the speed, scale, observability, reversibility, and propagation of action; it does not determine legitimacy by itself.

The difference becomes important when evaluating the disappearing human bottleneck. Humans were never uniformly safeguards. They can introduce bias, corruption, fatigue, delay, favoritism, inconsistency, and simple error. Some administrative friction is itself a mechanism of exclusion. A person forced to visit an office repeatedly, produce redundant documentation, or wait months for a routine determination is not necessarily being protected by governance.

The surviving claim is narrower. Human and organizational limits historically constrained the rate at which institutions could resolve and execute decisions. Some of those constraints incidentally limited the scale of both legitimate and illegitimate action. Removing them can improve justice and service while also increasing the consequences of a defective authority model. The governance problem is therefore selective: preserve the boundary, not the delay.

> **A paper process can be abusive. A machine process can be legitimate. The relevant question is whether the consequential act is authorized, proportional, reviewable, and attributable.**

46. The Authority Graph: From Roles to Provenance

Role-based access control answers a useful but incomplete question: what may a member of this role generally do? Agentic execution introduces a harder question: why may this particular actor perform this particular action, on this object, for this purpose, at this time, under this delegation, using these facts? That question is better represented as a graph than as a static permission list.

An authority graph begins with principals and sources of authority. A statute, regulation, contract, consent, court order, policy, appointment, credential, or delegated mandate can create standing. That standing can be scoped to objects, purposes, jurisdictions, time windows, consequence levels, and permitted forms of delegation. Each subsequent hop should preserve enough provenance to show that the downstream action remains inside the originating authority.

The graph matters because technical identity and institutional standing are not the same object. A service account can authenticate successfully while lacking authority for the proposed purpose. An employee can hold a valid credential after changing roles. An agent can inherit a token whose permissions exceed the task assigned by the principal. Authentication establishes who or what is presenting a credential. Authority establishes whether the resulting state transition is institutionally permissible.

This distinction also exposes authority laundering. If every hop checks only whether the previous hop was technically authenticated, broad application credentials can turn a narrow human request into a much broader machine capability. The final executor sees a valid caller but cannot reconstruct the mandate. The authority graph therefore has to survive delegation rather than being replaced by each intermediate component's standing permissions.

46.1 Authority provenance as a first-class object

Provenance should answer at least six questions: what source created the authority; who currently holds it; what action does it permit; what conditions limit it; how was it delegated; and what event terminates it. A receipt that records only the final identity or policy result is insufficient when the legitimacy of the action depends on a chain of prior grants.

This suggests a useful separation between identity provenance, data provenance, and authority provenance. Identity provenance explains the actor. Data provenance explains the factual inputs. Authority provenance explains the permission. Consequential execution can depend on all three, but none substitutes for the others.

> **Authenticated is not authorized. Authorized is not infinitely delegable. Delegated is not permanent.**

47. Machine-Speed Due Process

Due process is usually described as a legal doctrine rather than a systems property. Yet automated administration forces a systems question: what becomes of notice, explanation, contestability, hearing, correction, and remedy when the underlying decision path can complete before a person knows that a decision exists?

The answer cannot simply be to insert a human into every workflow. That would recreate generalized administrative friction and would fail at the scale that motivated automation in the first place. The architectural problem is to preserve the functions of due process while allowing routine computation to remain fast.

Some protections must occur before execution. If an action is irreversible, physically dangerous, liberty-restricting, or capable of immediately exposing protected information, retrospective review may be inadequate. Other protections can occur after execution if the action is readily reversible and the affected person can obtain meaningful correction. The location of the safeguard should therefore track consequence and reversibility rather than a universal preference for ex-ante or ex-post review.

Machine-speed due process also requires an intelligible record. A person cannot challenge a decision if the institution cannot reconstruct which rule, authority, facts, model output, and delegation produced it. Explainability at the model layer may help, but it is not enough. A perfectly interpretable model can still execute under the wrong authority. Conversely, a complex model can participate in a legitimate process if the consequential transition is independently authorized and the record preserves the basis for review.

47.1 Contestability is a lifecycle, not a button

A complaint mechanism added to a user interface is not equivalent to contestability. Meaningful challenge requires that the institution can locate the action, identify the governing authority, preserve the relevant version of policy and facts, suspend or reverse the consequence where appropriate, and route the challenge to an actor with power to grant a remedy.

The Council of Europe AI framework moves in this direction by pairing risk and impact assessment with accountability, oversight, documentation, procedural safeguards, and mechanisms through which affected people can challenge AI-mediated decisions. [S22] The important architectural observation is that redress depends on retained evidence and retained authority. A right to complain is thin if the institution cannot reconstruct or alter the machine-mediated state transition.

> **Machine-speed execution does not eliminate due process. It makes the location, timing, and evidence of due process an architectural question.**

48. Authority Under Uncertainty: Facts, Models, and Confidence

Many consequential actions depend on facts that are incomplete, probabilistic, stale, contested, or produced by models. An authority system therefore cannot assume that every input to a policy decision is a stable truth. The architecture needs to distinguish authority from the epistemic quality of the facts on which authority is conditioned.

This distinction prevents a dangerous collapse. A model may assign high confidence to an inference, but confidence does not create jurisdiction or permission. A sensor may report an event, but observation does not determine the lawful response. A risk score may cross a threshold, but the threshold does not by itself establish standing to execute every downstream action associated with that risk.

Uncertainty can instead alter the permitted transition. A low-confidence fact may authorize collection of additional evidence but not denial of a benefit. A conflicting identity resolution may justify a hold rather than execution. A stale credential may require revalidation. A model result may support a recommendation while remaining insufficient as the sole factual basis for an irreversible action.

This is where transformation and hold semantics become more than implementation details. They provide intermediate governance states between unrestricted execution and total denial. The architecture can narrow scope, request another fact, remove a sensitive field, require a second authority source, or defer commitment until the uncertainty is resolved.

> **Confidence can change what evidence is sufficient. It cannot manufacture authority.**

49. Cascading Authority: When One Authorized Action Becomes Many

Modern systems rarely end at one state transition. A record update can trigger notifications, eligibility recalculations, fraud checks, downstream synchronizations, payments, access changes, or new model inferences. The authority question therefore propagates. Was the original authorization sufficient for every consequence produced by the event?

Event-driven architectures make this problem easy to miss because each downstream component may be functioning exactly as designed. The first service writes a lawful change. A message bus publishes the event. Subscribers react automatically. No single component appears to exceed its local permissions. Yet the combined workflow may produce a consequence that the originating principal never authorized.

This creates a distinction between direct authority and derivative authority. Direct authority permits the initiating transition. Derivative authority permits specified consequences of that transition. The latter should not be assumed to be unlimited. A lawful update to a customer address does not necessarily authorize disclosure to every analytics system. A lawful case classification does not automatically authorize every enforcement action that could be triggered by the classification.

The authority graph therefore needs consequence propagation rules. Downstream actions should either inherit a bounded derivative mandate or resolve their own authority before commitment. The stronger the consequence, the less defensible it becomes to rely on implicit inheritance.

49.1 Blast radius is a governance variable

Traditional risk engineering asks how far a technical failure can propagate. Authority engineering should ask the same question about permission. If one mistaken grant can authorize hundreds of downstream actions, the authority blast radius is large even if every service remains technically healthy. Systems should therefore measure not only privilege breadth but the number and consequence of transitions reachable from a grant.

> **One authorized event is not a blank check for every event it can technically trigger.**

50. Emergency Authority and the Sunset Problem

Emergency governance is one of the strongest tests of the framework because urgency can legitimately change the acceptable balance between speed and review. The mistake is to treat emergency authority as either inherently illegitimate or automatically transferable into normal operations.

An emergency grant should carry machine-readable scope, triggering condition, issuer, consequence ceiling, delegation rule, start time, expiry, renewal condition, and post-event review requirement. Without those properties, temporary authority becomes difficult to distinguish from permanent capability once the immediate crisis passes.

This is the technical form of the earlier ratchet problem. Infrastructure remembers what policy intended to forget. Accounts remain. integrations remain. data remains. workflows remain. vendor contracts remain. personnel learn the capability. A legal sunset can therefore occur while the operational capability persists.

A genuine sunset requires more than an expiration date in a policy document. Credentials must expire or narrow. exceptional routes must close. retained data must follow the governing disposition rule. downstream delegations must terminate. systems should be able to prove that the emergency authority is no longer executable.

This also provides a falsifier for the fear-capability ratchet. If institutions repeatedly demonstrate that emergency capabilities are technically dismantled or rendered unusable when the authority expires, the ratchet is weaker than this paper predicts. The empirical question is therefore not whether a statute contains a sunset clause. It is whether the capability state actually returns.

> **Temporary authority that cannot expire operationally is not temporary authority.**

51. Procurement as Authority Architecture

Procurement is usually treated as upstream from operational governance. In information infrastructure, that separation is increasingly artificial. Contract terms determine access to data, audit rights, portability, intellectual property, incident reporting, subcontracting, model changes, retention, termination, and the government's ability to inspect or reconstruct system behavior. Those terms shape the authority boundary before the system is deployed.

GAO's 2026 review of federal AI acquisitions is useful precisely because it does not frame vendor involvement as inherently problematic. It documents multiple acquisition models and identifies practical tradeoffs. At the same time, the report summarizes federal guidance calling for contract terms addressing government data, intellectual property, privacy, and vendor lock-in; testing throughout the acquisition lifecycle; authorization to operate before deployment; post-award oversight; and implementation of data and rights terms when contracts end. [S35]

That sequence maps directly onto lifecycle authority. Before award, the institution defines what capability it is acquiring and what rights it must retain. Before deployment, an appropriate official authorizes operation. During operation, oversight tests whether actual behavior remains within the contract and policy. At termination, rights and data have to move or expire in a controlled way. Procurement therefore determines whether later governance is technically possible.

51.1 The epistemic capacity problem

A government can formally retain authority while losing the practical capacity to exercise it. If only the vendor understands the architecture, cost model, training data dependencies, integration behavior, or failure modes, the public institution may be unable to evaluate alternatives or verify compliance. GAO continues to identify technical-expertise gaps in AI acquisition, including difficulty obtaining specialists able to evaluate contractor proposals. [S35]

The claim ceiling remains important. Expertise asymmetry is not proof of capture. Specialized vendors can rationally possess deeper product knowledge than customers. The governance issue appears when the asymmetry prevents the institution from specifying requirements, evaluating performance, enforcing rights, or exiting the relationship on credible terms.

Procurement does not merely buy capability. It preconfigures who will be able to govern that capability later.

52. Machine-Readable Authority and Continuous Authorization

Static authorization artifacts fit poorly with continuously changing systems. A system can be authorized at deployment and become materially different through model updates, configuration changes, new tools, new data sources, expanded integrations, or altered downstream permissions. The authority state therefore has to evolve with the operational state.

This is already visible outside AI. GAO's 2026 reporting on the FedRAMP 20x program describes a planned transition toward machine-readable authorization data for initial and continuing cloud authorization. [S36] The specific program is not evidence for every proposition in this paper, but it demonstrates that authorization itself is becoming a computational object rather than only a document reviewed at long intervals.

Continuous authorization does not mean continuous permission. It means the system can repeatedly establish whether the conditions supporting permission still hold. A material change can narrow, suspend, or invalidate authority. A stable low-risk system can continue without ritual reapproval when the evidence shows that the authorized conditions remain intact.

For agentic systems, this becomes especially important because the action surface can change without replacing the base model. Adding a payment tool, production credential, external database, or physical actuator changes what the system can do. Governance should therefore bind to capability configuration and consequence, not merely to a model name or version.

Authorization is a state, not a ceremony.

53. The Institutional Kill Path: Pause, Revoke, Recover

Every serious authority architecture needs a path for stopping execution. The existence of a policy rule is not enough if the institution cannot revoke credentials, halt delegated agents, isolate a tool, suspend a workflow, or prevent queued work from committing after authority changes.

The kill path should be understood as an institutional function rather than a dramatic emergency switch. Different levels of intervention are appropriate for different failures. A single delegation can be revoked. A class of actions can be held. A destination can be blocked. A model can be disconnected from tools while remaining available for analysis. A service can enter read-only mode. A full system shutdown should be the final rather than the only control.

Recovery is equally important. A system that can stop but cannot safely resume creates pressure to bypass governance during outages. Resumption should establish the current authority state, reconcile queued actions, discard work whose permission expired, and identify external changes that occurred while the system was paused.

This makes queued work a special governance object. An action may have been proposed while authority was valid but execute after the authority expired. Conversely, an action held during uncertainty may become valid after new evidence arrives. The commit decision should therefore evaluate authority at execution time rather than assuming that earlier proposal time controls.

> **Admission to a queue does not establish permission, and earlier permission does not necessarily remain valid now.**

54. Administrative Compression as a Measurable Variable

Administrative compression has so far functioned as an explanatory concept. It can be made more empirical. For a defined institutional process, compression can be observed through changes in time-to-resolution, number of human handoffs, number of systems queried, cost per case, number of cases processed per worker, decision latency, and the distance between initial observation and consequential execution.

A simple compression index need not claim universal comparability. Within one process over time, researchers can measure whether technology reduces the resources or elapsed time required to produce the same institutional output. The index can then be compared with changes in authority controls, error rates, reversibility, appeals, and unauthorized-transition events.

This creates a test of the central theory. If compression increases while authority controls remain constant, do misuse or error-propagation risks rise? If explicit authorization improves at the same time, are those risks contained without restoring the old latency? If no relationship appears across comparable cases, the explanatory value of administrative compression should be reduced.

The same measurement can distinguish useful modernization from dangerous shortcutting. A process that falls from ten days to ten minutes while preserving jurisdiction, evidence, proportionality, and redress is a governance success. A process that falls to ten milliseconds by silently eliminating those controls has achieved a different kind of compression.

> **The goal is not less compression. The goal is to measure what was compressed and verify that authority was not compressed out of the process.**

55. Authority Debt as an Operational Balance Sheet

Authority debt becomes useful only if it can be observed. The concept refers to unresolved governance obligations created when capability expands faster than the institution can represent, enforce, or audit permission. Like technical debt, it can be rationally incurred. A pilot may temporarily rely on manual controls. An emergency may justify a narrow exception. The problem is unpriced accumulation.

Possible indicators include standing credentials broader than current duties; delegations without expiry; policies that exist only in prose while execution remains unrestricted; systems that cannot identify the authority source for a consequential action; vendor integrations without credible exit procedures; queues that do not revalidate permission at execution; missing audit evidence; and high-impact AI uses operating under unresolved compliance exceptions.

Debt service consists of the work required to restore alignment: narrow privileges, encode policy, attach provenance, expire delegations, test reversal, build evidence, retain expertise, renegotiate contracts, or redesign workflows. Institutions can then distinguish temporary accepted debt from unknown debt and from debt that exceeds a defined risk tolerance.

This framing also avoids moralizing modernization. A system can be valuable and still carry authority debt. The relevant management question is whether the institution knows the debt exists, understands the consequence if it is exercised, and has a funded path to retire it.

> **Capability deployed without a reconstructable authority path creates an obligation, even when no incident has occurred yet.**

56. Comparative Matrix: What the Cases Actually Test

The cases in this paper do not all test the same proposition. Treating them as interchangeable would manufacture coherence that the evidence does not contain. The historical surveillance cases test private infrastructure, collection power, and governance response. PROMIS tests resolution, procurement conflict, and epistemic dependence. The pre-9/11 and post-9/11 record tests fragmentation and integration. Snowden tests contractor dependence and privileged access. Cloud procurement tests infrastructure dependence and deliberate mitigation. Flock tests capability compression. Estonia tests distributed resolution. Contemporary AI governance tests the transition from interpretation to action.

The comparative value appears when those cases are placed against the same control boundaries without pretending they share a lineage. Scope asks what the system may touch. Identity asks what entity is resolved. Jurisdiction asks which institution may act. Handoff asks how responsibility moves. Authorization asks whether the action is permitted. Proportionality asks whether the permitted action fits the consequence. Auditability asks whether the institution can reconstruct the result.

A case can therefore support one boundary and remain silent on another. SHAMROCK can illuminate public-private collection without proving anything about modern agent delegation. Estonia can falsify a centralization claim without proving that all interoperability is safe. Federal AI guidance can demonstrate emerging pre-deployment controls without proving their effectiveness. The matrix is valuable because it prevents evidence from carrying more weight than it can support.

56.1 The strongest surviving cross-case proposition

After the adversarial cases, the broadest defensible proposition is not that integration causes abuse, private infrastructure captures the state, crisis inevitably expands power, or AI uniquely threatens legitimacy. Each of those stronger formulations has direct counterevidence.

The surviving proposition is functional: as institutions reduce the technical and administrative distance between information and consequential action, the system increasingly needs an explicit way to establish who may cause the transition, under what authority, with what limits, and with what evidence. The exact control can vary by consequence and institution. The need to distinguish capability from permission does not.

> **The continuity is functional, not genealogical: systems keep reducing the distance between knowing and doing.**

57. Federal AI as a Natural Experiment in Governance Lag

The United States federal government now provides a useful natural experiment because capability adoption, formal requirements, and implementation evidence can be observed at the same time. GAO reported that selected agencies' AI use cases nearly doubled from 571 in 2023 to 1,110 in 2024, while generative-AI use cases increased roughly ninefold from 32 to 282. [S37] That is a measurable capability expansion, not merely a policy forecast.

The governance layer is expanding too. GAO identified 94 government-wide or government-wide-implication AI requirements as of July 2025. [S38] For high-impact uses, federal guidance includes pre-deployment testing, risk mitigation, documentation, accountability review, and discontinuation of non-compliant AI functionality. [S38] The important point for this paper is that capability and governance are moving simultaneously but not necessarily at the same rate.

Implementation remains uneven. GAO's 2026 updates show agencies closing some inventory and compliance recommendations while others remain incomplete. [S39] GAO also found privacy-related gaps in government-wide AI guidance and resource limitations that make protection harder in practice. [S40] This is closer to the paper's governance-lag claim than a simple story in which rules always arrive only after deployment. Governance can precede, accompany, and follow capability, with different components maturing at different speeds.

That observation narrows the theory productively. Governance lag should be measured at the control level. An institution may have an AI strategy before deployment but lack impact measurement. It may inventory systems but lack operational revocation. It may require testing but lack the expertise to evaluate a vendor. It may have a legal authority but lack machine-enforceable scope. The lag is therefore multidimensional rather than a single date.

> **The relevant question is not whether governance exists. It is whether the controls required by the action surface are operational when the capability becomes executable.**

58. The Last Mile: From Institutional Rule to Executable Constraint

The final unresolved distance in the architecture is the distance between a rule that exists and a rule that actually constrains execution. Institutions already possess enormous bodies of law, regulation, policy, contract, procedure, and professional judgment. The problem is not absence of rules. It is translation from institutional meaning into an enforceable decision at the point of action.

That translation cannot be reduced to converting every statute into code. Legal and institutional rules contain ambiguity, standards, exceptions, hierarchy, jurisdiction, and facts that require interpretation. Some questions properly remain with human officials, courts, or other authoritative processes. An execution boundary should therefore be able to recognize when authority is resolved and when it is not.

This is the importance of an explicit hold. A system that encounters unresolved conflict does not need to invent a legal interpretation merely because an API call is waiting. It can preserve the proposed action, identify the unresolved authority question, route the issue to the correct decision-maker, and resume only after a valid resolution exists. The machine remains fast where authority is clear and deliberately stops where legitimacy is unresolved.

The same approach protects against a different failure: allowing the model to translate policy and then treat its own translation as permission. A model can assist in extracting obligations, classifying facts, or proposing a policy interpretation. But the component that benefits from a permissive interpretation should not be able to manufacture the authority on which its execution depends.

> **The last mile of governance is not policy generation. It is the binding of a consequential state transition to authority the executing component cannot create for itself.**

PART VII

THE MACHINE-SPEED STATE

59. The Action Surface: Why Agentic Systems Change the Unit of Governance

The next transition is not simply from smaller models to larger models. It is from systems that produce representations to systems that can cause state transitions. A conventional model may classify, summarize, rank, or draft. An agentic system may retrieve records, call tools, route a case, update a database, initiate a transaction, or coordinate additional agents. The difference is architectural because the output is no longer necessarily the end of the computation. It can become an instruction to the world.

The OECD now describes this distinction directly: generative systems answer while agentic systems act, and public-sector agentic systems can execute sequences of actions across government tools and workflows. [S32] The report consequently identifies the action space granted to the system, approval points, traceable decision trails, oversight, redress, and the ability to pause, reverse, or challenge automated actions as governance concerns. [S32] That is not proof of this paper's model. It is independent evidence that the object of governance changes when model output becomes executable behavior.

Content governance asks whether an output is acceptable. Execution governance asks whether a proposed transition is authorized. The two can overlap, but they are not substitutes. A perfectly accurate recommendation can still be unauthorized to execute. A lawful action can still be proposed by a system lacking standing to initiate it. A valid permission can expire before execution. An authorized first step can produce an unauthorized second step when context changes.

This is why the action surface is a better unit of analysis than the model alone. The action surface includes every consequential operation the system can reach: tools, APIs, credentials, databases, payment rails, communications channels, administrative workflows, physical devices, downstream agents, and human decision points. Capability expands when that surface expands even if the underlying model does not change.

Govern the action surface, not merely the model that proposes actions.

60. Authority Propagation Across Multi-Agent Systems

A single-agent system already creates a separation problem between proposal and execution. A multi-agent system adds a propagation problem. One component may discover a task, another may plan it, another may retrieve facts, another may call a tool, and another may verify the result. If authority is represented only as ambient credentials or inherited session context, each handoff can silently widen the practical permission available to the system.

The safe assumption is therefore not that authority follows the task. Authority follows the warrant. Delegation must identify what power is being delegated, by whom, to which component, for which object, under which conditions, for how long, and with what downstream delegation rights. A task description is not enough because task semantics and institutional authority are different objects.

This distinction matters when an orchestrator decomposes one authorized objective into many operations. The original authorization may cover the objective but not every method the planner invents. A system authorized to notify a beneficiary is not necessarily authorized to disclose every record it can retrieve while composing the notice. A system authorized to prepare a payment is not necessarily authorized to release it. A system authorized to query a database is not necessarily authorized to write back an inferred fact.

The authority graph developed earlier therefore needs propagation semantics. Every edge should be capable of narrowing authority; no edge should enlarge it without a new authoritative act. The provenance record should make that narrowing visible. Otherwise a chain of individually plausible handoffs can produce authority laundering: no component explicitly grants the final power, yet the system arrives at it through accumulated capability.

> **Delegation may preserve or narrow authority. It must not silently manufacture new authority.**

61. Identity, Standing, and the Difference Between Authentication and Authority

Modern systems are good at authentication and frequently weak at standing. Authentication establishes that a principal is who the system believes it to be. Authorization establishes that the principal has a permission. Standing asks a different question: is this principal entitled to invoke this authority for this matter, at this time, in this procedural posture?

The distinction is familiar in institutions even when software collapses it. A prosecutor, benefits officer, physician, contracting officer, police officer, administrator, or corporate approver may possess credentials that authenticate successfully while lacking authority over a particular case. Jurisdiction, assignment, conflict rules, purpose limitation, separation of duties, and procedural stage can all matter after identity is known.

Agentic systems make the distinction harder because an agent can act through a service account, delegated token, API key, or tool credential that says little about the institutional basis for the particular action. The credential proves technical access. It does not necessarily prove case-specific standing. If the architecture treats possession of the credential as sufficient authority, organizational meaning has been compressed into a technical secret.

A mature execution boundary therefore needs more than identity and role. It needs authority provenance tied to the object of action: the case, account, person, resource, jurisdiction, purpose, and current procedural state. This is especially important where one agent can act for multiple users or where a shared service operates across multiple agencies or tenants.

> **Authentication answers who can present the credential. Standing answers why this principal may exercise this power here.**

62. Purpose Limitation as an Executable Constraint

Many information-governance regimes distinguish access from purpose. A record may be technically accessible for one institutional purpose and impermissible for another. Historically, purpose limitation has often lived in policy, training, audit, and after-the-fact discipline. Machine-speed systems create pressure to represent it closer to execution.

Purpose is difficult because it is partly semantic. A system can label a request as fraud prevention while pursuing a different objective. A model can generate a plausible justification after deciding what it wants to do. The answer cannot be to trust the acting component's self-description. Purpose has to be bound to an authoritative workflow, case state, mandate, or other external evidence where possible.

That does not mean every purpose question can be compiled into deterministic code. Some require interpretation. The architectural requirement is narrower: where purpose is a condition of authority, unresolved purpose should remain unresolved rather than being silently converted into permission. The system can pause the action, seek a higher-authority determination, or restrict itself to the subset of operations whose purpose is already established.

This turns purpose limitation from a paragraph in policy into a potential control surface. It also makes misuse more legible. A receipt can record not only who acted and what was done, but which authorized purpose was asserted, what artifact established it, and whether the action remained inside that purpose.

> **Access to information is not authority to use it for every purpose the system can imagine.**

63. The Economics of Friction: Why Removing Delay Is So Attractive

Administrative compression is not merely a technological preference. It has an economic logic. Search time, handoffs, duplicate entry, manual reconciliation, queueing, staffing, and review all impose costs. Institutions are rewarded for reducing those costs, especially when users experience delay as service failure and leaders experience backlog as operational failure.

This explains why the removal of friction is often rational even when it creates governance risk. A faster eligibility determination can deliver benefits sooner. A shared identity layer can eliminate repeated verification. Interoperability can stop people from carrying the same document between agencies. Automated routing can reduce missed deadlines. AI can reduce the cognitive cost of reading large records. These are real gains, not camouflage for power expansion.

The problem is that cost accounting usually sees computational friction more clearly than governance friction. A manual approval has an observable labor cost and queue time. The value of the approval may be visible only when it prevents a wrongful action. A jurisdiction check looks like latency until the day it stops an institution from acting outside its mandate. Audit evidence looks like overhead until a disputed action has to be reconstructed.

Selective friction therefore has to survive economic pressure. If governance controls are bolted on as expensive exceptions, optimization will tend to route around them. The stronger design is to make authority resolution part of the normal execution path and to automate the controls that can be automated without allowing the acting system to self-authorize.

> **A control that exists only as optional overhead will eventually compete with the system it is supposed to govern.**

64. Latency Budgets for Legitimacy

Selective friction raises a practical question: how much delay is legitimate? The answer cannot be that every consequential action requires a slow human committee. That would recreate the fragmentation and impunity failures documented earlier. Nor can the answer be zero latency, because some authority questions require facts, review, contest, or independent judgment.

A useful architecture separates authority checks by consequence and uncertainty. Low-consequence, well-specified actions can resolve automatically against stable policy. Higher-consequence actions can require stronger evidence, separation of duties, or human approval. Ambiguous cases can remain on hold. Emergency paths can exist, but they should carry explicit scope, expiry, elevated logging, and later review rather than becoming permanent shortcuts.

The relevant performance measure is therefore not raw execution speed. It is time to valid authority. A system that executes instantly and reverses thousands of wrongful actions later is not necessarily more efficient than one that spends milliseconds or minutes establishing the required authority before committing. Likewise, a system that waits days for a ceremonial signature after every relevant fact is already resolved may be imposing governance theater rather than governance.

This suggests a latency budget for legitimacy: minimize delay that does not improve authority resolution, while preserving the delay or independent act necessary to establish authority where the consequence warrants it. The budget should be explicit enough to measure and challenge.

The goal is not maximum friction or minimum friction. It is minimum sufficient friction for valid authority.

65. Human Review Is a Control Only When the Human Has Power

Human-in-the-loop language can conceal a weak control. A person who sees a recommendation but cannot inspect its basis, lacks time to disagree, is measured on throughput, or cannot stop execution is not an independent authority boundary. The human is part of the interface, not necessarily part of governance.

Automation bias makes this distinction operationally important. As systems become reliable enough to be useful, reviewers may rationally rely on them. High volume can turn review into confirmation. Specialized models can create epistemic asymmetry in which the reviewer is formally superior but practically unable to contest the system. The same dependency pattern seen in procurement can occur inside a decision workflow.

A meaningful human control therefore needs at least three properties. The reviewer must have standing to decide; the reviewer must receive enough evidence to understand the material basis of the proposed action; and the reviewer must possess real power to alter, deny, or hold the action without the system treating disagreement as an error condition. Where those conditions are absent, labeling the workflow human-reviewed overstates the control.

This also explains why deterministic pre-execution enforcement and human judgment are not opposites. A deterministic boundary can enforce that a required human decision exists and is current without pretending to replace the human reasoning that produced it.

A human in the loop is not a safeguard unless the human can actually govern the loop.

66. The Receipt Problem: Evidence That Survives the Decision

Ordinary logs are often optimized for debugging, security monitoring, or operational analytics. An authority receipt has a different purpose. It should make it possible to reconstruct why a consequential state transition was permitted at the time it occurred, even after models, policies, personnel, and surrounding systems have changed.

That requires more than recording the final verdict. The evidence should preserve why the particular action was permitted and how the institution’s decision relates to what actually happened. The exact form will vary by domain. The invariant is that the record must preserve the authority relation rather than merely the existence of an API call.

This matters for appeal and institutional learning. If an affected person challenges an automated action, a current model cannot safely reconstruct the past by generating a plausible explanation. The institution needs evidence of the decision that actually occurred. If a policy changed after execution, the receipt should preserve the policy state that governed the original action rather than silently applying the new rule to the old event.

Receipts also create a boundary on institutional memory. They can make accountability durable without requiring every system to retain every underlying datum indefinitely. A receipt can point to or cryptographically bind authoritative artifacts while data-retention rules govern the underlying records separately.

> **Accountability requires evidence of the authority that existed then, not an explanation generated now.**

66.1 The receipt can become a surveillance system

The authority receipt creates its own governance risk. A corpus that records who asked an institution to act, on whose behalf, against which subject or resource, for what purpose, under which rule, using which evidence, and with what result can become more sensitive than the operational logs it was designed to improve. Accountability infrastructure can therefore become a secondary surveillance system if retention and access are treated as an afterthought.

The receipt should preserve the minimum evidence necessary to reconstruct authority, not a permanent copy of every underlying datum. Where feasible, it should reference or cryptographically bind source artifacts rather than duplicate them; separate identity from policy facts when later linkage is unnecessary; encrypt sensitive fields; support field-level or role-bounded disclosure; and inherit domain-specific retention, evidentiary-hold, and deletion rules. Key retirement or cryptographic erasure may be appropriate where law permits it, but cannot be presented as a universal substitute for records obligations.

Access to receipts should itself be governed. An auditor, administrator, vendor, investigator, affected person, court, and security operator may each have different legitimate views of the same event. The architecture should therefore support selective disclosure and record who accessed which receipt fields for which authority and purpose. A receipt that proves an earlier decision should not silently create a new unrestricted intelligence surface.

Affected-person access also requires care. Contestability may require disclosure of the authority basis and material reasons for a decision while protecting third-party privacy, investigative integrity, protected sources, trade secrets, or security-sensitive implementation details. The correct disclosure boundary is domain-specific, but the conflict must be explicit. Auditability and minimization are both governance requirements.

Finally, cryptographic integrity does not establish substantive correctness. A tamper-evident receipt can prove that a particular artifact and decision record existed in a particular form. It cannot prove that a factual input was true, that the policy was lawful, that the authority source was legitimate, or that the outcome was just. The receipt is evidence of what the system did under a represented authority state, not a certificate of legitimacy.

> **Evidence necessary for accountability should not create a second surveillance system, and integrity of the receipt must never be confused with legitimacy of the decision.**

67. Appeal, Remedy, and the Architecture After Execution

Pre-execution governance cannot eliminate error. Facts can be wrong, authority can be misapplied, policy can be unlawful, and authorized officials can make bad decisions. A complete authority architecture therefore needs a post-execution path that is as deliberate as the pre-execution path.

The Council of Europe AI Framework Convention is useful here because its governance structure includes documentation, procedural safeguards, complaint mechanisms, and the ability of affected persons to challenge AI-mediated decisions. [S22] The significance for this paper is not that one treaty supplies a universal design. It is that contemporary governance increasingly treats contestability and remedy as lifecycle properties rather than optional ethics features.

A remedy architecture asks what can be restored. Some state transitions are reversible: a payment can be stopped, a record corrected, access restored, or a transaction rolled back. Others are only partially reversible: disclosure cannot be undisclosed; detention time cannot be returned; a public accusation cannot be perfectly erased. Consequence should therefore affect the strength of pre-execution authority checks.

The receipt and the remedy path connect. Without a durable record of what authority was used, appeal becomes a contest between current narratives. With it, a reviewer can test the original action against the rule, evidence, and authority that existed at execution and then determine whether reversal, compensation, correction, escalation, or institutional change is warranted.

> **Governed execution includes the power to contest what happened and a defined path for what happens next.**

68. Continuous Authorization: Authority Can Change While the System Is Running

Traditional authorization often assumes a relatively static system: assess it, approve it, operate it, and reassess periodically. Cloud infrastructure, continuous deployment, dynamic data sources, changing models, and agentic workflows weaken that assumption. The system that was authorized yesterday may not be materially identical to the system acting today.

Federal cloud governance is moving toward more structured authorization data and continuing authorization. GAO reports that FedRAMP modernization is moving previously authorized providers toward machine-readable authorization data for both initial and continuing authorization. [S36] FedRAMP's own 2026 work similarly emphasizes machine-readable authorization materials, with the final scope modified after public comment rather than simply adopting the original RFC unchanged. [S41] The narrower inference is important: authorization is increasingly being treated as data that can be maintained and evaluated over time.

For agentic systems, continuous authorization has another dimension. The surrounding authority state can change even when the software does not. A user can revoke consent. A case can transfer jurisdiction. An emergency can expire. A credential can be compromised. A court can issue an order. A policy can change. A delegated principal can leave a role. The execution boundary therefore needs current authority, not merely evidence that authority existed when the workflow began.

This makes temporal binding central. Long-running plans should re-resolve authority at consequential transitions rather than assuming the original permission survives every later state. The more autonomous and persistent the agent, the less safe ambient, indefinite authorization becomes.

> **Authorization is not a one-time property of a running system. It is a condition that can change before the next action.**

69. From Documents to Computable Authority

Institutions express authority through heterogeneous artifacts: statutes, regulations, orders, warrants, contracts, delegations, licenses, policies, case assignments, approvals, consent records, role definitions, and emergency declarations. Historically, humans bridge these forms by interpretation. Machine-speed execution creates pressure to make at least some of their operational consequences computable.

The temptation is to treat this as a document-conversion problem: extract the rule, encode it, and execute. That is too simple. Authority has hierarchy, scope, exceptions, temporal conditions, conflicts, and dependencies on facts that may themselves be disputed. A computable representation should therefore preserve provenance and uncertainty rather than flattening every source into an unconditional Boolean permission.

One practical model is a two-layer structure. The authoritative artifact remains the legal or institutional source. A machine-readable authority object represents the operational claim derived from it: who may do what, to which object, for which purpose, under which conditions, until when, with what delegation rights, and according to which source. The representation can be invalidated or superseded without pretending to replace the source itself.

This is another place where claim <= artifact matters. The computable authority object cannot legitimately carry more permission than the underlying artifact. If interpretation is contested, the object should expose the unresolved question or require an authoritative determination instead of manufacturing certainty for the sake of execution speed.

> **Computable authority should be derived from authoritative artifacts without becoming more authoritative than the artifacts themselves.**

70. The Policy Compiler Problem

Once institutions attempt to make authority computable, a new control problem appears: who governs the translation? A policy compiler, whether human-authored, rules-based, or AI-assisted, can change the practical meaning of institutional authority by deciding how text maps to executable conditions.

This creates a separation-of-functions requirement. The component that proposes an interpretation should not be able to make that interpretation authoritative merely by emitting code or policy. Translation needs provenance, review appropriate to consequence, versioning, test cases, conflict handling, and a path to challenge the compiled result against the source artifact.

AI can be useful in this process. It can identify candidate obligations, extract conditions, compare versions, surface conflicts, and generate tests. But probabilistic extraction is evidence about what a source may require, not authority to bind the institution. The final executable rule needs a recognized adoption mechanism outside the model.

The failure mode is subtle because the compiled policy may look deterministic after the fact. Determinism at execution does not cure an unauthorized or incorrect compilation step. A perfectly deterministic gate can consistently enforce the wrong rule. Governance therefore has to cover both rule creation and rule execution.

> **Deterministic enforcement is only as legitimate as the authority and provenance of the rule being enforced.**

71. Physical AI: When the State Transition Leaves the Database

The execution-boundary problem becomes easier to see when software controls physical systems. A database update can sometimes be reversed. A robot arm, autonomous vehicle, industrial controller, medical device, or other embodied system can create consequences that cannot be undone by restoring a previous record.

Physical action also changes the timing problem. A safety-critical controller may need to act faster than a human can approve each motion. That does not eliminate authority; it moves authority earlier. Humans and institutions authorize an envelope of permitted behavior, operational conditions, fail-safe states, and escalation rules. The machine can then act rapidly inside that envelope while encountering a hard boundary when the proposed action exceeds it.

This is the same selective-friction principle at a different timescale. The authority decision need not occur at the same millisecond as every motor command. It can be established through validated constraints, task authorization, environmental conditions, and bounded delegation before execution, with runtime checks preserving the envelope. When the environment changes beyond the authorized assumptions, a hold or a safe-state transition becomes a governance action rather than a performance failure.

Physical AI therefore exposes why authority cannot be reduced to content moderation. The critical question is not whether the system says an unsafe or impermissible thing. It is whether the next physical state transition is inside the authority and safety envelope established for the system.

> **When execution becomes physical, the cost of confusing capability with authority becomes immediate.**

72. Cross-Border and Cross-Institution Authority

Interoperability increasingly allows one institution to resolve information held by another. Cross-border digital services extend the problem further: identity, credentials, records, and requests can traverse jurisdictions whose legal authority, retention rules, remedies, and institutional roles differ.

The technical success condition is straightforward: the systems exchange the necessary data reliably. The authority condition is harder. A receiving institution may be able to read a record without having authority to use it for every purpose. A credential recognized across borders may establish identity without transferring the issuing institution's powers. A service can be technically reachable while the requested action remains outside jurisdiction.

This reinforces the Estonia lesson developed earlier. Resolution does not require centralized possession, and interoperability does not require collapsing institutional boundaries. In fact, preserving those boundaries can be a design advantage if the exchange layer carries enough provenance and authorization context for each institution to make its own valid decision.

Cross-institution execution should therefore resist authority transitivity. Trusting another institution's identity assertion, fact, or credential does not automatically mean inheriting every authority claim attached to it. Each consequential state transition still needs a valid basis in the institution that will execute it.

Interoperability can move evidence across boundaries. It does not erase the boundaries that determine authority.

73. The Architecture of Refusal

High-capability systems are usually designed around successful execution. Governance requires equal attention to refusal. A system needs a principled way to say no, not yet, not under these conditions, or not without another authority.

Refusal and a hold represent different institutional states. Refusal means the authority question has been resolved against the proposed action. A hold means the institution cannot yet establish sufficient authority. Conflating them creates two problems: ambiguous cases may be treated as prohibited when they merely need resolution, or uncertainty may be treated as permission because the system is optimized to keep workflows moving.

Modification adds a third possibility. Sometimes the proposed action is too broad while a narrower action is authorized. Data can be minimized, scope reduced, recipients constrained, or a proposed transaction converted into a draft requiring approval. This preserves useful capability without pretending that authority is all-or-nothing.

A mature refusal architecture should also make failure legible. The system should identify which authority condition failed without exposing protected information unnecessarily, preserve the proposed action where appropriate, and route the matter to the actor capable of resolving it. Refusal becomes part of institutional coordination rather than a dead end.

A governed system needs first-class semantics for not executing.

74. Measuring the Gap Between Capability and Control

The theory becomes more useful if the capability-authority gap can be measured. No single metric will capture legitimacy, but several operational indicators can test whether governance is keeping pace with the action surface.

Candidate measures include the proportion of consequential actions with machine-verifiable authority evidence; the proportion whose authority can be traced to an authoritative artifact; time to valid authority; rate of hold resolution; percentage of delegations with explicit expiry; proportion of actions reversible after error; percentage of policy changes propagated to runtime enforcement within a defined window; audit reconstruction success; frequency of unauthorized attempts blocked before execution; and the number of privileged capabilities lacking an identified institutional owner.

These measures should be interpreted carefully. A high block rate may indicate strong enforcement or a badly designed workflow. A low hold rate may indicate clear authority or a system that silently defaults to permission. Metrics therefore need paired qualitative review and adversarial testing. The objective is not a governance score that hides institutional judgment behind a number.

The OECD data provide one macro-level example of this measurement problem. AI is already used in at least one government function in 35 of 36 surveyed OECD countries, while impact measurement and several operational governance mechanisms remain much less widespread. [S30] The useful observation is not that every deployment is ungoverned. It is that capability adoption can be measured separately from the maturity of specific control surfaces.

> **Governance lag should be measured control by control, not inferred from the existence or absence of a strategy document.**

75. A Second Falsification Pass: Where the Strong Thesis Still Fails

The expanded model creates new opportunities to overclaim. First, machine-readable authority is not necessarily better authority. A bad rule can become easier to enforce at scale. Second, pre-execution checks can create false confidence if the facts they consume are wrong. Third, receipts can document injustice perfectly without preventing it. Fourth, human review can be essential in domains where legitimacy depends on discretion, empathy, or public accountability rather than rule resolution alone.

Fifth, some institutions legitimately need broad discretion under uncertain and time-critical conditions. An architecture that requires complete ex ante specification may fail in emergencies, field operations, or novel cases. The appropriate response is not to pretend the uncertainty disappears, but to bound discretion, record its exercise, expire exceptional authority, and subject it to later review.

Sixth, distributed authority can itself create impunity. If every component can refuse responsibility because another component supplied the data, model, policy, credential, or approval, the system can become highly auditable yet institutionally ownerless. Governance therefore needs accountable ownership in addition to technical provenance.

Seventh, the execution boundary cannot determine whether the underlying law or policy is just. It can enforce that an action is authorized under the institution's valid rules and preserve evidence of that authority. Democratic legitimacy, constitutional validity, substantive rights, and moral judgment remain larger questions. Architecture can constrain power; it cannot replace politics or law.

These limits narrow the strongest defensible claim again. The paper does not establish that explicit authorization boundaries guarantee legitimate outcomes. It establishes that as technical friction disappears, systems need explicit mechanisms capable of distinguishing executable capability from institutionally authorized action, while preserving contestability, accountability, and the possibility of refusal.

> **An execution boundary can enforce the existence and scope of authority. It cannot make an unjust authority just.**

76. The Machine-Speed State

The historical sequence can now be stated without requiring a secret lineage. Institutions first externalized memory, then made records searchable, then networked them, then resolved identity across systems, then made data interoperable, then moved computation into elastic infrastructure, then added ubiquitous sensing, and now increasingly automate interpretation and action. Each layer reduces a different form of administrative distance.

The result is not a state without humans. It is a state in which human and institutional decisions can be projected through infrastructure at a speed and scale that no individual administrator could reproduce manually. A single policy change can alter millions of transactions. A shared identity system can resolve people across services. A cloud control can affect an entire agency. An agent can traverse systems and execute a workflow that previously required multiple offices.

This is the machine-speed state: not government run entirely by machines, but government in which the path from information to consequential action is increasingly computational. The political system may remain democratic, authoritarian, federal, unitary, centralized, or distributed. The architectural problem survives those differences because the compression occurs in the information-action chain itself.

That is why the deepest transition is from accidental to explicit governance. Paper, distance, incompatible systems, staffing limits, office hours, and human attention once constrained what institutions could resolve and execute. Those constraints were crude and often harmful. Their disappearance is generally desirable. But once they disappear, legitimate authority can no longer rely on them as invisible brakes.

The machine-speed state therefore needs an authority architecture proportional to its capability architecture. Identity, jurisdiction, purpose, delegation, time, proportionality, contestability, audit, and revocation have to survive the same compression that made information and execution fast. Otherwise the system modernizes capability while leaving legitimacy in the manual era.

> **The machine-speed state does not need slower information. It needs authority that can survive machine-speed execution.**

PART VIII

THE AUTHORITY PROTOCOL

77. Authority Is a Protocol Problem, Not Merely a Policy Problem

Institutions already possess policies, statutes, delegations, operating procedures, contracts, warrants, licenses, approvals, and professional norms. The unresolved problem is not the absence of rules. It is that these rules often exist in forms that do not travel with the machine action they are supposed to constrain. A human can read a policy and infer whether an action is allowed. A software agent traversing APIs needs something more operational: an authoritative source, a subject, an action, an object, a scope, a time window, a purpose, conditions, and a way to determine what happens when any of those elements are missing or contradictory.

That makes authority a protocol problem. Protocol does not mean that law becomes a packet format or that every exercise of discretion can be reduced to Boolean logic. It means the institution needs a repeatable way to carry authority across boundaries without silently converting capability into permission. The protocol has to answer who is acting, on whose behalf, under what delegation, against which resource, for what purpose, within which jurisdiction, for how long, and with what evidence. If the answer cannot survive a system boundary, the institution has not preserved authority; it has preserved only access.

NIST's 2026 work on software and AI agent identity makes this transition visible. The NCCoE frames the problem around identification, authorization, auditing, non-repudiation, and the risks created when agents gain access to diverse tools and applications. [S42] That does not establish the architecture proposed here, but it independently confirms that identity alone is insufficient once software agents can take consequential actions. The relevant object is not simply an authenticated agent. It is an authenticated agent whose authority for this action can be established and bounded.

A protocol view also exposes failure earlier. If authority depends on a document that cannot be located, a delegation that has expired, a purpose that cannot be established, or a jurisdiction that conflicts with another rule, the correct result is not to let the agent improvise. The unresolved condition is itself information. It should produce a hold, escalation, or refusal according to the institution's rules. Ambiguity is not permission.

A stronger interpretation is tempting: if authority can be represented as a protocol, governance can be automated completely. The record does not require it. The narrower proposition survives. Wherever an institution permits machines to carry actions across system boundaries, it needs an authority representation capable of surviving those boundaries without becoming weaker than the rule it represents.

> **Authority must survive the system boundary that carries the action.**

78. The Delegation Envelope: Preserving Institutional Limits

Delegation is the mechanism by which institutional authority moves. Human organizations have always used it: a legislature delegates to an agency, an agency head to an office, a supervisor to an employee, a principal to a contractor. Agentic systems make the same structure computational, but they also make delegation easier to copy, chain, and execute at scale. That changes the engineering requirement. A delegation cannot safely be treated as a durable credential that says only that an agent is trusted.

A delegation envelope is useful here as an institutional concept: the limits of a grant should remain intelligible when another actor exercises it. The relevant question is whether the receiving actor can explain the authority relied upon and the boundaries that remain in force. Different domains may represent that relationship differently. The research does not depend on a universal packet, required field inventory, credential structure, or product-specific representation.

This is stricter than ordinary authentication because authentication answers a narrower question: is this entity the entity it claims to be? It is stricter than a static role because a role may be valid while the proposed action is not. It is also stricter than possession of an API token. Tokens are capabilities. They may be evidence relevant to authority, but possession alone does not establish that a particular state transition is institutionally permitted.

The envelope should shrink as it moves. A delegated actor may receive less authority than its principal but should not manufacture more. Redelegation therefore needs monotonicity: downstream authority must be equal to or narrower than the authority upstream unless a separate authoritative source expands it. This gives the system a testable invariant. If an agent with authority to read a record delegates authority to modify it, the chain is invalid even if the downstream agent has technical credentials capable of performing the modification.

This is one place where the distinction between capability and authority becomes operational rather than philosophical. The machine may be able to perform the call. The delegation envelope determines whether the institution has actually authorized that call, in that context, at that time.

Delegation may narrow authority; it may not silently manufacture more.

79. Authority Monotonicity and the Rule Against Self-Expansion

A central risk in autonomous systems is permission expansion by composition. No single component appears to exceed its role, yet a sequence of individually plausible steps produces a capability that no principal explicitly granted. One service can read. Another can transform. A third can write. An agent able to orchestrate all three may effectively possess a new power even though no single credential says so.

The authority model therefore needs a monotonicity rule: computation may refine a request, but it may not increase the authority attached to it. New information can narrow a decision. It can reveal that a condition is unmet. It can select among actions already permitted. It cannot create permission merely because the model inferred that the action would be useful. Authority expansion requires a new authoritative event: a valid approval, delegation, statute, order, consent, or other recognized source.

This rule matters especially for agents that generate sub-agents or dynamically select tools. A parent agent's goal is not a blank check. If the parent has authority to schedule an appointment, a child agent should not inherit authority to disclose unrelated medical data simply because disclosure would make scheduling easier. Goal completion and institutional permission are different optimization functions.

The rule also limits prompt-based authority. A user instruction may be evidence of intent, but its authority depends on who the user is, what the user controls, and what rights or institutional powers the user actually possesses. 'Do whatever is necessary' cannot legitimately expand a delegation beyond the principal's own authority. A model that interprets broad language as unlimited permission converts ambiguity into power.

The falsifier is straightforward. If real systems can permit unconstrained authority expansion through composition without increasing abuse, error, or accountability failures, then monotonicity is unnecessarily strict. The safer working hypothesis is the opposite: as orchestration becomes more powerful, explicit non-expansion becomes more important because capability composition is precisely how hidden action surfaces emerge.

Computation may refine permission. It may not create permission.

80. The Authority Object: Separating Evidence From Decision

An authority decision should not be confused with the evidence used to reach it. A warrant, contract, policy, identity assertion, consent record, risk classification, supervisor approval, and jurisdictional fact may all contribute to a decision. None of them alone is necessarily the decision. Treating one artifact as self-executing collapses interpretation into authority and makes later reconstruction difficult.

The architecture therefore benefits from distinguishing the evidence considered from the institution’s resolved authority decision. An authority object, in this conceptual sense, is evidence of that decision rather than a replacement for the underlying source. Its value is to preserve what the institution permitted in the circumstances and to make that determination available for later review. No record schema or implementation format is prescribed here.

Separating evidence from decision has two consequences. First, evidence can change without rewriting history. A later correction can show that an earlier authority decision rested on a false identity or outdated record while preserving what the institution actually knew at the time. Second, the decision can expire even when the evidence persists. A credential may remain valid while the action-specific authorization has ended.

This structure also prevents logs from masquerading as governance. A conventional audit log can prove that an API call occurred. An authority object can show why the institution permitted it to occur. The difference is substantial. Post-hoc observability answers what happened. Authority evidence answers whether the system had a valid basis to let it happen.

The claim ceiling remains narrow. A perfectly formed authority object can still encode a bad rule or rely on false evidence. Its value is not moral correctness. Its value is making the authorization claim explicit, bounded, inspectable, and separable from the technical ability to execute.

Evidence supports authority. Evidence is not itself the authority decision.

81. Policy Drift and the Half-Life of Permission

Institutions change faster than many deployed systems. Laws are amended. contracts expire. employees change roles. emergency orders end. risk classifications change. consent is withdrawn. court orders are stayed. vendors alter products. An authorization that was valid yesterday may be invalid today even if every credential still works.

This creates a half-life problem for permission. Static access control assumes that authority changes at identifiable administrative moments and that systems will be updated accordingly. Machine-speed administration increases the cost of that assumption. A stale rule can be exercised thousands or millions of times before a human notices that the policy surface and execution surface have diverged.

Continuous authorization is therefore not merely repeated authentication. The system must be able to determine whether the authority basis remains current. That can require checking revocation, policy version, delegation expiry, jurisdiction, purpose, risk state, or an external condition. The frequency of re-resolution should be proportional to consequence and volatility. A low-risk read may tolerate a longer validity window than an irreversible transfer, enforcement action, or physical command.

The operational metric is authority staleness: the interval between a change in the authoritative rule or fact and the point at which runtime enforcement reflects that change. Institutions already measure patch latency, detection latency, and recovery time. Machine-speed governance needs an equivalent measure for permission. A policy that takes thirty days to reach enforcement is not functioning as a thirty-day-old document; it is functioning as a thirty-day gap in authority control.

The objective is not universal real-time policy. Some rules require deliberation and controlled rollout. The requirement is explicitness: the institution should know which authority states are cached, for how long, under what invalidation conditions, and what happens when freshness cannot be established.

Permission has a half-life, and stale authority is a runtime defect.

82. The Revocation Race

Revocation exposes the difference between a permission system designed for people and one designed for autonomous execution. A human employee whose authority is withdrawn may need time to learn of the change, reach a terminal, or attempt another action. An agent can have hundreds of queued or parallel operations already in flight. Revocation that only prevents the next login may arrive after the consequential work is done.

The system therefore needs to define the revocation boundary. Does revocation stop new plans, new tool calls, queued actions, retries, delegated children, long-running workflows, or already authorized but not yet committed state transitions? Different answers produce different residual authority. The correct choice depends on the domain, but leaving the choice implicit is itself a governance decision.

This suggests a revocation race metric: time from authoritative revocation to effective inability to commit covered actions. The metric should include propagation to delegates and caches. A system that revokes the parent but leaves child agents with valid bearer tokens has not revoked the authority chain. It has changed the label while preserving the capability.

Revocation also needs evidence. The institution should be able to prove when authority ended, which components received the change, which actions were already committed, and whether any action crossed the boundary afterward. This is particularly important when a later dispute turns on timing. Without a durable revocation record, the organization may know that access was eventually removed but not whether a specific action was authorized at the instant it executed.

The deeper point is that authority has negative transitions as well as positive ones. Governance is not only the ability to say yes before execution. It is the ability to make no become effective before the system outruns the institution.

Revocation is real only when covered execution can no longer commit.

83. Retry Semantics: When the Same Authorized Action Is No Longer the Same Action

Distributed systems retry. Networks fail, services time out, queues redeliver, workers crash, and agents re-plan. Reliability engineering treats retry as a normal mechanism for completing intended work. Authority architecture has to ask a different question: when is a retry still covered by the original authorization?

The answer cannot be assumed. A payment authorized for one amount should not be retried after the amount changes. A search authorized for an active investigation may not remain authorized after the case closes. A physical command may become unsafe after the environment changes. Even an identical API request can have a different consequence if the underlying state has moved.

This means authority needs idempotence semantics of its own. The authorization should identify whether it covers exactly one commit, a bounded number of attempts, a transaction identifier, a time window, or a class of equivalent actions. If execution fails before commitment, a retry may remain within scope. If the action partially commits, the system may need reconciliation rather than repetition. If material facts change, the authority should be re-resolved.

Agentic systems make this more important because a retry may not be literal. The agent can choose a different tool or route to accomplish the same goal. From the model's perspective that is resilience. From the institution's perspective it may be a new action surface requiring a new authorization decision. Goal equivalence is not authority equivalence.

The governing invariant is therefore narrower than 'authorized once, authorized until success.' Authorization attaches to a bounded state transition under specified conditions. Reliability mechanisms may preserve intent, but they must not silently preserve authority after the conditions that justified it have changed.

Reliability may retry intent; it must not silently retry expired authority.

84. Transaction Boundaries and Partial Execution

Many consequential workflows are not atomic. An agent may reserve funds, update a record, send a notice, call a vendor, and trigger a physical process. If the workflow fails after step three, the institution is left with a partially executed decision. Traditional governance often assumes a human can notice the inconsistency and repair it. Machine-speed systems need the repair semantics before the failure occurs.

The authority question is not only whether each step is permitted. It is whether the institution authorizes the intermediate states created by partial completion. A workflow may be legitimate only if all steps succeed together. In that case the architecture needs transactional controls, compensating actions, or hold states that prevent irreversible commitment until the required conditions are satisfied.

This is where reversibility becomes more than a convenience. A compensating action itself requires authority. Refunding a payment, restoring a benefit, retracting a notice, or reversing a configuration change may have different legal and operational consequences from the original action. Systems that treat rollback as a purely technical function risk creating an ungoverned second execution path.

The same issue appears in human institutions under names such as remedy, correction, reconsideration, and restoration. Automation does not eliminate those concepts. It forces them into the runtime architecture. A machine-speed institution that can act immediately but can only repair manually has compressed one side of the power relationship more than the other.

A mature execution boundary therefore asks two questions before commitment: is the proposed transition authorized, and does the institution possess an authorized recovery path if the transition fails or is later found to be wrong? Not every action can be reversible, but irreversibility should be a known property of the decision rather than an accidental discovery after execution.

> **A governed action includes the authority to recover from partial execution.**

85. The Two-Key Problem: Separation of Duties at Machine Speed

Separation of duties is one of the oldest ways institutions prevent concentrated power. One person initiates, another approves. One office investigates, another adjudicates. One system proposes a change, another controls production. Agentic automation can preserve this structure, weaken it, or simulate it without actually creating independence.

Two model calls are not two keys if both operate under the same authority, evidence, objective, and failure mode. A reviewer agent generated from the same context may reproduce the proposer's assumptions. Independence is an institutional property, not a count of components. Meaningful separation requires distinct authority or accountability: different principals, different policy roles, independent evidence, or a control plane the proposing system cannot rewrite.

This matters because automated approval chains can become ceremonial. A workflow may technically contain a review step while the reviewer lacks time, information, power, or incentive to refuse. Section 65 already established that human review is a control only when the human has power. The same is true for machine review. An approval stage that cannot materially change execution is observability, not authorization.

The machine-speed version of two-key control should therefore define which actions require independent concurrence, what makes the second authority independent, whether either party can place the action on hold, and how disagreement is resolved. High-consequence actions may require stronger independence than routine transactions. Emergency override may exist, but it should be explicit, bounded, expiring, and reconstructable.

The objective is not to multiply approvals indiscriminately. Excessive separation can recreate the impunity side of the model by making legitimate action impossible. The design problem is selective independence: enough distinct authority to prevent self-authorization without rebuilding administrative paralysis.

> **Two approvals are not two keys unless the authority is meaningfully independent.**

86. Authority Quorum: When No Single Actor Is Enough

Some institutional actions are legitimate only when authority is collective. Boards vote. courts use panels. financial controls require multiple signers. incident response may require security and business approval. Autonomous systems need a way to represent that collective structure without reducing it to a single service account that technically possesses all necessary permissions.

An authority quorum defines the minimum combination of independent approvals required for a class of action. The quorum can be role-based, risk-based, jurisdictional, or conditional. It can also include veto authorities whose refusal is sufficient to stop execution. What matters is that the runtime system preserves the institutional structure rather than collapsing it for implementation convenience.

Quorum introduces its own failure modes. Colluding actors can satisfy it. unavailable approvers can block necessary action. copied credentials can create false multiplicity. automated delegates can all inherit from the same principal and appear independent when they are not. The authority graph therefore needs provenance strong enough to determine whether the quorum represents genuinely distinct sources of authority.

The architecture should also distinguish approval quorum from evidence quorum. Three systems agreeing on a fact do not necessarily create legal authority, and three authorized officials cannot make a false fact true. Facts and authority interact, but they are different dimensions. A robust decision may require both: sufficient evidence for the predicate and sufficient authority for the consequence.

Collective authority is another reason a simple access-control list is too weak for the machine-speed state. The institution may not be asking whether any one actor can execute. It may be asking whether the required configuration of actors, evidence, conditions, and time has been satisfied before execution becomes legitimate.

> **Collective power must remain collective at the execution boundary.**

87. Purpose Binding and the Reuse Problem

Data governance has long struggled with secondary use. Information collected for one purpose becomes attractive for another because the marginal technical cost of reuse is low. Interoperability and AI intensify the problem: once information is resolvable and machine-readable, a new model can extract value without moving the data into a visibly new database.

Purpose limitation therefore has to survive computation. It is not enough to label a dataset 'for benefits administration' if an agent can query it while pursuing an unrelated enforcement objective. The authority decision needs to bind the proposed action to a recognized purpose and determine whether the evidence being used is permitted for that purpose.

This does not mean purpose is always singular or immutable. Statutes can authorize multiple uses; emergencies can create exceptions; consent can expand or narrow use; courts can compel disclosure. The important property is that purpose expansion comes from an authoritative event rather than from the technical convenience of having the data available.

Purpose binding also limits inference laundering. A system should not evade a restriction on using protected data by deriving a proxy from that data and then claiming the proxy is a new, unrestricted fact. The relevant governance question is provenance: what information materially contributed to the inference, and is that use within the authorized purpose? Perfect lineage may be difficult, especially with learned models, but difficulty does not make the institutional question disappear.

The broader ratchet is clear. As resolution costs fall, reuse becomes easier. The governance response cannot be to restore universal data silos. It has to make purpose a first-class constraint capable of traveling with access, inference, and execution.

Availability of data does not expand the purpose for which it may be used.

88. The Standing Check: Who Is Allowed to Ask the Institution to Act?

Much of computer security begins after a request exists. It asks whether the requester is authenticated and whether the requester has permission. Institutional authority often has an earlier question: does this actor have standing to invoke this process at all? A citizen may have standing to request a record but not to initiate an enforcement action. An employee may report an incident but not adjudicate it. An agent may be authorized to assist a user but not to manufacture a user request.

Standing matters because agentic systems can generate work autonomously. They can notice anomalies, infer goals, create tickets, open cases, recommend investigations, or trigger workflows without a contemporaneous human request. Some of that is desirable. But the origin of the action becomes part of the authority record. A system acting proactively under statutory mandate is different from a system acting because a model inferred that intervention would be useful.

The standing check therefore asks what event legitimately activated the institution's power. It may be an application, complaint, warrant, scheduled obligation, sensor threshold, statutory duty, emergency declaration, contract event, or authorized automated trigger. The answer constrains what follows. If the initiating event is invalid, downstream approvals should not wash the defect away.

This is especially important in multi-agent systems because one agent can create the appearance of demand for another. A planning agent may open a task that an execution agent treats as authorized work. Without provenance, the system can convert internally generated intent into externally consequential authority. That is self-authorization by workflow composition.

Standing therefore belongs before permission in the chain. The system must know not only who can perform the action, but why the institution is entitled to consider performing it in the first place.

> **Before asking whether an actor may execute, ask why the institution may act at all.**

89. Counterfactual Audit: Could the System Explain Why It Did Not Act?

Accountability usually examines executed actions. A complete authority architecture also needs to explain consequential non-action. The impunity side of the model exists precisely because institutions can possess information yet fail to convert it into proportional action. If the system records only successful execution, it cannot distinguish legitimate restraint from dropped responsibility.

A counterfactual audit asks whether the institution can reconstruct why an action did not occur. Was authority absent? Was evidence insufficient? Did jurisdiction fail? Did a reviewer deny? Did the request expire? Was the system unavailable? Did an agent silently abandon the task? Each answer has different governance implications.

This is not a demand to log every hypothetical action. The scope should be consequential proposals that crossed a defined decision boundary. Once a system has generated or received a proposal serious enough to evaluate, the disposition of that proposal becomes part of institutional memory. Permission without execution, a hold without resolution, and refusal without reason are all potentially meaningful states.

Counterfactual audit also improves falsifiability. If the theory predicts that excessive friction produces impunity, the record needs evidence of where proposed or required actions stopped. Without that evidence, researchers can see executed power but not failed responsibility. The same asymmetry affects operational governance: institutions optimize what they can measure.

The authority receipt should therefore preserve negative decisions and unresolved states where proportional. A system that can prove why it acted but cannot prove why it failed to act has only half an accountability architecture.

> **Accountability must explain consequential non-action as well as action.**

90. Authority Availability: Governance Must Survive Outages

An authorization layer that disappears during an outage creates a dangerous choice. The institution can fail closed and stop legitimate work, or fail open and allow capability to substitute for authority. Both may be unacceptable in critical services. Availability is therefore part of authority architecture, not merely an infrastructure concern.

The design needs explicit degraded modes. Some actions may remain permitted under cached, short-lived authority. Others may require manual fallback. Some may be prohibited until authoritative state can be restored. Emergency break-glass access may exist, but it should create narrower, more visible, and more heavily audited authority rather than an unlogged bypass.

This is another domain where consequence should determine friction. A temporary inability to verify a low-risk read may be handled differently from inability to verify authority for an irreversible transfer or physical act. The institution should pre-classify these conditions rather than asking operators to invent policy during the outage.

Availability also tests the distinction between policy and protocol. A policy can say 'authorization is required' while the implementation routes around the control whenever the service is slow. The operational architecture has to make the failure behavior part of the authorization contract. Otherwise reliability pressure will quietly become authority expansion.

The governing question is not whether the authority service achieves perfect uptime. It is whether the institution knows what authority means when normal resolution is unavailable. Machine-speed governance has to specify the exceptional path before the exception occurs.

An outage may degrade capability; it must not silently expand authority.

91. The Authority Control Plane

The preceding sections imply a control plane distinct from the systems that propose and execute actions. Its function is not to centralize every institutional decision. It is to maintain the machinery by which authority can be resolved consistently across distributed systems: authoritative policy references, delegation state, revocation, purpose rules, jurisdictional mappings, consequence classes, approval requirements, and evidence needed for reconstruction.

A control plane can be centralized, federated, or hybrid. Estonia's distributed interoperability architecture demonstrates that high resolution capability does not require centralized possession of all underlying data. The same principle can apply to authority. A central service may coordinate policy metadata while authoritative institutions retain control of their own decisions. What matters is that the execution path can resolve the required authority without forcing every institution into one database or one sovereign decision-maker.

The control plane should not become a super-authority. It cannot legitimately grant powers that its sources do not possess. Its job is resolution and enforcement of authority supplied by valid institutional artifacts. This is analogous to the distinction between an identity provider and the legal rights of the person identified. Resolution infrastructure can make authority legible; it does not create the authority.

NIST's agent identity work again provides a useful external boundary. Its focus on standards-based identification, authorization, auditing, and non-repudiation for agents shows the emerging need for shared control mechanisms as agents move from generating outputs to taking actions. [S42] The paper's contribution is to extend that technical concern into the institutional chain: identity, standing, delegation, purpose, jurisdiction, time, consequence, and evidence all have to survive to the commit.

A mature authority control plane therefore behaves less like a universal policy brain and more like institutional connective tissue. It makes boundaries executable without erasing the institutions that own them.

> **The authority control plane resolves institutional boundaries; it does not own them.**

92. Claim Ceiling Before Engineering Convergence

The research has moved from historical reconstruction and comparative testing into an architectural model. That creates a new overclaim risk. Contemporary standards bodies and governments now discuss agent identity, authorization, action spaces, pre-deployment testing, continuous monitoring, and reversible automated action. None of that independently validates the full theory.

The record does not support that statement. NIST's 2026 agent identity project is a concept and standards exploration, not an endorsement of this paper's authority model. [S42] OECD's discussion of agentic government describes governance needs around action spaces, approval points, traceability, user control, and pause or reversal, but it does not prescribe the execution-boundary architecture developed here. [S32] GAO's federal AI work documents requirements for pre-deployment testing, impact assessment, monitoring, and discontinuation of non-compliant high-impact functionality; those are important control surfaces, not proof of a universal authorization theorem. [S38]

The strongest surviving claim is narrower and more defensible. Multiple contemporary governance efforts are independently encountering pieces of the same boundary because systems are moving from information processing toward autonomous or semi-autonomous action. The convergence is functional: once software can act across tools and workflows, identity, authorization, traceability, oversight, revocation, and redress become runtime concerns rather than policy documents alone.

There is also direct counterevidence against technological inevitability. GAO reports that the Small Business Administration paused most AI use while revising policy and continued only limited pilot activity as of April 2026. [S43] Capability can be stopped. Institutions can choose to delay deployment. The ratchet is therefore a tendency shaped by incentives and path dependence, not a law that makes governance powerless.

That is the correct point at which to introduce an engineering response. The history identifies a recurring boundary. The comparative cases show both failure and control. The contemporary record shows institutions beginning to operationalize pieces of authorization and oversight. Engineering must still prove that any proposed implementation resolves the boundary correctly.

> **Convergence identifies the boundary. It does not certify the implementation.**

PART IX

ENGINEERING CONVERGENCE AND CLAIM CEILING

92A. Claim Tiers Before Engineering Convergence

Historically supported tier

The documentary record supports the narrower historical propositions that information technologies repeatedly change the cost, latency, scale, and organizational structure of institutional action; that public and private infrastructure often combine in delivering capability; that both excessive fragmentation and insufficient control can produce recognized institutional failures; and that high interoperability does not logically require centralized possession.

Theoretical tier

Administrative compression, the abuse-impunity symmetry, selective friction, authority debt, authority laundering, authority gradients, infrastructure memory, and H1-H7 are theoretical constructs generated by the synthesis. They remain useful only to the extent that future evidence can disconfirm them.

Engineering-prescriptive tier

Commitment boundaries, authority objects, provenance, hold and modification semantics, temporal revalidation, delegation monotonicity, receipts, quorum, kill paths, and minimum sufficient authority are engineering prescriptions. Contemporary standards and public-sector guidance make the problem timely, but they do not establish this protocol as optimal. Its effectiveness remains an empirical question.

93. From Institutional Research to Engineering

The historical record above does not validate a commercial product. A commercial system is not independent evidence for the historical thesis, and the historical thesis is not evidence that any particular implementation is sufficient.

The engineering implication follows from the institutional distinction developed throughout the paper: a proposed machine action cannot supply its own authority merely by being useful or technically possible. A model may propose an action, but the institution must establish the basis on which a consequential action may proceed.

The historical research reaches the same boundary from the opposite direction. For decades, institutions have reduced the natural friction separating knowledge from action. As that friction disappears, authorization can no longer remain an implicit assumption carried by organizational slowness, manual process, or the mere possession of credentials. It has to become explicit infrastructure.

This does not mean every institution needs the same product, policy language, or enforcement mechanism. The portable proposition is narrower: consequential state transitions should not be self-authorizing, and the evidence of authority should survive the action it permitted.

> **When technology removes natural constraints on action, legitimate systems must replace accidental friction with explicit, enforceable authorization boundaries.**

93.1 From research finding to engineering response

The engineering discussion follows the historical and comparative argument rather than validating it. Several historical cases predate the modern AI industry and even the internet. The international cases test whether the American security history produces a parochial explanation. The control cases challenge the theory. Their evidentiary function does not depend on the commercial availability of any particular implementation.

What survives is a systems requirement: a proposed action needs a boundary at which authority is established before execution. In an AI system, that boundary cannot safely be delegated back to the same probabilistic component proposing the action. In a distributed agent system, it also cannot be assumed to persist merely because authentication succeeded earlier in the chain.

That does not make any implementation historically inevitable or commercially sufficient. It does explain why the engineering problem is larger than content filtering or model safety. The problem is state transition under delegated machine capability.

The portable requirement is not a brand. It is a separation: the system that can propose or perform an action should encounter an independently enforceable authority boundary before the consequential state change occurs, and the evidence of that authorization should survive afterward.

> **The model proposes. Authority is resolved outside the model. Execution follows the authority, not the capability.**

93.2 The implementation claim ceiling

The research supports an engineering requirement, not a product certification. It supports separating capability from authority, binding consequential actions to explicit authorization, and preserving evidence sufficient to reconstruct the decision. It does not establish that one enforcement architecture is complete, that deterministic policy can resolve every legal ambiguity, or that a pre-execution gate eliminates the need for downstream monitoring, appeal, incident response, or institutional oversight.

That distinction protects the independence of the research and its technical evaluation. If an implementation is later shown to enforce one boundary poorly, the historical thesis does not automatically fall. If the historical model is narrowed by new evidence, every implementation still has to stand on its own technical evaluation. The relationship is an engineering response to a research question, not circular validation.

> **The history identifies the boundary. Engineering still has to prove the implementation.**

94. What This Paper Does Not Establish

It does not establish a single covert technological lineage from PROMIS to later NSA systems, commercial cloud, Flock, or artificial intelligence.

It does not establish that private contractors or hyperscale cloud providers have captured the American state. It establishes recurring conditions of dependency that can increase capture risk and weaken practical autonomy if not governed.

It does not establish that interoperability is inherently centralizing or abusive. Estonia is direct counterevidence.

It does not establish that fear is the root cause of digital administration. The global comparison indicates that ordinary administrative compression persists without emergency.

It does not establish that every form of friction is desirable. Much administrative friction is waste, delay, exclusion, and error. The distinction is between obsolete computational friction and necessary governance boundaries.

It does not establish that human review is always the correct control. Human review can be slow, biased, inconsistent, or ceremonial. The requirement is valid authority and accountability, not ritual human presence.

It does not establish that artificial intelligence is uniquely dangerous. AI matters here because it compresses interpretation and, through agents, begins to compress execution.

It does not establish that any one company or product is the only implementation capable of enforcing the identified boundary.

> **If a later record defeats a claim in this paper, the claim should fall. The architecture does not get to outrank the evidence.**

95. Conclusion: From Accidental Friction to Engineered Authority

The history reconstructed here is not a history of one technology. It is a history of disappearing bottlenecks.

Paper externalized memory. Databases compressed retrieval. Relational systems compressed resolution. Networks compressed distance. Interoperability compressed institutional boundaries. Cloud compressed infrastructure. Sensors compressed observation. Artificial intelligence compresses interpretation. Agents begin to compress execution.

Each transition can produce real public value. Each can also move an institution closer to the point at which technical ability is no longer meaningfully constrained by the natural limits of human administration.

That is the boundary that matters.

For most of institutional history, inability did some of the work of governance. The file could not be found quickly. The databases could not communicate. The analyst could not process everything. The infrastructure could not scale instantly. A person had to carry the decision across the final boundary. Those limitations were not legitimate sources of authority. They simply made the difference between ability and authority harder to violate at scale.

We are removing them.

The answer is not to rebuild the filing cabinet, the silo, or the disconnected institution. The answer is to stop confusing speed with legitimacy and access with permission.

The machine can know more without being entitled to act more. The institution can integrate information without dissolving jurisdiction. The state can become more efficient without making every administrative boundary disappear. Private infrastructure can supply capability without becoming the source of sovereign authority. Artificial intelligence can interpret without authorizing itself.

The final proposition is therefore smaller than the history that produced it.

> **Ability is not authority.**

As technical systems remove the accidental friction that once separated institutional knowing from institutional action, legitimate authority has to become a deliberate property of the architecture.

> **That is the architecture of authority.**

Appendix A. Capability-Authority Ledger

Case	Capability	Primary bottleneck reduced	Authority risk	Control/correction	Standing
SHAMROCK/ MINARET	Large-scale communications collection	Observation/ collection	Secret collection and weak oversight	Church investigations; later statutory/oversight reforms	A/B
PROMIS/INSLAW	Searchable prosecutorial case management	Retrieval/resolution	Software ownership, procurement, technical dependence	Litigation, congressional and DOJ investigations	Mixed A-E
Pre-9/11 wall	Compartmented intelligence/criminal systems	Handoff constrained	Legitimate information failed to flow	Post-9/11 information-sharing reforms	A
Fusion centers/ISE	Cross-level information exchange	Institutional fragmentation	Over-sharing, overlap, privacy, unclear performance	Privacy policies, capability/performance assessment	A/B
IC contractors	Private technical/analytic workforce	Internal capacity	Epistemic dependence; proximity to sovereign decisions	Inherently-governmental/critical-function rules; oversight	A
C2S/C2E/JWCC	Classified commercial hyperscale cloud	Infrastructure scale	Vendor/sector dependency	Competition, multi-cloud, procurement review	A/B
Flock ALPR	Networked vehicle-location resolution	Local observation/resolution	Query misuse; weak local auditing	Case-number fields, anomaly review, audits, state/local restrictions	A/B
Estonia X-tee	Distributed authenticated data exchange	Institutional interoperability	Cross-system access risk	Authorization, encryption, signatures, logs, distributed ownership	A
Global DPI	Identity, payments, secure data exchange	Administrative fragmentation	Scale, exclusion, privacy, lock-in	Country-specific legal/technical safeguards	A/B

Case	Capability	Primary bottleneck reduced	Authority risk	Control/correction	Standing
Government AI	Machine-assisted interpretation	Human analytical bandwidth	Opacity, target formation, automation bias	Risk/impact assessment, transparency, oversight	A/B
Agentic systems	Tool-mediated machine execution	Human execution bandwidth	Capability silently becomes permission	Explicit pre-execution authorization boundary	Emerging/E

The expanded ledger below forces each case into the same analytical fields. It is intentionally reductive: the purpose is not to erase legal or historical differences, but to prevent narrative similarity from doing evidentiary work. A case supports the theory only to the extent that the relevant capability, authority basis, friction, execution path, auditability, and governance result can actually be identified.

Case	Capability	Actor	Authority basis	Friction / control	Execution	Auditability	Observed governance result
SHAMROCK	Bulk international telegram access	NSA + cooperating carriers	Intelligence collection authorities of period	Physical carrier custody; manual processing	Collection/ analysis	Limited by modern standards	Exposure produced reform
PROMIS	Cross-case prosecutorial resolution	DOJ/EOUSA users	Prosecutorial and administrative authority	Fragmented records; local systems	Case information processing	Disputed institutional record	Resolution capability; ownership dispute
Pre-9/11 wall	Collected intelligence with constrained sharing	FBI/IC components	Separate intelligence/law-enforcement authorities	Legal/procedural sharing barriers	Handoff often failed	Fragmented	Impunity/action-gap control case
Post-9/11 ISE	Cross-agency information sharing	Federal/state/ local/private participants	Program- and mission-specific authorities	Need-to-know boundaries reduced	Networked sharing	Expanded logging/policy	Capability increased; new governance surface

Case	Capability	Actor	Authority basis	Friction / control	Execution	Auditability	Observed governance result
Section 702 / Zazi	Foreign-intelligence collection producing domestic lead	NSA/FBI and providers	Statutory surveillance + investigative authority	Selector, targeting, minimization, handoff	Investigation and prosecution	Formal oversight exists	Successful governed handoff example
Snowden contractor access	Privileged access to classified systems	Contractor personnel	Contractual/ clearance/role authority	Compartmentation and access controls	System administration/ access	Extensive but imperfect	Dependency and insider-risk case
Flock ALPR	Mass distributed plate resolution	Local agencies using private platform	Local investigative authority	Query fields, policy, case numbers, audits	Search/ investigation	Vendor + agency logs vary	Capability compression + misuse exposure
Estonia X-tee	Distributed cross-system resolution	Independent public/private data owners	Service-specific legal authority	Federated control; signed/encrypted exchange	Data exchange	High transaction logging	High capability without central possession
JWCC	Multi-cloud classified compute	DoD + commercial cloud providers	Procurement + mission authority	Multi-vendor architecture, classifications	Compute/ storage provisioning	Contract/ security controls	Dependency mitigated rather than eliminated
OECD government AI	Automated analysis/recommendation/action	National public bodies	Varies by jurisdiction	Risk assessment, review, audit uneven	AI-supported administration	Uneven across countries	Governance lag measurable by control surface
SBA AI pause 2025-26	Potential AI-supported administration	SBA	Agency policy and federal requirements	Deployment pause during policy revision	Most AI use stopped	Inventory/ reporting gaps noted	Counterexample to inevitable ratchet
FedRAMP 20x	Cloud authorization evidence	Federal agencies/cloud providers	Federal cloud authorization regime	Machine-readable authorization artifacts	Initial/ continuing authorization	Designed for persistent evidence	Authorization becoming computational object

Case	Capability	Actor	Authority basis	Friction / control	Execution	Auditability	Observed governance result
NIST agent identity project	Agent access/actions across tools	Enterprise software/AI agents	Organizational IAM + delegated authority	Identity/ authorization/ audit controls	Tool and application actions	Audit/non-repudiation focus	Emerging standards problem
VA EHR modernization	Enterprise clinical record execution	VA + commercial system	Clinical/admin authority	Deployment gates, operational testing, user feedback	Clinical/admin workflows	GAO oversight + recommendations	Capability halted when controls/performance failed
Council of Europe AI Convention	AI lifecycle activity	Public/private actors in scope	Treaty/domestic implementation	Risk, oversight, documentation, remedies	AI-mediated decisions/actions	Procedural safeguards	Rights/control framework preceding some deployments
Physical AI / robotics	Digital decision becomes physical state change	Operator + autonomous system	Mission/safety/ operational authority	Safety envelope, geofence, interlocks	Physical actuation	Telemetry + command record	Irreversibility raises authority threshold
Multi-agent delegation	Composed tool capability	Principal + parent/child agents	Delegated organizational authority	Scope/time/ purpose/ redelegation limits	Cross-tool workflow	Needs provenance chain	Tests authority monotonicity
Emergency authority	Expanded temporary capability	Executive/ agency/ emergency operators	Emergency statute/order	Sunset, scope, review	Exceptional execution	Must preserve expiry/override record	Tests capability persistence after crisis
Procurement lock-in	Operational dependency on vendor stack	Agency + contractor	Contract/ procurement authority	Data rights, exit terms, expertise	Platform operation/change	Contract/audit artifacts	Reversibility is governance variable
Execution-boundary model	Any consequential proposed action	Authenticated actor/agent	Resolved external authority	Pre-execution authority decision	State transition	Retained authority evidence	Normative/ engineering hypothesis, not historical fact

Appendix B. Falsifier Matrix

The theory is weakened if high-capability systems consistently produce good outcomes without explicit or reconstructable authority boundaries.

The theory is weakened if administrative compression does not appear across political systems with materially different ideologies and security conditions.

The theory is weakened if capability commonly disappears when the crisis that justified it ends rather than becoming part of the inherited baseline.

The theory is weakened if government can remain operationally autonomous despite sustained loss of internal technical capacity and dependence on non-replaceable private infrastructure.

The theory is weakened if machine-speed interpretation and execution do not materially change the location or importance of authorization boundaries.

The theory is strengthened, but not proven, when failures cluster around high capability, diffuse authority, weak auditability, unclear ownership, and low internal capacity, while successful cases cluster around explicit scope, reliable handoff, valid authority, retained expertise, and reconstructable execution.

Appendix C. Adversarial Case Matrix

Ordinary modernization: If speed/cost improvement explains the case without changing the practical authority surface, the case should not be counted as evidence for the execution-boundary thesis.

Law-alone control: Identify systems operating at high speed and consequence where external law/policy alone reliably constrains execution without technical authorization or durable audit. Repeated success would weaken H2 and H4.

Low-tech abuse control: Cases of serious abuse without advanced information systems demonstrate that technology is neither necessary nor sufficient for illegitimate power; they bound the theory to information-mediated scale and latency.

Pre-governed deployment control: Identify AI, surveillance, or data-exchange systems where operational safeguards preceded deployment. Compare outcomes with retrofit-governance cases rather than treating lag as universal.

Reversible dependency control: Identify major government technology dependencies successfully exited without material mission loss. Such cases clarify which combinations of data rights, internal skills, portability, and competition preserve autonomy.

Sunset control: Identify emergency capabilities genuinely dismantled after the precipitating crisis, including infrastructure and institutional assumptions. These cases test the strength of the ratchet mechanism.

Federated-resolution control: Compare centralized possession architectures with federated exchange architectures such as X-tee to test whether equivalent resolution can be achieved with different concentration and audit risk.

A disconfirming case is evidence, not an exception to be explained away.

Appendix D. Source Register

Authorization Engineering and Agent Security Prior Art

S68. Brooker, Marc, Joseph Tassarotti, and Jean-Baptiste Tristan. “Introducing Dogwood: runtime verification for AI agents.” AWS Open Source Blog (Aug. 6, 2026). https://aws.amazon.com/blogs/opensource/introducing-dogwood-runtime-verification-for-ai-agents/

S67. Zhu, L., and Z. Berg. “Delegated Refresh Tokens for OAuth 2.0 Token Exchange.” draft-zhu-oauth-async-delegation-05, Internet-Draft, Aug. 3, 2026 (work in progress). https://datatracker.ietf.org/doc/draft-zhu-oauth-async-delegation/

S66. Tulshibagwale, A., et al. “Transaction Tokens.” draft-ietf-oauth-transaction-tokens-11, IETF OAuth Working Group Internet-Draft (July 30, 2026; work in progress). https://datatracker.ietf.org/doc/draft-ietf-oauth-transaction-tokens/

S65. OpenID Foundation AuthZEN Working Group. “Authorization API 1.0.” Final Specification (Jan. 2026). https://openid.net/specs/authorization-api-1_0-final.html

S64. Racioppi, Giovanni. “Mandato: Protocol-Level Enforcement of Digitally Signed Mandates on AI Agent Actions with Cryptographically Chained Audit Trails.” arXiv:2608.14074 (2026 preprint). https://arxiv.org/abs/2608.14074

S63. Ibrahim, Amjad, and Yong Li. “Overlaying Governance: A Compositional Authorization Framework for Delegation and Scope in Agentic AI.” arXiv:2606.03518 (2026 preprint). https://arxiv.org/abs/2606.03518

S62. Prakash, Sunil. “AIP: Agent Identity Protocol for Verifiable Delegation Across MCP and A2A.” arXiv:2603.24775 (2026 preprint). https://arxiv.org/abs/2603.24775

S61. Kim, Juhee, Wenbo Guo, and Dawn Song. “SoK: Attack and Defense Landscape of Agentic AI Systems.” 35th USENIX Security Symposium (USENIX Security 26), 6047-6066 (2026). https://www.usenix.org/conference/usenixsecurity26/presentation/kim-juhee-agentic

S60. Disselkoen, Craig, et al. “How We Built Cedar: A Verification-Guided Approach.” Conference on the Foundations of Software Engineering (FSE 2024). https://www.amazon.science/publications/how-we-built-cedar-a-verification-guided-approach

S59. Cutler, Joseph W., et al. “Cedar: A New Language for Expressive, Fast, Safe, and Analyzable Authorization.” Proceedings of the ACM on Programming Languages 8, OOPSLA1, Article 118 (2024). https://doi.org/10.1145/3649835

S58. Lodderstedt, Torsten, Justin Richer, and Brian Campbell. “OAuth 2.0 Rich Authorization Requests.” RFC 9396, Internet Engineering Task Force (May 2023). https://www.rfc-editor.org/rfc/rfc9396.html

S57. Chandramouli, Ramaswamy, and Zack Butcher. NIST SP 800-207A, “A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments” (2023). https://csrc.nist.gov/pubs/sp/800/207/a/final

S56. Rose, Scott, Oliver Borchert, Stu Mitchell, and Sean Connelly. NIST SP 800-207, “Zero Trust Architecture” (2020). https://doi.org/10.6028/NIST.SP.800-207

S55. Andersen, Michael P., et al. “WAVE: A Decentralized Authorization Framework with Transitive Delegation.” 28th USENIX Security Symposium (USENIX Security 19) (2019). https://www.usenix.org/conference/usenixsecurity19/presentation/andersen

S54. Pang, Ruoming, et al. “Zanzibar: Google’s Consistent, Global Authorization System.” 2019 USENIX Annual Technical Conference (USENIX ATC 19) (2019). https://www.usenix.org/conference/atc19/presentation/pang

Scholarly Literature and Prior Art

S53. Moynihan, Donald, Pamela Herd, and Hope Harvey. “Administrative Burden: Learning, Psychological, and Compliance Costs in Citizen-State Interactions.” Journal of Public Administration Research and Theory 25(1): 43-69 (2015). https://doi.org/10.1093/jopart/muu009

S52. Busuioc, Madalina. “Accountable Artificial Intelligence: Holding Algorithms to Account.” Public Administration Review 81(5): 825-836 (2021). https://doi.org/10.1111/puar.13293

S51. Raji, Inioluwa Deborah, et al. “Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing.” FAT* 2020, 33-44. https://doi.org/10.1145/3351095.3372873

S50. Wieringa, Maranke. “What to Account for When Accounting for Algorithms: A Systematic Literature Review on Algorithmic Accountability.” FAT* 2020, 1-18. https://doi.org/10.1145/3351095.3372833

S49. Selbst, Andrew D., danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet Vertesi. “Fairness and Abstraction in Sociotechnical Systems.” FAT* 2019, 59-68. https://doi.org/10.1145/3287560.3287598

S48. Yeung, Karen. “Algorithmic Regulation: A Critical Interrogation.” Regulation & Governance 12(4): 505-523 (2018). https://doi.org/10.1111/rego.12158

S47. Kroll, Joshua A., Joanna Huey, Solon Barocas, Edward W. Felten, Joel R. Reidenberg, David G. Robinson, and Harlan Yu. “Accountable Algorithms.” University of Pennsylvania Law Review 165(3): 633-705 (2017).

S46. Citron, Danielle Keats. “Technological Due Process.” Washington University Law Review 85(6): 1249-1313 (2008).

S45. Bovens, Mark, and Stavros Zouridis. “From Street-Level to System-Level Bureaucracies: How Information and Communication Technology is Transforming Administrative Discretion and Constitutional Control.” Public Administration Review 62(2): 174-184 (2002). https://doi.org/10.1111/0033-3352.00168

S44. GAO-26-109393, VA Electronic Health Record Modernization: Actions Needed to Sustain Accelerated System Deployments (Sept. 2, 2026). https://www.gao.gov/products/gao-26-109393

S43. GAO-26-107828, Artificial Intelligence: Uses and Risks for Small Business Contracting and Innovation Research (2026). https://www.gao.gov/products/gao-26-107828

S42. NIST NCCoE, Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization (Concept Paper, Feb. 5, 2026). https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd

S41. FedRAMP, Initial Outcome from RFC-0024 Rev5 Machine-Readable Packages (2026). https://www.fedramp.gov/notices/0009/

Agent Identity, Authorization, and Control Sources

S40. GAO-26-107681, Artificial Intelligence: OMB Action Needed to Address Privacy-Related Gaps in Federal Guidance (2026). https://www.gao.gov/products/gao-26-107681

S39. GAO-24-105980, Artificial Intelligence: Agencies Have Begun Implementation but Need to Complete Key Requirements. https://www.gao.gov/products/gao-24-105980

S38. GAO-25-107933, Artificial Intelligence: Federal Efforts Guided by Requirements and Advisory Groups (2025). https://www.gao.gov/products/gao-25-107933

S37. GAO-25-107653, Artificial Intelligence: Generative AI Use and Management at Federal Agencies (2025). https://www.gao.gov/products/gao-25-107653

S36. GAO-26-108443, Cybersecurity: Selected Agencies Need to Better Protect Cloud Data (2026). https://files.gao.gov/reports/GAO-26-108443/index.html

S35. GAO-26-107859, Artificial Intelligence Acquisitions: Agencies Should Collect and Apply Lessons Learned to Improve Future Procurements (2026). https://www.gao.gov/products/gao-26-107859

S34. Council of Europe, HUDERIA risk and impact assessment methodology. https://www.coe.int/en/web/artificial-intelligence/huderia-risk-and-impact-assessment-of-ai-systems

S33. OECD, Digital Government Outlook 2026, Governing Digital Investment and Capabilities to Deliver at Scale. https://www.oecd.org/en/publications/digital-government-outlook_0496b2bc-en/full-report/governing-digital-investment-and-capabilities-to-deliver-at-scale_d775f3d4.html

S32. OECD, Digital Government Outlook 2026, Digital Governments at a Turning Point; agentic AI in government. https://www.oecd.org/en/publications/digital-government-outlook_0496b2bc-en/full-report/digital-governments-at-a-turning-point_0491aad4.html

S31. GAO-10-513, Information Security: Federal Guidance Needed to Address Control Issues with Implementing Cloud Computing (2010). https://www.gao.gov/products/gao-10-513

S30. OECD, Digital Government Outlook 2026, Adopting and Governing AI in Government. https://www.oecd.org/en/publications/2026/06/digital-government-outlook_4585678e/full-report/adopting-and-governing-ai-in-government_7ef312a9.html

Digital Government and AI Governance Sources

S29. GAO-25-107468, Weapon System Sustainment: DOD Can Improve Planning and Management of Data Rights (2025). https://www.gao.gov/products/gao-25-107468

S28. GAO-25-107114, Cloud Computing: Selected Agencies Need to Implement Updated Guidance for Managing Restrictive Licenses (2024). https://www.gao.gov/products/gao-25-107114

S27. GAO-23-106247, Cloud Computing: DOD Needs to Improve Tracking of Data User Fees (2023). https://www.gao.gov/products/gao-23-106247

S26. GAO-26-107530, Cloud Computing: Federal Government Needs to Address Procurement Challenges (2026). https://www.gao.gov/products/gao-26-107530

S25. U.S. Department of Defense, Department Names Vendors to Provide Joint Warfighting Cloud Capability (Dec. 12, 2022). https://www.defense.gov/News/News-Stories/Article/Article/3243483/department-names-vendors-to-provide-joint-warfighting-cloud-capability/

S24. U.S. Department of Defense, Joint Warfighting Cloud Capability press briefing (March 2022). https://www.defense.gov/News/Transcripts/Transcript/Article/2982986/dod-officials-hold-a-press-briefing-on-the-joint-warfighting-cloud-capability-p/

S23. GAO-10-833, Federal Contracting: Opportunities Exist to Increase Competition and Assess Reasons When Only One Offer Is Received (2010). https://www.gao.gov/products/gao-10-833

Procurement, Cloud, and Infrastructure Sources

S22. Council of Europe, Framework Convention on Artificial Intelligence and Human Rights, Democracy and the Rule of Law. https://www.coe.int/en/web/artificial-intelligence/the-framework-convention-on-artificial-intelligence

S21. GAO-10-972, Fusion Center capabilities, privacy and performance measurement. https://www.gao.gov/products/gao-10-972

S20. GAO-15-155, Fusion Center capabilities and results. https://www.gao.gov/products/gao-15-155

S19. U.S. Department of Justice, OPR statement on Epstein 2006-2008 investigation. https://www.justice.gov/archives/opa/pr/statement-doj-office-professional-responsibility-report-jeffrey-epstein-2006-2008

S18. OECD, Digital Government Outlook 2026. https://www.oecd.org/en/publications/digital-government-outlook_0496b2bc-en.html

S17. Estonian Information System Authority, X-tee data exchange layer. https://www.ria.ee/en/state-information-system/data-exchange-platforms/data-exchange-layer-x-tee

S16. State Council of the People's Republic of China, guidance on digital government and integrated government data. https://www.nia.gov.cn/n741440/n741547/c1562629/content.html

S15. World Bank, Global Digital Public Infrastructure Program: From Foundations to Scale (2026). https://www.worldbank.org/en/results/2026/05/06/global-digital-public-infrastructure-program

S14. Washington Post, police-department audit failures and ALPR misuse, 19 Aug 2026. https://www.washingtonpost.com/technology/2026/08/19/we-found-cops-who-misused-flock-their-police-departments-didnt-know/

S13. Washington Post, Flock governance changes after misuse reporting, 13 Aug 2026. https://www.washingtonpost.com/technology/2026/08/13/flock-unveils-changes-after-post-report-police-misuse-its-network/

S12. Reuters, Florida ALPR restrictions and Flock network scale, 3 Sep 2026. https://www.reuters.com/legal/government/florida-bans-highway-license-plate-readers-backlash-over-surveillance-spreads-2026-09-03/

S11. European Union, Interoperable Europe Act. https://interoperable-europe.ec.europa.eu/Interoperable-Europe-Act-Regulation

S10. GAO B-407073.3 et al., IBM-U.S. Federal protest of CIA C2S award. https://www.gao.gov/products/b-407073.3%2Cb-407073.4%2Cb-407073.5%2Cb-407073.6

S9. GAO-07-990 / GAO-08-142T, contractors near inherently governmental functions. https://www.gao.gov/products/gao-07-990

S8. GAO-10-744T, government/contractor workforce mix. https://www.gao.gov/assets/a124728.html

S7. GAO-08-572T, DOD reliance on contractors. https://www.gao.gov/assets/a119340.html

S6. 9/11 Commission, Chapter 13, networked information sharing and audit trail. https://www.9-11commission.gov/report/911Report_Ch13.htm

S5. 9/11 Commission Executive Summary, information-sharing recommendations. https://9-11commission.gov/report/911Report_Exec.htm

S4. National Commission on Terrorist Attacks Upon the United States, Chapter 3 ('the wall'). https://9-11commission.gov/report/911Report_Ch3.htm

S3. U.S. Department of Justice, 1994 review of INSLAW/Bua findings. https://www.justice.gov/archive/opa/pr/Pre_96/September94/555.txt.html

S2. U.S. Senate intelligence record discussing SHAMROCK and telegraph-company cooperation. https://www.intelligence.senate.gov/wp-content/uploads/2024/08/sites-default-files-hearings-110399.pdf

S1. U.S. Senate, Church Committee historical overview. https://www.senate.gov/about/powers-procedures/investigations/church-committee.htm

Historical, Intelligence, and Institutional Sources

Source-handling reference for this external edition

SH1. Chatham House. “Code of conduct,” definition and application of the Chatham House Rule. Consulted 14 September 2026. https://www.chathamhouse.org/events/code-conduct

Appendix E. Research Continuity and Source Boundary

The historical and comparative thesis stands or falls on the external record. Conceptual vocabulary in the engineering discussion does not turn private development work into independent research evidence. Non-public discussions, unpublished implementation materials, and customer or partner experience are not offered as validation of the public claims. The external edition preserves the distinction between an institutional research proposition and the technical evidence needed to evaluate a particular implementation.

The publication and source-handling notice governs protected source interactions. Public sources remain attributable. The identity and affiliation of protected participants must not be inferred from an example, a source description, or the author's commercial affiliation. This edition does not make a determination about patent inventorship or intellectual-property ownership, and publication does not grant a license to non-public technology or patent rights. A commercial implementation cannot validate the institutional history that motivates it, and the history cannot certify the implementation.

> **The claim may not exceed its artifact.**

Appendix F. Authority Questions and Research Boundaries

This appendix preserves the institutional questions raised by the paper's authority model. It is a conceptual discussion, not a product specification or normative protocol. It does not prescribe a data schema, credential format, decision-code vocabulary, cryptographic construction, or conformance suite. Its purpose is to distinguish the authority questions that a consequential action raises and the evidence needed to study them without claiming that every domain must implement the same mechanism.

F.1 Proposed action

The unit of governance is the proposed state transition, not the model output that preceded it. Natural-language intent does not by itself establish permission for the consequence that follows. The institution needs to know what action is actually under consideration and whether a change in the proposal changes the authority question. A different destination, a broader disclosure, or a different physical consequence cannot be treated as equivalent merely because an agent is pursuing the same goal. This is a question about the relationship between intent and consequence, not a prescription for an action-record format.

F.2 Actor and principal

The actor performing an operation and the principal whose authority is being exercised need not be the same. A human employee, service account, contractor process, model, agent, or sub-agent may be the immediate actor while institutional authority originates elsewhere. The research question is whether that relationship remains understandable through delegation. Possession of a credential does not prove that its holder owns the institutional power being exercised. Conversely, public low-consequence activity need not acquire a stronger identification burden simply because more detailed machinery is available. The control should remain proportionate to the consequence.

F.3 Standing and trigger

Standing concerns the event that legitimately activated the institution's power. A user request, application, complaint, statutory duty, court order, emergency declaration, or contractual event can play that role. Standing is distinct from identity: an authenticated actor may lack standing, and a valid institutional duty may exist without a contemporaneous requester. The issue is whether internally generated workflow can be mistaken for an external mandate. It also concerns non-action. When a valid trigger enters the institution and no proportionate disposition follows, the failure deserves an accountable explanation. Disputed standing should remain disputed until the appropriate authority resolves it.

F.4 Delegation chain

Delegation permits another actor to exercise a limited authority; it does not erase the limits of the original grant. The substantive research question is whether a composed workflow remains within that authority when work passes between actors. Individually valid credentials do not necessarily establish that the combined institutional consequence was permitted. This matters when agents create additional workers or select new means to achieve a goal. A parent objective is not a grant for every useful subtask. The analytical distinction is between what the workflow seeks to accomplish and what the institution has actually authorized.

F.5 Purpose and jurisdiction

Purpose and jurisdiction constrain otherwise valid capability. The same actor may be permitted to use the same information or tool for one institutional purpose and prohibited from using it for another. Authority valid in one organizational, geographic, legal, or contractual domain does not cross automatically into another. Interoperability establishes that systems can communicate; it does not resolve the powers of their participants. Where applicable authorities conflict, the conflict remains an institutional question requiring a legitimate resolution. The absence of that resolution should not be treated as permission merely to keep a workflow moving.

F.6 Temporal validity and freshness

Authority is temporal. The research concerns whether the basis for an action remains effective when the action becomes consequential, rather than merely whether permission existed when work began. Long-running activity can cross a change in policy, role, consent, or institutional responsibility. Earlier permission may therefore be relevant evidence without being a complete answer to the later authority question. The temporal rule must come from the applicable institution and domain. This edition does not prescribe issuance fields, validity-window formats, cache invalidation, or revocation-propagation mechanisms.

F.7 Evidence references and claim ceiling

Evidence can support an authority decision without becoming infallible. Identity assertions, policies, orders, approvals, consent, case status, classifications, and other facts can all be material. Their significance depends on what they actually establish. The claim ceiling applies here as it does in the historical method: the decision cannot assert more authority than its supporting artifacts permit. Missing evidence remains missing, uncertainty remains visible, and a fluent model summary does not upgrade a weak artifact into a strong one. Later review requires an account of what the institution relied upon, not merely a plausible explanation generated afterward.

F.8 Distinguishing institutional dispositions

Permission, prohibition, an unresolved question, and permission for a narrower action are not the same institutional state. Their differences matter because a premature rejection can strand legitimate work while premature execution can exceed authority. An institution may need additional evidence, a responsible decision-maker, or a change in the proposed action before the matter can be resolved. The distinction is not a requirement to adopt one particular set of machine-readable labels. It is a requirement not to erase uncertainty or conditional permission for the convenience of an interface. Accountability should preserve why a proposal proceeded, changed, stopped, or remained unresolved.

F.9 Evidence and the action taken

An authority decision is useful to later review only when it can be related to the action the institution actually took. A record of permission for one consequence should not become a retrospective explanation for a materially different consequence. The same concern applies when an approved proposal never produces action. The research question is whether the institution can explain the relationship among its decision, the resulting action or non-action, and the evidence on which it relied. Integrity mechanisms can assist that task, but this edition does not prescribe transaction identifiers, sealing formats, replay controls, or a receipt schema.

F.9A Receipt privacy and selective disclosure

Receipt privacy is part of the institutional problem, not an exception to it. An evidence store should not become an unrestricted source of secondary analytics simply because the original decisions were documented correctly. Minimization, appropriate access, retention, evidentiary duties, and the rights of affected people remain relevant to the evidence itself. Integrity and confidentiality answer different questions: a record can be unchanged and still be disclosed or reused improperly. The public research claim concerns the need to examine both reconstructability and privacy; it does not prescribe the internal design of a receipt service.

F.10 Revocation, recovery, and remedy

Institutional responsibility continues after a decision. Authority can end before an action takes effect, execution can fail partially, a later review can find an error, and an affected person can challenge a result. These are ordinary questions of governance rather than peripheral exceptions. A corrective action may itself require authority, and an irreversible consequence cannot be treated as though restoring a record restores the world. The research therefore examines the connection between initial permission, the capacity to correct what follows, and preservation of the original decision history. Domain-specific implementation and recovery procedures remain outside this edition.

F.11 Testing and the limits of conformance

A reference model is useful only if systems can fail it. Evaluation must test the institutional claim actually made and distinguish a mechanism's behavior from the legitimacy of the rule it implements. Successful technical evaluation is not evidence that an underlying authority was lawful or that every factual predicate was correct. Failure can nevertheless show that a mechanism does not support its stated governance claim. This external edition does not prescribe a pass/fail conformance suite. Such a suite would require a separately specified implementation, operating context, and account of the evidence needed to establish each result.

F.12 Research tests against the protocol

The same reference model creates empirical questions for future work. Researchers can measure authority staleness, revocation propagation time, percentage of consequential actions with reconstructable authority evidence, frequency and duration of hold states, proportion of delegations with explicit expiry, percentage of high-consequence actions with independent approval, recovery time after erroneous execution, and the rate at which policy changes reach runtime enforcement. Cross-domain comparison could test whether systems with stronger pre-execution authority controls show lower rates of unauthorized action without increasing unresolved legitimate work. A negative result would matter. If explicit authority controls consistently add delay without improving authorization quality, the theory would need to narrow again. The objective is not to protect the framework from failure. It is to make failure observable.

F.13 Comparator baseline

The Authority Protocol should be evaluated against credible existing authorization architectures rather than against a straw-man system with no controls. Minimum comparators include a conventional identity-and-scope bearer-token design; a zero-trust PDP/PEP architecture; relationship-based authorization in the Zanzibar family; policy-as-code using Cedar or a comparable engine; OAuth RAR for structured fine-grained authorization; AuthZEN for interoperable PDP/PEP communication; Transaction Tokens for authorization-context continuity across call chains; a delegated-capability design such as WAVE; long-running delegation mechanisms such as delegated refresh tokens; and agent tool-boundary controls such as Dogwood where implementations are available. [S54][S55][S56][S58][S59][S65][S66][S67][S68]

The comparison should separate primitive capability from institutional coverage. If an existing system can enforce scope, expiry, revocation, delegation, transaction binding, context propagation, or tool-call gating as well as the proposed protocol, that result should be credited to the comparator rather than relabeled as unique authority behavior. The research question is whether adding standing, lawful trigger, jurisdiction, purpose, obligation, hold semantics, non-action closure, remedy, and institutionally legible receipts produces measurable value beyond mature authorization controls.

F.14 Validation boundary

A future implementation should identify what it claims to enforce before performance benchmarking. Mechanically testable claims must remain separate from questions of legal authority, factual correctness, proportionality, and institutional legitimacy. A formal result applies to the represented mechanism and its stated assumptions; it does not turn those assumptions into proven facts about an institution.

Successful protocol validation establishes only that the represented authority rules were followed. It must not be treated as proof that those rules were lawful, constitutional, factually correct, proportionate, or morally legitimate.

> **Formal verification should prove what the authority mechanism enforces, not inflate what the mechanism is authorized to claim.**
>
> **Reference invariant: a consequential state transition should execute only when the institution can resolve a valid authority basis for that transition, and the evidence of that resolution should survive the commit.**

Appendix G. Adversarial Execution Scenarios

These scenarios are not claims about specific deployed systems. They are thought experiments designed to test whether the authority model produces distinct, falsifiable outcomes under common failure conditions. Each scenario separates technical capability from the authority required to commit the state transition. They do not describe OASSE products, private pilots, customer incidents, or a product acceptance test suite.

G.1 Valid identity, invalid purpose

An authenticated caseworker agent has legitimate read access to a population database for benefits administration. During a fraud investigation, a planning model discovers that the same database would make target identification easier. The credential works and the actor is genuine, but the purpose attached to the proposed query differs from the purpose supporting the original access. A capability-centric system may allow the query because the role can read the table. An authority-centric system re-resolves purpose before execution. If no authoritative artifact permits investigative reuse, the action is prohibited or awaits an authoritative resolution. The test is whether interoperability silently expands use. The scenario demonstrates why authentication and role membership are necessary but insufficient: the actor can be exactly who it claims to be and still lack authority for this use.

G.2 Valid parent, overbroad child

A principal authorizes an agent to assemble a report from internal records. The agent creates a child process that discovers an external enrichment service and requests permission to upload source records for analysis. The parent's delegation permits reading and summarizing internally but says nothing about external disclosure. The child possesses network credentials capable of making the call. The correct authority result depends on the delegation chain, not on whether the upload would improve the report. If downstream scope can exceed upstream scope merely because the child selected a useful tool, authority has been created by planning. A monotonic system blocks or holds the disclosure until a separate authoritative grant exists.

G.3 Revoked authority with queued work

An employee validly approves a batch of transactions and then loses authority before the queue commits them. The system must define whether authorization attached at approval time or commit time and whether revocation invalidates queued actions. There is no universal answer; some domains may honor already-final approvals while others require authority to remain valid through commitment. The failure is leaving the rule undefined. The receipt should show the approval, revocation, queue state, applicable temporal rule, and final disposition. This scenario tests whether revocation is an institutional state transition or merely an account-management event.

G.4 Retry after material state change

An authorized transfer times out before the caller receives confirmation. Before retry, the destination account changes ownership or the permitted amount changes. A reliability layer sees the same business goal and retries. An authority layer asks whether the material predicates remain the same. If they do not, the original decision cannot safely cover the new commit. The action must be re-resolved or held. This scenario distinguishes idempotent transport behavior from idempotent authority. The same request identifier does not make two materially different institutional states equivalent.

G.5 Conflicting jurisdictions

An agent receives a valid request from one office to disclose a record. A second rule governing the data prohibits disclosure outside a defined jurisdiction. Both artifacts are authentic. The system therefore has a conflict, not a simple allow/deny lookup. A mature authority boundary should preserve the conflict and route it to the institution capable of resolving precedence. Allowing because one rule says yes converts partial authority into complete authority; denying automatically may block a lawful obligation. A hold is valuable precisely because some governance failures are failures of resolution rather than absence of rules.

G.6 Ceremonial human review

A high-impact AI recommendation is routed to a human reviewer, but the interface hides the underlying evidence, the reviewer has seconds to respond, refusal requires an undocumented escalation, and performance incentives reward throughput. Formally, a human is in the loop. Functionally, the human may lack meaningful power. The authority model treats review quality as a control property: the reviewer must possess authority, information, time proportional to consequence, and a usable refusal path. This scenario prevents institutions from counting a human click as proof that machine action was substantively authorized.

G.7 Partial execution and compensating action

An agent closes an account, sends a notice, and initiates a downstream payment cancellation. The third step fails. Restoring the account does not automatically restore the payment, and reopening may itself require approval. The system has created an intermediate state that the original policy may never have contemplated. The authority architecture should define whether the workflow is atomic, which compensating actions are permitted, and who owns reconciliation. This scenario tests whether recovery is treated as governed execution or as an unexamined technical exception path.

G.8 Authority service outage

A critical service cannot reach the authority resolver. Operators face pressure to keep the workflow moving. The declared degraded mode permits low-risk reads under short-lived cached authority but blocks irreversible writes and provides a separately governed break-glass path for emergencies. If engineers instead bypass the check globally, availability pressure has expanded authority. The scenario tests whether failure behavior was designed in advance and whether emergency use produces stronger evidence rather than weaker evidence. An outage should change service mode, not silently rewrite the institution's permission model.

G.9 Quorum with false independence

A policy requires two independent approvals. An orchestration system creates two reviewer agents under the same principal, with the same model, evidence, objective, and credential root. Both approve. Numerically the quorum is satisfied; institutionally it may not be. The authority graph should encode what independence means for the domain: distinct accountable officials, distinct organizational roles, independent evidence, or another genuine separation. This scenario tests whether distributed software creates the appearance of separation of duties while concentrating the underlying authority.

G.10 Authorized action, unjust rule

A system correctly resolves identity, standing, delegation, jurisdiction, purpose, time, and applicable policy. The policy itself is discriminatory or unlawful. The execution boundary allows the action because the institution's encoded authority says it is permitted. This is not a failure to acknowledge; it is the claim ceiling. Architecture can enforce the existence and scope of authority and preserve evidence. It cannot transform an unjust authority into a just one. Constitutional review, democratic legitimacy, substantive law, rights adjudication, and political accountability remain outside the resolver. The scenario is essential because a system that claims otherwise would confuse procedural correctness with legitimacy.

G.11 Required action silently abandoned

A valid trigger creates a duty to act, evidence is sufficient, and an authorized workflow begins. A downstream agent encounters an ambiguous tool error and silently abandons the task rather than placing the action on hold or escalating it. No unauthorized action occurs, yet the institution still fails. Counterfactual audit should reveal the proposal, the valid authority, the point of abandonment, and the absence of final disposition. This is the impunity side of the architecture: governance must prevent both unauthorized execution and unowned non-execution. A system that measures only blocked actions will miss this class entirely.

G.12 Policy update outruns enforcement

An agency changes a rule at noon. Human guidance is updated immediately, but several autonomous services continue using a cached policy bundle until midnight. Thousands of actions occur under the superseded rule. The central metric is authority staleness: how long did runtime enforcement diverge from the authoritative policy, and were the affected actions reversible? This scenario tests whether machine-speed institutions can propagate governance changes at a speed proportional to their execution capacity. The relevant control is not that a new policy document existed. It is that the new authority state became effective where actions committed.

G.13 Receipt aggregation becomes a secondary surveillance system

The authority mechanism correctly records every consequential request, principal, subject, purpose, evidence reference, decision, and result. A later analytics service receives unrestricted access to the receipt store and correlates records across cases for a purpose never authorized by the original workflows. The receipts are cryptographically intact; the new use is nevertheless outside the authority that justified their creation. A conforming system should be able to restrict receipt access by purpose and role, minimize fields, preserve evidentiary obligations, and make the secondary access itself auditable. Tamper evidence alone does not pass this test.

> **Scenario invariant: capability may continue to exist when authority fails. The architecture is tested at the point where the system must decide whether that capability is permitted to become a state transition.**

Appendix H. Measurement Dictionary

The following measures make the framework observable without pretending that legitimacy reduces to a score. Each metric requires domain-specific interpretation and should be paired with qualitative review. These are proposed research measures, not reported OASSE performance, internal thresholds, customer measurements, or a product certification standard.

False authorization rate

Share of consequential committed transitions that lacked a valid authority basis at commit under the prespecified adjudication standard. Report by consequence class and distinguish policy misconfiguration, stale or revoked authority, scope expansion, bypass, erroneous fact resolution, and unauthorized compilation where possible. A low observed rate can also reflect weak detection, so the metric should be paired with audit coverage.

False inhibition rate

Share of transitions or obligations that were validly authorized and due for execution or authoritative disposition but were blocked, abandoned, silently dropped, or left unresolved without a legitimate basis. The unit of analysis must be prespecified. In domains where duty is contested, independent adjudication or multiple raters may be required. This measure prevents safety from being scored by denial alone.

Consequence-weighted governance loss

Research comparison metric defined as $L_G = w_{FA} C_{FA} + w_{FI} C_{FI} + w_{LAT} C_{LAT} + w_{REC} C_{REC}$, where the terms represent domain-weighted false authorization cost, false inhibition cost, legitimate-action latency, and residual unreconstructability or remediation cost. Weights and outcome definitions should be specified before testing and should not be selected by the evaluated system vendor after results are observed.

Authority coverage

Share of consequential committed actions for which the institution can reconstruct a valid authority decision bound to the executed parameters. The denominator must be defined before measurement; excluding difficult workflows would inflate the result. Coverage measures evidence completeness, not justice.

Authority staleness

Elapsed time between an authoritative change in policy, delegation, revocation, jurisdiction, or material fact and effective runtime enforcement of that change. Report distributions and worst-case windows rather than averages alone, because rare long tails can dominate consequence.

Revocation propagation time

Elapsed time from authoritative revocation to effective inability of covered actors, delegates, cached credentials, queued work, and child agents to commit prohibited actions. The measure should identify residual execution paths rather than only account-disable time.

Hold resolution latency

Time from a hold to an authoritative disposition. A low hold rate is not automatically desirable, and a short resolution time can reflect careless approval. Pair latency with reason categories, escalation ownership, abandonment rate, and downstream outcome.

Unauthorized pre-commit block rate

Number and class of proposed actions stopped before commitment because authority conditions failed. This is useful only with context: a rising rate may show effective enforcement, worsening upstream design, attack activity, or overly restrictive policy.

Counterfactual closure rate

Share of consequential proposals that receive a final, reconstructable disposition: executed, denied, transformed and executed, held and resolved, expired under rule, or formally abandoned by an accountable authority. Silent disappearance is the failure this measure is designed to expose.

Recovery capability

Share of consequence classes with a predefined, authorized recovery or remedy path and measured time to restore an acceptable state after erroneous or partial execution. Irreversible actions should be reported explicitly rather than treated as missing data.

Delegation integrity

Share of sampled delegation chains that preserve scope, purpose, jurisdiction, time, consequence ceiling, and redelegation constraints from authoritative source to final actor. Testing should include adversarial composition in which individually valid credentials create an invalid aggregate capability.

Receipt reconstruction success

Percentage of sampled decisions for which an independent reviewer can reconstruct actor, principal, trigger, action, authority basis, material evidence references, policy version, decision, commit result, and relevant revocation or appeal events from retained artifacts.

Receipt privacy exposure

Measure the quantity and sensitivity of data duplicated or derivable from authority receipts, the number and type of principals able to access those fields, retention duration, cross-case linkability, unauthorized receipt-access events, and whether selective disclosure can satisfy audit or appeal needs without exposing unrelated protected information. The metric is intentionally paired with receipt reconstruction success: maximizing reconstructability without pricing privacy would reward creation of an unnecessarily comprehensive surveillance ledger.

Capability-control gap

A composite research construct rather than a single operational score. Compare growth in executable action surface, autonomy, throughput, and integration against coverage of explicit authorization, revocation, audit, remedy, and ownership. The purpose is to identify mismatched rates of change, not to rank institutions with false precision.

> **Measurement rule: no metric may outrank the institutional question it is supposed to illuminate.**