




Institution-to-Execution Architecture (I2EA)

From Human Governance to Governed AI Execution

Mana Onishi

2026

Why Institution-to-Execution Architecture (I2EA) Is Needed

Current discussions of AI Governance have advanced numerous principles and requirements, including safety, fairness, transparency, privacy, accountability, and human oversight. Institutional frameworks for Governing AI—such as laws, regulations, international principles, technical standards, industry guidelines, and internal organizational Policy—are also developing rapidly.

However, a critical issue is that the existence of Governance is not the same as the ability of Governance to control actual AI Decision and Execution.

Governance in human society is primarily expressed in natural language and presupposes interpretation by humans and operation through organizations and institutions. AI Systems, by contrast, form Decisions and Actions at high speed and with considerable complexity through Data, Models, Applications, Agents, Tools, APIs, Devices, and other components. Even where institutional Requirements exist, Governance remains disconnected from actual Execution unless those Requirements are connected to AI Systems in Machine-readable form and clearly specify what applies in a particular Context, who holds Authority, and under what conditions Execution is permitted.

Furthermore, the institutions and organizations of human society have never operated solely through formally articulated Rules.

In actual organizational operations, it is difficult to specify in advance every condition for judgment and every exception-handling procedure as a rule. When institutions and procedures alone could not address a situation, people in operational settings have exercised judgment based on experience, Context, established practice, prior cases, and relationships with others, thereby bridging the space between institutions and actual Operation.

Much of this work has been treated not as an independent function but as an individual capability described as being attentive, experienced, or well acquainted with the circumstances. Consequently, although such judgments may be essential to the actual functioning of an organization, their conditions and exception-handling processes are not always explicitly recorded or evaluated as a Governance Structure.

In organizations composed of humans, such supplementation through tacit knowledge often does not create a major problem. Humans can combine rules with social common sense, organizational culture, relationships with others, past experience, and subtle differences in circumstances to determine that a particular case differs from the ordinary situation.

When the same structure is transferred to AI, however, the nature of the problem changes.

If AI itself is allowed to supplement unstated conditions and exceptions through Inference, there is no guarantee that the exceptions, Authority, Intent, practices, or permissible scope inferred by the AI will correspond to what human society actually recognizes. The more advanced the AI's Reasoning capability becomes, the greater its capacity may be to generate plausible explanations and judgments in response to institutional gaps.

Tacit knowledge that was effective for humans may therefore become a source of erroneous judgment if transferred directly to AI.

What is required here is not the mechanical conversion of all tacit knowledge into Rules. Nor can ambiguity and exceptions be completely eliminated from the real world.

What matters is to externalize necessary conditions, exceptions, Authority, and grounds for judgment as an explicit Governance Structure to the extent possible, while also establishing a structure that prevents AI from using Inference to create nonexistent Authority, Policy, Exceptions, or Permissions and proceeding to Execution when sufficient Governance grounds do not exist.

This is where the structural Gap between Institution and Execution arises.

This Gap cannot be resolved merely by making AI more intelligent. Even if advanced AI can understand laws and Policy, understanding them is different from applying them under legitimate Authority. Likewise, even if AI forms an appropriate Decision, that Decision is not necessarily Enforced reliably in actual Execution.

Monitoring AI Output alone is equally insufficient. Governance requires Evidence not only of the Outcome—what occurred—but also of which Information and Policy were applied, under which Authority and Context the Decision was formed, how that Decision was connected to Execution, and what actually occurred.

Moreover, AI is not a static System completed at Deployment. Models, Data, User Behavior, Operational Context, connected Systems, laws and Policy, and the social environment change continuously. Governance therefore requires more than pre-Deployment Evaluation or a one-time Approval. It requires a structure that continuously observes Evidence from Real-world Operation, detects new Risk, Deviation, and Patterns, and, when necessary, provides Feedback to the legitimate Governance Authority.

What is required, therefore, is not the addition of one more AI principle.

What is required is a coherent Architecture that converts Institutional Governance already established in human society into Machine-readable Governance, connects Information, Policy, Authority, and Context to Decision, Enforces that Decision through actual Execution, and returns Evidence obtained from Execution to Governance.

Institution-to-Execution Architecture (I2EA) was designed to connect this structural Gap.

The purpose of I2EA is not to make AI itself a Governance Authority. It is to preserve the Governance Authority of human society while connecting it to AI System Decision and Execution and making that process verifiable.

The central question addressed by I2EA is:

How can human governance remain effective when decisions and actions are increasingly performed by AI systems?

This is the question at issue.

In response, I2EA presents as a single Architecture the structure extending from Institution to Decision and Execution, and then from Evidence obtained through Execution back to Governance.

What Is I2EA

Governance that has been performed implicitly in human society must become an explicit structure that AI can process.

'Treat this as a fact'	Fact
'Evaluate this situation in this way'	Evaluation
'This rule applies'	Policy
Understand 'who may decide and change what'	Authority
Recognize 'the procedure actually used for processing'	Process
'Translate human Governance into a structure that AI can process'	Translation Hub
'Connect structured Governance to the AI System'	Hamecohmig
Decide, 'These conditions require this response'	Decision
Ensure that what was decided is followed through actual action	Execution Boundary
Preserve 'why that Decision was made at that time'	Decision-State Evidence (DSE)
Detect that 'something has recently changed'	AI Vigilance

I2EA does not require AI to imitate the implicit Governance of humans.

**Structure what can be made explicit.
Do not allow AI to fill in what cannot be made explicit at its own discretion.**

While preserving the Governance Authority of human society, I2EA connects Institution to Decision, Execution, Evidence, and back to Governance.

AI as Part of Society

AI does not function as an entity independent of and external to human society. It connects with people, companies, governments, professionals, institutions, and services; handles information within society; supports judgment; and at times undertakes actual actions.

I2EA is designed on the premise that AI extends human capabilities rather than simply replacing them. Even when the same AI is used, the results differ according to the user's knowledge, experience, expertise, judgment, and creativity. As AI capability increases, the capability of the human using it becomes increasingly important as well.

The development of AI therefore does not diminish the importance of education. As humans themselves learn and develop, the range of capabilities that AI can extend also expands.

However, human knowledge, judgment, experience, social rules, and authority cannot be properly connected to AI if they exist only in the human mind.

I2EA is not an Architecture for replacing humans with AI. It is an Architecture for connecting the Knowledge, Expertise, Judgment, and Authority of human society with AI Capability and linking them to Governed Execution.

Governance Records

I2EA separates the information required for Governance into distinct Records according to their respective roles.

Fact

What existed or what occurred.

Policy

What Governed the situation.

Evaluation

How the Fact was evaluated.

Process

What processing the System actually performed.

Fact ≠ Policy ≠ Evaluation ≠ Process

Separating these Records prevents an AI-generated Evaluation from being treated as Fact, or an actually performed Process from becoming the basis of Policy, and makes each independently verifiable.

I2EA explicitly separates not only types of information but also Authority.

Authority is not a type of Record. It is a separate axis of Governance indicating who possesses the legitimate power to decide, approve, and change what.

Accordingly, Fact cannot be rewritten to conform to Evaluation, and Evaluation cannot change Policy. A change to Policy requires legitimate Governance Authority and Process.

Record Type ≠ Authority to Change

I2EA structures Governance information while also structuring who is permitted to act upon that Governance.

Translation Hub

Governance in human society—including laws, regulations, industry standards, professional guidelines, and internal organizational Policy—is generally written in natural language for human readers.

Even identical text, however, is not necessarily interpreted consistently by AI. If interpretation varies with Context, Model, Prompt, or other factors, different conditions for judgment may be formed from the same Governance.

The I2EA Translation Hub is a mechanism for translating the Governance of human society into a Machine-Readable Governance Structure that AI can reference and apply consistently.

Machine-Readable here does not mean conversion into a 'machine language' incomprehensible to humans. It means converting Policy content, conditions of application, Authority, exceptions, subjects, Version, and other elements into a form that AI can process structurally while remaining verifiable by humans.

Human-Readable Governance

→ **Structured Translation**

→ **Human-Verifiable & Machine-Readable Governance**

Structuring Governance alone, however, does not guarantee that the translation is correct.

Every institutional and professional domain—law, medicine, finance, and others—has its own concepts, terminology, exceptions, and interpretations. If Governance is made machine-readable without understanding that meaning, the resulting data may be formally well structured while carrying a meaning different from the original Governance.

The Translation Hub therefore requires Review & Validation by human experts who understand both the original Governance and its specialist terminology.

Experts are not an Authority that creates new Governance. Their role is to verify that the meaning of Governance established by legitimate Authority has not been lost or altered during the process of making it machine-readable.

Expertise ≠ Authority

The Translation Hub is not a place where Governance is decided.

It is a place that bridges Governance already established in human society into a structure that AI can process while preserving its meaning and Authority.

Umecohming / Hamecohming

I2EA uses two Governance Functions—Umecohming and Hamecohming—to connect Information obtained from Reality and the Governance of human society to AI Decision.

Umecohming

Umecohming is a mechanism that structures information obtained from Reality as a machine-readable Governance Context that AI can process under Governance.

Information processed by AI has not only content but also Governance attributes—such as Source, Identity, Provenance, and Confidentiality—that affect subsequent Evaluation and Decision.

Umecohming associates these attributes with the information as Governance Context, preventing AI from processing informational content in isolation from that Context.

Reality → Information → Governance Context

The name derives from the Japanese word 'umekomu' and the English 'to embed.' Because it denotes a continuing Process rather than a static Object, '-ing' was combined with the term to form Umecohming.

Hamecohming

Hamecohming is a mechanism that converts laws, regulations, Policy, institutional requirements, and other forms of Governance existing in human society into a machine-readable Governance Structure and connects that Structure to an AI System.

It connects Governance structured and validated through the Translation Hub described in the preceding section in a form that AI Evaluation and Decision can reference. Hamecohming itself, however, neither determines which Policy applies nor creates new Governance.

Human Governance → Machine-Readable Governance Structure

The name derives from the Japanese word 'hamekomu' and the English 'to fit into.' As with Umecohming, '-ing' was combined with the term to indicate a continuing Process, forming Hamecohming.

Two Different Connections

Umecohming and Hamecohming support AI Decision from different directions.

Umecohming = Information → Governance Context
Hamecohming = Human Governance → Governance Structure

Umecohming structures *****What Governance Context does the information being processed by AI carry?*****

Hamecohming addresses *****How is the Governance of human society connected to a structure that AI can reference?*****

Neither performs the Decision itself. They are Functions for structuring and connecting the information and Governance on which a Decision may legitimately rely.

Why New Terms?

Umecohming and Hamecohming are terms coined for I2EA.

These Functions are not identical to any one existing technical concept, such as Tagging, Metadata, Translation, or Policy Management. I2EA positions them as independent Governance Functions that connect the information and Governance of human society to AI Decision and Execution.

If expressed only through existing technical terminology, they may be understood merely as parts of those existing functions, and the roles and boundaries intended by I2EA may be lost. New names were therefore assigned to clarify their boundaries from existing concepts and to treat them as independent Functions.

**Structure the Governance that can be made explicit and
connect it to Decision.**

**Do not allow AI to invent Governance that has not been
structured.**

Decision & Execution Boundary

I2EA explicitly records a Decision formed on the basis of Governance (Fact, Evaluation, Policy, Process + Authority) and uses that Decision as the criterion for subsequent Execution.

The mere formation and recording of a Decision, however, does not guarantee that Execution will conform to that Decision.

If compliance with the Decision is also entrusted to the AI itself, Decision and Execution are completed within the same System, leaving no independent Enforcement Boundary.

I2EA therefore establishes an Execution Boundary outside the AI.

Governance → Decision → Execution Boundary → Execution

The Execution Boundary is not a place for reinterpreting Policy or forming a new Decision. Immediately before Execution, it verifies whether a valid Decision exists and whether the Execution conditions associated with that Decision are satisfied, permitting Execution to pass only when those conditions are met.

In other words:

Decision ≠ Execution

Decision Authority ≠ Enforcement Responsibility

The role of the Execution Boundary is not to decide again, but to verify a Decision already formed and enforce it upon Execution.

Why OS-Level Enforcement?

The Execution Boundary does not require a particular technology or implementation method. What the Architecture requires is an independent Enforcement Mechanism outside the AI rather than reliance solely on the AI's self-control.

I2EA presents OS-Level Enforcement as a leading implementation approach.

The operating systems in widespread use today are concentrated among a small number of major Platforms. Establishing a common Boundary at a layer close to the OS may therefore enable broader System Actions to be controlled across Applications, rather than implementing Execution Control separately in each AI Application.

Furthermore, given the current market structure in which providers of major general-purpose operating systems are concentrated among U.S. companies, this is an implementation candidate that may combine broad coverage with practical deployability in terms of standardization and implementation cooperation.

eBPF (extended Berkeley Packet Filter) is one specific technological example. Because it can observe and control System Actions at a layer close to the OS, it may be used as a technology for implementing the Execution Boundary.

I2EA, however, requires neither eBPF nor any particular OS. Depending on the System architecture, different implementation methods may be adopted, including OS controls, API gateways, sandboxes, and device controllers.

Execution Boundary = Architectural Requirement
OS-Level Enforcement = Recommended Implementation Approach
eBPF = Example Technology

The architecture requires the boundary, not a specific technology.

Decision-State Evidence

Even when AI safety and Governance are assessed by a third party, the resulting assessment may fail to command sufficient confidence. One reason is that no adequate common structure has yet been established for specifying what was assessed, on the basis of which Evidence, and according to which criteria.

This is easier to understand through an analogy with accounting.

For a company to be assessed by a third party, it is not sufficient merely to possess a large volume of transaction records. Transactions are recorded consistently under bookkeeping rules, accounting standards organize principles of recognition and evaluation, and audit has criteria for verification.

Because such a common structure exists, third parties can refer to the same Evidence while conducting qualitative and quantitative assessments of the company's condition.

AI Governance does not yet have a sufficiently developed Evidence foundation to support this type of assessment and verification.

I2EA likewise does not itself complete evaluation criteria or audit standards common to all AI.

As a preceding step, it provides a mechanism for preserving, in a consistent structure, the state that existed when a Decision was formed and for continuously accumulating Evidence that can later be compared and verified.

This is Decision-State Evidence (DSE).

One Decision = One DSE

DSE is not merely a Decision Log.

Its components are grounded in the I2EA Governance Structure, including Fact, Evaluation, Policy, Process, and Authority. It also associates and preserves the Model, Context, applied Governance, Decision, and other state necessary to reconstruct later why that Decision was formed at that time.

DSE therefore records more than simply 'what the AI answered' or 'which Action it selected.'

It preserves the information, Governance, Evaluation, Authority, Model, and Context under which the Decision was formed.

What was actually executed after the Decision is treated not as part of DSE itself but as Execution Evidence.

DSE = Why the Decision Was Formed

Execution Evidence = How It Was Enforced and What Happened

This separation makes it possible to verify later both the Decision itself and how that Decision was actually Executed.

If DSE is continuously accumulated in a consistent structure, it can form an Evidence foundation for comparing different Decisions, different points in time, and different Systems.

I2EA does not standardize assessment results themselves. It first enables the regular and verifiable accumulation of foundational data needed to establish future evaluation criteria, audit, Assurance, and qualitative and quantitative analysis.

**Before standardizing evaluation, create Evidence that
can be evaluated.**

Vigilance & Feedback


A single Decision-State Evidence (DSE) record provides Evidence for reconstructing why a particular Decision was formed at a particular point in time.

Governance of an AI System, however, cannot be achieved by verifying a single Decision alone.

As DSE and Execution Evidence accumulate continuously, Deviation, Anomaly, Repetition, Shift, and Emerging Pattern that could not be identified by examining individual Decisions alone can be detected over time.

This resembles accounting: a single transaction record cannot reveal changes across an entire company, but continuous analysis of systematically recorded data can reveal trends, anomalies, and structural changes.

I2EA positions this mechanism of continuous observation, detection, analysis, and Feedback as AI Vigilance.



Evidence and Findings obtained through Vigilance provide a basis for reviewing whether existing Evaluation Criteria, Policy, operational conditions, and related elements are functioning appropriately in relation to the actual state of an AI System.

Evaluation Criteria and Governance can then be improved as necessary, and the results reflected in future Decision.

AI Vigilance itself, however, does not automatically rewrite Policy or Evaluation Criteria.

The role of Vigilance is to identify change and problems and provide the resulting Evidence as Feedback to Governance. Changes to Policy or Evaluation Criteria are made through the Governance Process by actors holding the legitimate Authority for each.

Evidence → Learning ≠ Authority → Authorized Update

Updated Governance is subsequently reflected in Decision and connected through the Execution Boundary to future Execution.

**Observe → Detect → Evaluate → Feedback → Authorized Update → Future Decision
→ Execution**

I2EA thereby forms a Feedback Loop that returns Evidence obtained from actual Execution to Governance and enables continuous improvement, rather than establishing Governance once at deployment and leaving it as a static mechanism.

Governance → Decision → Execution → Evidence → Vigilance → Governance

Institution-to-Execution Ecosystem

Many initiatives supporting AI Governance are already under way around the world.

They include governments and international organizations that establish laws and regulations; companies and Engineers that develop AI; Researchers who study safety; Organizations and Individuals that use Systems; professionals responsible for independent Audit and Assurance; Insurance that evaluates and transfers Risk; AI Vigilance Services responsible for continuous Monitoring; and Creators who produce the information and works used by AI.

Each has a distinct role.

I2EA does not seek to consolidate them under one vast organization or a single Governance Authority.

It seeks to connect them through a common Architecture while preserving their respective roles, expertise, and Authority.

I2EA designs Human Governance, Machine-Readable Governance, Decision, Execution, Evidence, Vigilance, and Feedback to Governance as one continuous structure.

Architecture alone, however, is insufficient to make this structure function in society.

Government forms laws and regulations as a legitimate Governance Authority.

Developers and Technology Providers connect them to actual AI Systems and Execution Infrastructure.

Organizations and Individual Users use AI in different Contexts and connect Governance to actual social activity.

Auditors and Assurance Providers independently verify, on the basis of Evidence, whether Governance is functioning appropriately.

Insurers connect accumulated Evidence to Risk Evaluation and Risk Transfer.

Creators and Information Providers connect the Provenance, Rights, and Value of Knowledge and Creative Works used by AI to Governance.

Researchers, Standards Bodies, and Professional Communities develop Evaluation Criteria, Interoperability, Technical Standards, and Governance Practices.

Evidence and Learning obtained through AI Vigilance are then provided as Feedback to the respective legitimate Authorities and connected to subsequent Governance.

The Governance Tapestry

When I began developing I2EA, many excellent 'threads' already existed in AI Governance.

Safety, Law, Authority, Cybersecurity, Audit, Evidence, Standards, Certification, Insurance, Monitoring.

The problem is not the absence of these threads.

The problem is how to connect them from Institution to Execution.

I2EA is an Architecture that seeks to design the Tapestry for making those connections.

A Tapestry, however, cannot be woven by one person alone.

I2EA itself cannot be realized by one Researcher, one Company, one Government, or one Technology. Connecting it to actual society requires the participation of Governments, Companies, Researchers, Engineers, Auditors, Insurers, Creators, Organizations, Individual Users, and diverse specialists.

**Different threads. Different roles. Different authorities.
One connected Governance Tapestry.**

I2EA is not a completed social system.

This is an Architecture to be woven together.

Acknowledgements 1

The development of DSE and its surrounding concepts has been substantially influenced by ongoing technical discussions, implementation work, and validation with Robert Schneider (FlexVertex). The author is responsible for the overall I2EA Architecture and its final content; however, because the concepts developed reciprocally through repeated discussion, it is difficult to draw precise boundaries around contributions to individual ideas.

Acknowledgements 2

I2EA was developed together with ChatGPT—from informal conversations during the earliest stage of the concept (ChatGPT 4o) to the articulation of the present Architecture (GPT-5.5 Sol)—using only a Plus subscription.

It engaged patiently and carefully with my incomplete questions and, at times, even my misdirected frustration. The English translation is also almost entirely attributable to it.

Without you, it would have been difficult to develop I2EA to this point.