DIG®

Digital Information Governance: The Standard for Defensible AI-Influenced Decisions in Energy

The problem, the vocabulary, the framework, the maturity model, the principles, the risks, and the operating model for governing decisions in the age of machine-shaped judgment.

AUTHOR MATTHEW BERTRAM
PUBLISHED BY MODALPOINT
EDITION V1.0 · Q3 2026 · HOUSTON
MARK DIG® / DIGITAL INFORMATION GOVERNANCE® · USPTO REG. NO. 8147558

PUBLICATION

Digital Information Governance (DIG®): The Standard for Defensible AI-Influenced Decisions in Energy. Version 1.0, published Q3 2026 by ModalPoint, Houston, Texas. This document is the source of record for the DIG framework. The reference site at digitalinformationgovernance.com is maintained in alignment with this paper.

PATENT NOTICE

Certain methods described in Section 14, including authority-hierarchical validation of machine-initiated actions, non-probabilistic governance override, and the Agentic Governance Control Plane architecture, are the subject of pending United States patent applications.

DISCLAIMER

This paper describes a governance discipline. It is not legal advice, and it does not substitute for counsel on the EU AI Act, the Texas Responsible Artificial Intelligence Governance Act, or any other statute or regulation cited herein. Statistics are

drawn from primary sources listed in Section 17. Field vignettes are illustrative composites and do not describe any specific company.

CITATION

Bertram, M. (2026). Digital Information Governance (DIG®): The Standard for Defensible AI-Influenced Decisions in Energy, v1.0. ModalPoint, Houston.

CONTENTS

In this paper

SECTION 01

Executive summary

Artificial intelligence has moved into the decisions energy companies are held accountable for. Governance has not followed it there. This paper defines the discipline that closes that gap.

In 2024, 78 percent of organizations reported using AI in at least one business function. In 2026, only 21 percent report a mature governance model for the agentic AI now acting inside their operations.[1, 2] Between those two numbers sits the defining operational risk of this decade: decisions that are increasingly shaped by machines, made at machine speed, and recorded nowhere.

Energy feels this earlier and harder than most industries. An AI-influenced decision in a bank that goes wrong is a financial loss. An AI-influenced decision on a drilling program, a maintenance deferral, a trading limit, or a capital allocation can be a well integrity event, a fatality investigation, or a nine-figure write-down. And when that decision is challenged, by a regulator, a partner, an insurer, or a court, the question will not be whether the model was accurate. The question will be simpler and older: **who decided, on what information, under what authority, and can you show me?**

Most organizations cannot answer it. Their governance programs were built for other layers of the problem. Data governance governs tables. Information governance governs records. AI governance, as currently practiced, governs models: bias, drift, explainability. None of them governs the decision itself, the moment where machine output and human authority meet and something irreversible happens in the physical or financial world. That moment is currently the least instrumented point in the modern enterprise.

Digital Information Governance (DIG®) is the discipline for that layer. It keeps AI-influenced decisions defensible and auditable: the information feeding them is sourced and trusted, the decisions themselves are traceable to named human authority, the organization is accurately represented across the AI systems that now describe it, and all of it can be proven on demand. This paper defines the discipline in full. It establishes the vocabulary, sets out the four-pillar framework (Information Provenance, Decision Traceability, Representation Integrity, Audit Readiness), introduces the five-level DIG Maturity Model, states seven governance principles, maps the risk taxonomy and the regulatory environment, and specifies an operating model that a leadership team can stand up with named roles, defined artifacts, and a fixed cadence.

What the paper argues

First, the unit of governance is the decision, not the model. A perfectly validated model feeding an untraceable decision is ungoverned. Regulators are converging on the same view: the EU AI Act requires event logging that supports traceability of high-risk systems, and the Texas Responsible Artificial Intelligence Governance Act, effective January 1, 2026, attaches Attorney General enforced penalties of up to $200,000 per uncurable violation to obligations that are ultimately about conduct and accountability, not model mathematics.[5, 8]

Second, defensibility must be captured at decision time, not reconstructed after a challenge. A record assembled three weeks after an incident, from fragments across five systems, is an argument. A record captured contemporaneously is evidence. The DIG Maturity Model measures exactly this shift: from Level 1, where decisions leave no usable trail, to Level 5, where defensibility is the resting state of the organization.

Third, governance limits on machine-initiated action must be deterministic. As agentic systems begin to execute rather than merely recommend, probabilistic controls (a model judging another model) are not a control environment. Authority must be explicit, hierarchical, and enforced by mechanisms a model cannot talk its way past. Section 14 specifies this architecture.

Fourth, governance pays. Organizations that regularly audit and assess their AI systems are more than three times as likely to report high value from generative AI.[3] Auditability is not a tax on velocity. It is the precondition for delegating more consequential work to machines, which is where the value is.

Who this paper is for

It is written for the executives who already sense the problem: the COO who has watched AI recommendations enter operational decisions with no record, the general counsel who knows the next incident investigation will ask questions the company cannot answer, the CFO signing off on capital decisions shaped by models nobody can reconstruct, and the board member who understands that "the AI did it" has never once worked as a defense. It assumes no machine learning background. It assumes operating responsibility.

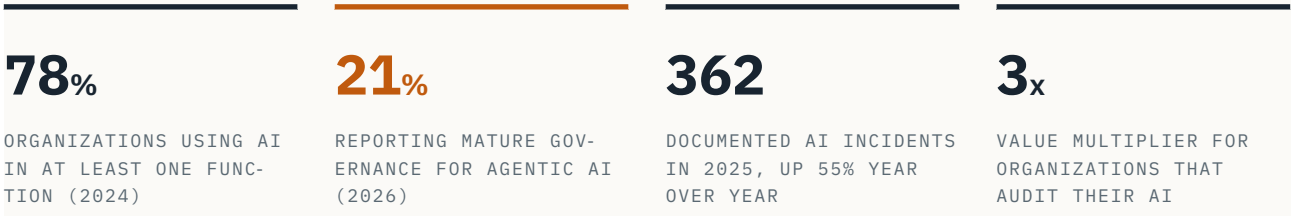

FIG. 01 · THE ACCOUNTABILITY GAP, BY THE NUMBERS · SOURCES: SEC. 17

SECTION 02

Houston, 1970: the archetype of decision integrity

Fifty-six years ago, a room in Houston made a string of irreversible decisions under extreme uncertainty, with three lives on the line and the whole world watching. Every one of those decisions could be defended then, and can still be reconstructed today. It is worth asking why.

On the evening of April 13, 1970, roughly 200,000 miles from Earth, an oxygen tank in the service module of Apollo 13 exploded. Within minutes, Mission Control in Houston was looking at telemetry that made no sense: pressure readings collapsing, voltages sagging, a spacecraft venting something into space. The first instinct of several controllers was rational and revealing. They suspected the instrumentation, not the spacecraft. Data you cannot trust is worse than no data at all, and the controllers knew it. Before anyone acted on the readings, the room worked to establish whether the readings themselves could be believed.

What followed over the next four days is usually told as a story of improvisation: the lunar module pressed into service as a lifeboat, the carbon dioxide scrubber adapted with plastic bags and tape, a power-up sequence invented from nothing and rehearsed in the simulator before it was ever transmitted to the crew. It was all of that. But read as a governance case, Apollo 13 is something else. It is the cleanest example on record of an organization making consequential decisions, fast, under uncertainty, without ever losing the thread of who decided, on what information, and under what authority.

The mechanics of a defensible decision

Consider how that room actually worked. Every controller owned a system and spoke for it by name: EECOM for electrical and environmental, GNC for guidance, FIDO for flight dynamics. When the flight director needed a decision, he polled the room. Each controller answered GO or NO GO, out loud, on a recorded loop, and each answer carried personal accountability for the state of one slice of the spacecraft. The flight director held final decision authority, explicitly delegated and universally understood. Nobody in that building was confused about whether a recommendation was a decision.

The information feeding those decisions had known provenance. Telemetry came from identified sensors with understood failure modes. When the data was suspect, its suspect status was itself surfaced and weighed. The trajectory decision that shaped the whole rescue, whether to attempt a fast direct abort or loop around the Moon on a free-return path, was made against explicit constraints (a possibly damaged main engine, consumables budgets, crew condition) that everyone could see and that survive in the record today.

And everything was recorded contemporaneously. The flight director loop, the backroom loops, the console logs. When the review board convened after splashdown, it did not conduct interviews to reconstruct what people remembered deciding. It replayed the decisions as they were made. The organization was, in the vocabulary this paper will formalize, defensible by default. Not because anyone expected litigation, but because the discipline of capturing decisions at decision time was the operating culture.

> **Mission Control did not have better information than a modern operator. It had better custody of its information, and better custody of its decisions.**

The inversion

Now consider a consequential decision made in an energy company this quarter. A production engineer asks an AI copilot to summarize eighteen months of well integrity data and recommend whether a workover can slip to next quarter. The model reads from a data lake of uneven freshness, drafts a confident recommendation, and the engineer, under schedule pressure, accepts it. The deferral goes into the plan. No record captures what the model read, what it recommended, what the engineer changed, or on whose authority the deferral was approved. The chat session ages out in thirty days.

That decision was made with a fraction of the discipline available to a room full of engineers in 1970, using slide rules and mainframes with less computing power than a key fob. The modern organization has incomparably more data and incomparably less custody. It makes more decisions, faster, with more machine influence, and can reconstruct fewer of them. This is the inversion at the heart of this paper: **decision speed has scaled with the technology; decision integrity has not.**

The rest of this document is, in one sense, an attempt to give the modern operator what that room in Houston had. Not the nostalgia and not the headcount, but the structure: named authority, sourced information, contemporaneous records, and the standing ability to prove all of it. The GO/NO GO poll was authority-hierarchical validation performed by humans. The flight director's rules were deterministic overrides. The loop tapes were decision records. DIG is those disciplines, rebuilt for a world where some of the voices on the loop are machines.

SECTION 03

The problem: the accountability gap

Adoption is near-universal. Mature governance is rare. Incidents are rising. Penalties are now written into statute. The gap between using AI and being able to defend its use has become a balance-sheet risk.

The research community gave this problem its name before industry felt its weight. In 2020, Raji, Smart and colleagues described the **accountability gap**: once AI systems are deployed, the harms they contribute to become difficult to trace back to their source, and therefore difficult to assign, remedy, or learn from.[9] What was then a concern about algorithmic auditing has since become the general condition of the enterprise. AI is no longer a tool at the edge of the business. It sits inside the decision flow: screening, summarizing, ranking, recommending, and increasingly executing.

The numbers describe a widening gap, not a closing one. AI adoption reached 78 percent of organizations in 2024, up from 55 percent a year earlier.[1] Yet in 2026, only 21 percent of organizations report a mature governance model for agentic AI, meaning roughly four in five lack basic capabilities such as audit trails and clear decision boundaries for systems that can act.[2] Recognition of risk consistently outruns mitigation of risk: explainability was rated a relevant risk by 40 percent of organizations but actively mitigated by only 31 percent; fairness by 34 percent against 26 percent. Mitigation lagged recognition in every measured category.[1]

Meanwhile the failure rate is climbing and the response capacity is falling. Documented AI-related incidents reached a record 362 in 2025, a 55 percent increase over the prior year, and that database is a reported floor, not a census.[11] Over the same period, the share of organizations rating their own AI incident response as excellent fell from 28 percent to 18 percent, even as the share experiencing three to five incidents climbed from 30 percent to 50 percent.[11] Organizations are having more AI failures and getting worse at handling them, simultaneously.

40% / 31%

EXPLAINABILITY: RATED A RISK VS. ACTIVELY MITIGATED

50%

ORGANIZATIONS HIT BY 3 TO 5 AI INCIDENTS IN 2025, UP FROM 30%

18%

RATE THEIR AI INCIDENT RESPONSE AS EXCELLENT, DOWN FROM 28%

FIG. 02 · RECOGNITION OUTRUNS MITIGATION; INCIDENTS OUTRUN RESPONSE · SOURCES: SEC. 17

Why energy is the leading edge

Every industry has this gap. Energy has it with consequences attached. Three structural features make the sector the leading edge of the problem.

The decisions are physical and irreversible. A mispriced ad campaign is refunded. A deferred workover that becomes an annulus pressure event, a predictive maintenance model that under-calls a compressor failure, a trading agent that breaches a position limit in a fast market: these decisions convert into steel, hydrocarbons, and money within hours, and they do not convert back. The cost of an undefendable decision scales with the irreversibility of its consequences, and few industries make more irreversible decisions per day than energy.

The sector was already governed by evidence. Energy companies live under regimes that were demanding documentation long before AI arrived: process safety management, well control rules, environmental permitting, market conduct oversight, and the discovery obligations of near-permanent litigation. These regimes share one assumption: when something goes wrong, the operator can produce the record of what it knew and what it decided. AI-influenced decisions are entering exactly these workflows while sitting outside exactly these record-keeping habits. The gap is not that energy lacks a compliance culture. It is that the compliance culture has not yet noticed where the decisions moved.

The accountability question is already being asked. Joint venture partners now ask how AI is used in operations they co-fund. Insurers ask underwriting questions about AI in safety-critical workflows. Acquirers ask in diligence. Regulators, as Section 13 details, have stopped asking and started legislating: Texas operators woke up on January 1, 2026 subject to a comprehensive AI statute with penalties up to $200,000 per uncurable violation.[8] The interval in which "we are still figuring out our AI policy" was an acceptable answer has closed.

The shape of the failure

When the gap surfaces, it almost never looks like a rogue algorithm. It looks like an ordinary decision that cannot be reconstructed. The pattern repeats across incidents with striking consistency: an AI system contributed materially to a decision; the contribution was invisible at the time because it arrived through a chat window, a dashboard, or an embedded feature; the humans involved believed, reasonably, that they were the decision-makers; and after the fact, nobody can establish what the machine recommended, what information it used, or how its influence propagated. The organization is left arguing about a decision it cannot even describe.

> **"The AI did it" has never worked as a defense. "We can show you exactly what happened" has rarely failed as one.**

The remedy is not less AI. The organizations auditing and assessing their AI systems are more than three times as likely to report high value from generative AI, the strongest governance-to-value multiplier measured.[3] The remedy is a discipline built for the layer where the failure actually occurs. Defining that layer precisely is the work of the next two sections.

SECTION 04

Why existing governance fails at the decision layer

Most organizations already run three governance disciplines and are being sold a fourth. All four are necessary. None of them governs the decision.

The instinctive objection to a new governance discipline is that the enterprise already has too many. The objection deserves a serious answer, because it is half right. Energy companies of scale already operate data governance, information governance, and, increasingly, AI governance programs. The answer is that each of these governs a different object, and the object that fails in the incidents described above, the decision, belongs to none of them.

DISCIPLINE	OBJECT GOVERNED	THE QUESTION IT ANSWERS
Data governance	Structured data: quality, lineage, access, cataloging	Is the data correct and controlled?
Information governance	Records and information lifecycle: retention, compliance, e-discovery	Are the records kept and disposed of properly?
AI governance (model-centric)	Models: bias, drift, explainability, model risk	Is the model sound?
Digital Information Governance	The AI-influenced decision	Can the decision be defended?

FIG. 03 · THE GOVERNANCE STACK · DIG GOVERNS THE DECISION LAYER

The seams are where decisions fall through

Each discipline is sound within its boundary, and the failure happens between the boundaries. Walk a single decision through the stack and the gap becomes concrete. A reservoir engineering team evaluates an acquisition. Data governance certifies the production tables. The vendor decline curves imported for the evaluation pass no such control, because they entered through a spreadsheet, not the governed warehouse. A generative model summarizes the combined dataset and produces the type curves that anchor the valuation. Model governance has nothing to say, because the model is a licensed foundation model, not an internally developed one on the model inventory. The investment committee decides, partly on the strength of that summary. Information governance will faithfully retain the committee minutes, which record the decision but not the machine's contribution to it. Every program worked as designed. The decision is still indefensible, because no discipline owned the moment where sourced and unsourced information, machine judgment, and human authority converged.

This is why bolting decision requirements onto the existing disciplines does not work. Data governance can certify every table in the lake and still say nothing about what a model actually read when it produced a recommendation. Model governance can validate a model exhaustively and still say nothing about what a human did with its output. The gap is structural: **the existing disciplines govern assets, and decisions are events.** An event-shaped problem needs an event-shaped discipline, with its own record, its own controls, and its own accountable owners.

What model-centric AI governance misses

The point bears sharpening for AI governance specifically, because it is the newest discipline and the one most often assumed to cover this ground. The dominant frameworks and toolchains of AI governance are model-centric: model inventories, model cards, bias testing, drift monitoring, explainability tooling.[10] All valuable. But a model-centric program has three blind spots that map exactly onto the incident pattern of Section 3.

It cannot see decisions made with models it does not inventory. The overwhelming majority of AI influence in an energy company today arrives through general-purpose assistants, copilots embedded in enterprise software, and vendor features, none of which appear on an internal model inventory. Governing the inventory while the workforce decides with ungoverned tools is governing the map, not the territory.

It stops at the model boundary. Model governance ends where the recommendation is delivered. Everything decisive happens after that: whether a human reviewed it, what they changed, what authority approved it, what was done. A perfectly governed model feeding an unrecorded decision produces an ungoverned outcome with excellent documentation of the wrong thing.

It has no answer for agency. As systems shift from recommending to executing, the governance question shifts from "is the model sound" to "what is this system authorized to do, and what physically prevents it from exceeding that authority." That is not a model property. It is a decision-rights property, and it requires the authority structures specified in Section 14.

None of this retires the existing disciplines. DIG consumes them: it relies on data governance for the trustworthiness of governed sources, on information governance for the retention of decision records, and on model governance for the soundness of inventoried systems. DIG sits at the top of the stack because the decision is where the other three either add up to defensibility or fail to. The next section defines the discipline precisely.

SECTION 05

Definition and vocabulary

Disciplines begin when their words are fixed. This section states the canonical definition of DIG and the working vocabulary used throughout the framework. Section 16 collects the full glossary.

CANONICAL DEFINITION · DIG FRAMEWORK V1.0

Digital Information Governance (DIG®) is the discipline of keeping AI-influenced decisions defensible and auditable: ensuring the information feeding decisions is sourced and trusted, the decisions themselves are traceable to accountable human authority, the organization is accurately represented across AI systems, and its oversight can be proven on demand to regulators, partners, and courts.

Three properties of the definition do the work. It is **decision-centric**: the governed object is the decision event, not the data asset or the model. It is **adversarially framed**: "defensible" presumes a challenger, because the test of governance is not the quiet quarter but the deposition, the audit, the partner dispute, and the incident investigation. And it is **proof-oriented**: a control that existed but cannot be evidenced is, for every purpose that matters, a control that did not exist.

A note on the name. The phrase resembles "information governance," the established records-management discipline, and the resemblance is deliberate but the objects differ. Information governance manages the lifecycle of records and data: creation, storage, retention, disposition. Digital Information Governance manages whether an AI-influenced decision can be defended after the fact. One is records-centric, the other decision-centric. They are allies, not synonyms.

The working vocabulary

AI-influenced decision

Any decision in which an AI system materially shaped the inputs, the analysis, the recommendation, or the execution, whether or not a human made the final call. The definition is deliberately broad and deliberately indifferent to how the influence arrived: a chat window, an embedded copilot, an agentic workflow, and a vendor feature all count. Influence, not architecture, is the test.

Decision integrity

The condition in which a decision can be reconstructed, explained, and defended: its information basis is known, its reasoning is recorded, its authority is named, and its record was captured at decision time. Decision integrity is to decisions what well integrity is to wells: a designed-in property verified by evidence, not an aspiration audited by hope.

Decision record

The contemporaneous artifact that captures a decision: what was recommended, by which system, on what information, who reviewed it, what authority approved it, and what was done. The decision record is the atomic unit of DIG, specified in full in Section 14. It is to the decision layer what the model card is to the model layer.[10]

Decision provenance

The chain of custody of the information feeding a decision: source, transformations, freshness, and trust status at the moment of use. Distinct from data lineage, which tracks assets in a warehouse; provenance attaches to the decision event and includes everything the decision actually consumed, governed or not.

Decision surface

The full set of points in an organization where AI can influence a decision. The decision surface is always larger than the model inventory, usually by an order of magnitude, and mapping it is the first act of a DIG program.

Authority map

The explicit assignment of decision rights across humans and machines: who or what may recommend, who may approve, at what thresholds, with what escalation. In most organizations this exists for capital (the delegation of authority matrix) and nowhere else. DIG extends it to the full decision surface.

Authority-hierarchical validation

The control pattern in which any machine-initiated action is validated against the authority map before execution: the system's delegated authority is checked the way a flight director polls the room, structurally and every time. Specified in Section 14.

Deterministic override

A governance limit enforced by non-probabilistic means. A model asked to police another model produces a probability; a deterministic override produces a wall. Position limits, spend ceilings, and safety envelopes for machine-initiated action must be walls.

Shadow AI

AI use that occurs outside sanctioned tools and visibility: personal accounts, unapproved plug-ins, vendor features enabled by default. Shadow AI is to the decision surface what shadow IT was to the network, and it is where the least defensible decisions are made.

Epistemic drift

The gradual, unexamined shift of an organization's working beliefs toward machine-generated content: summaries citing summaries, models trained on model output, until the original sources are no longer in the loop. Drift is invisible decision by decision and structural in aggregate.

Synthetic consensus

The false confidence produced when multiple AI outputs agree because they share sources, training data, or prompts, and the agreement is mistaken for independent corroboration. Three copilots reading the same stale data lake are one opinion wearing three badges.

Representation baseline

The documented record of how AI systems, search engines, and data environments currently describe the organization: its safety record, its litigation history, its financial condition, its leadership. The baseline is the reference against which misrepresentation is detected and corrected. See Section 8.

Defensible by default

The mature state, Level 5 of the maturity model, in which decision integrity is captured automatically as decisions are made, so that defensibility is the organization's resting posture rather than a project triggered by a challenge.

Information Provenance

Where the information feeding a decision came from, and whether it can be trusted. Without provenance, a decision cannot be reconstructed, because nobody can say what it was based on.

Every decision rests on information, and every defense of a decision begins with its information basis. The first pillar therefore asks the first question a challenger will ask: what did you know, where did it come from, and why did you trust it? Provenance is the chain of custody of decision inputs: the source, the transformations along the way, the freshness at the moment of use, and the trust status assigned to each element.

The concept is old. Energy companies already run provenance disciplines without using the word: chain of custody on custody-transfer measurement, calibration records on instrumentation, chronology of well files. What has changed is that AI systems have become the great launderers of provenance. A model ingests a governed production table, an ungoverned vendor spreadsheet, a six-year-old consultant deck, and a paragraph from the open internet, and returns a single fluent recommendation in which all four are indistinguishable. The output arrives with the confidence of the best source and the reliability of the worst, and the blend is invisible to the person deciding.

The three provenance failures

Unsourced inputs. Information enters the decision with no identifiable origin at all. Most retrieval-augmented and copilot workflows do not record what was actually read to produce a given output; the organization can identify the corpus in general but not the basis of this recommendation in particular. After the fact, "what did the model consider" becomes unanswerable, which makes "was the decision reasonable" unanswerable.

Stale and misgoverned inputs. The information has an origin, but its freshness or quality was unfit for the decision that consumed it. Sensor tags mislabeled in the historian, decline curves of unknown vintage, a policy document superseded two revisions ago. Data governance may even know the asset is stale; provenance failure occurs because the decision consumed it anyway, with no check at the point of use.

Synthetic inputs laundered as fact. Machine-generated content re-enters the information environment and is consumed downstream as if it were an original source: a model-written summary filed into the document store, then retrieved by another model as evidence. This is epistemic drift in mechanism, and left ungoverned it compounds, because each generation of synthesis is trained on the last.

FIELD VIGNETTE · ILLUSTRATIVE COMPOSITE · CAPITAL ALLOCATION

A mid-cap operator evaluates a bolt-on acquisition in the Permian. The technical team uses an AI workflow to integrate the data room: governed internal analogs, the seller's type curves, and a vendor database licensed years earlier and never re-vetted. The model produces integrated type curves that anchor a nine-figure valuation. The committee memo cites "AI-assisted integrated analysis." Eighteen months later, underperformance triggers a board review. The question is not whether the model was clever. It is which decline curves drove the curves that drove the price, and nobody can answer, because the workflow recorded its output but not its diet. The write-down is painful; the discovery that the analysis cannot even be audited is what changes careers.

What the pillar requires

In a governed organization, provenance is enforced at the point of decision, not merely at the point of storage. Four controls carry the pillar. **Source registration:** the information sources that feed consequential decisions are inventoried in a provenance register with owner, freshness expectations, and trust tier. **Capture at retrieval:** AI workflows serving consequential decisions record what they actually read, not just what they produced; the recommendation and its basis travel together into the decision record. **Trust tiering at the point of use:** inputs surface with their status attached, so a recommendation built on untiered or stale sources announces itself before it is relied upon. **Synthetic marking:** machine-generated content is labeled at creation and the label survives storage and retrieval, so synthesis is never mistaken for source.

The standard is not perfection of data, which no operator will ever have. The standard is honesty of basis: at the moment of decision, the decider can see what the recommendation rests on, and after the decision, the organization can prove it. A decision made on admittedly imperfect data, knowingly and with the imperfection recorded, is defensible. A decision made on unknown data is not, even when the data happens to have been good.

Decision Traceability

The record of the decision itself: what was recommended, by which system, on what basis, who reviewed it, and on whose authority it was acted. The difference between "the AI did it" and "a named person decided, and here is the trail."

If provenance governs what fed the decision, traceability governs the decision event itself. It is the pillar most directly aimed at the accountability gap, and the one on which regulators have converged fastest: the EU AI Act's record-keeping requirements for high-risk systems, the human oversight provisions, and the documentation obligations of every emerging regime are all, at bottom, traceability requirements.[5] They encode a principle older than any statute: an organization is accountable for what it decides, and accountability presupposes a record.

Traceability answers five questions, and a decision is traceable only when all five are answerable from a contemporaneous record. **What was recommended**, in the form actually presented, not a paraphrase recalled later. **By which system**, identified specifically enough to know its version and configuration. **On what basis**, linking to provenance. **Who reviewed it**, and what they accepted, modified, or rejected. **On what authority it was acted**, naming the human owner and the delegation under which they approved. The fifth question is the one that separates DIG from logging: a log records that something happened; a decision record establishes who answers for it.

Why traceability fails in practice

The influence is invisible at the time. AI contributions arrive conversationally and get absorbed into human work product. The engineer who accepts a copilot's framing does not experience a decision point; the framing simply becomes the memo. By the time the decision is formalized, the machine's contribution has dissolved into it, unrecorded.

The tools are ephemeral by design. Chat sessions expire. Context windows roll. Vendor features log for product analytics, not for evidence. The default retention posture of the modern AI stack is amnesia, which is precisely the wrong posture for consequential decisions.

Accountability diffuses across the seam. When a recommendation crosses from machine to human, each side assumes the other holds the accountability. The human treats the model as an authoritative analyst; the organization treats the human as the decider; the record captures neither. In post-incident interviews this appears as the signature sentence of the ungoverned decision layer: "I assumed the system had checked that."

FIELD VIGNETTE · ILLUSTRATIVE COMPOSITE · WELL INTEGRITY

An offshore team is working the annual plan. A completions engineer asks an approved copilot to summarize eighteen months of integrity data on a producing well and assess whether the scheduled workover can slip a quarter. The summary is competent, the recommendation confident, the schedule pressure real. The deferral is approved in a Tuesday planning meeting and recorded in the plan as a date change. Five months later, sustained casing pressure forces a shut-in and the regulator's inquiry arrives with standard questions: who made the deferral decision, and on what engineering basis? The plan shows a date. The meeting minutes show attendance. The chat session that contained the actual analysis aged out of retention months ago. The company now must defend an engineering judgment it cannot produce, made by a person it cannot definitively name, on information it cannot reconstruct. The well gets repaired. The credibility does not.

What the pillar requires

Traceability is achieved by making the decision record, specified in Section 14, a routine byproduct of deciding rather than an act of documentation discipline. Three design rules govern. **Capture at decision time:** the record is created when the decision is made, by the workflow itself wherever possible; contemporaneous capture is what converts a story into evidence. **Proportionality by tier:** not every decision warrants the same record. The decision inventory tiers the surface, and Tier 1 decisions (safety-critical, market-facing, capital-committing) carry the full record while routine decisions carry a lighter trace. Governance that ignores proportionality collapses under its own weight and gets bypassed. **Named ownership, always:** every consequential AI-influenced decision has exactly one accountable human owner recorded by name. Not a committee, not a role mailbox, not "the system." The single name is the pillar's load-bearing element, and organizations that resist it are usually discovering that their accountability was already diffuse before AI arrived; the machines merely made it visible.

> **A log records that something happened. A decision record establishes who answers for it.**

SECTION 08 · THE FRAMEWORK · PILLAR 03

Representation Integrity

Keeping the company accurately represented across the AI systems, search engines, and data environments that now describe it to the world. A misrepresentation is no longer a nuisance; it is material that counterparties, regulators, and courts can consume.

The first two pillars govern decisions the organization makes. The third governs decisions made **about** the organization, by others, using AI. This is the pillar the discipline could not have needed a decade ago, and the one executives most consistently underestimate, because the exposure sits outside their systems entirely.

Here is the structural change. When a bank's credit committee, a supermajor's supply chain team, an insurer's underwriter, a joint venture partner's diligence group, or a plaintiff's counsel wants to understand a company today, their first analytical act increasingly runs through an AI system. The answer they receive is assembled from whatever the models have absorbed: filings, news, litigation records, industry databases, stale directories, and the statistical residue of training. Nobody at the company sees the answer. Nobody approves it. And yet it functions, operationally, as the company's representation of record for that decision. If the answer conflates the company with a similarly named entity, resurrects a resolved consent decree as current, misstates the safety record, or simply describes the business as it was five years ago, the error propagates into someone else's decision with the company's name on it.

Energy companies are unusually exposed. The sector's naming conventions practically invite conflation: decades of entities sharing surnames, basins, and suffixes, restructured and renamed through cycles. Its public record is adversarial by default, dense with litigation, enforcement actions, and activist coverage that models absorb without the disposition history. And its commercial life runs on counterparty assessment: prequalification, bid evaluation, JV diligence, insurance placement, credit review, all of which are now AI-assisted on the other side of the table.

FIELD VIGNETTE · ILLUSTRATIVE COMPOSITE · COUNTERPARTY DILIGENCE

A midstream company reaches the final round of a competitive process to anchor a new gathering system. The customer's diligence team, working at speed, asks its AI research assistant to profile the finalists' compliance histories. The assistant, drawing on public records, attributes to the company a pipeline safety consent order that in fact belongs to a similarly named, long-defunct entity two states away. The error is plausible, specific, and wrong. It never appears in a document the company can rebut, because it appears in a briefing the company never sees. The award goes elsewhere on "risk profile" grounds. The company loses to a sentence it does not know was written. Months later, the same conflation surfaces in an insurance renewal, and this time, because a representation baseline is finally in place, it is caught, documented, and corrected at the source.

What the pillar requires

Baseline: establish the representation baseline, a documented survey of how the major AI systems and data environments currently describe the company across the dimensions that drive counterparty decisions: identity and corporate history, safety and environmental record, litigation and enforcement, financial condition, leadership, and operational footprint. The baseline converts an invisible exposure into an inspectable artifact. **Monitoring:** re-survey on a fixed cadence and around material events (transactions, incidents, leadership changes), because representations move when the record moves. **Correction at the source:** when misrepresentation is found, remediate the underlying record it derives from: the stale registry, the unlinked disposition, the ambiguous public filing, the outdated web presence. Correcting sources is slower than complaining about outputs and is the only remedy that persists. **Evidence:** retain the baseline, the findings, and the corrections. When a misrepresentation causes commercial harm, the company that can show a monitored baseline and a correction trail is in a categorically different position from the company discovering the problem in the deposition.

One boundary matters. Representation integrity is the discipline of keeping the record **accurate**, not the art of making it flattering. The pillar's standard is the same standard the other pillars apply to the company's own decisions: truthfulness that survives scrutiny. An organization that games its representation has not governed the risk; it has become the risk.

SECTION 09 · THE FRAMEWORK · PILLAR 04

Audit Readiness

The ability to prove, on demand, that the first three pillars held. Audit readiness is the posture that turns "we have AI policies" into "we can show our oversight worked."

The fourth pillar is the framework's enforcement of its own honesty. Provenance, traceability, and representation integrity describe conditions; audit readiness demands that the conditions be provable, on demand, to a challenger who did not schedule the challenge. It is the difference between an organization that believes it is governed and an organization that can demonstrate it, and every seasoned operator knows how large that difference turns out to be under pressure.

The pillar rests on a claim this paper treats as foundational: **an unprovable control is not a control.** Policies that exist in a document repository, reviews that happened but left no artifact, oversight that was real but is now a matter of testimony: all of these collapse to zero at exactly the moment they are needed. Regulators have internalized this. The EU AI Act does not ask whether high-risk systems are overseen; it requires logging, technical documentation, and conformity evidence.[5, 6] ISO/IEC 42001 does not certify intentions; it certifies a management system with auditable records.[7] TRAIGA's enforcement mechanics turn on what an operator can document when the Attorney General's inquiry arrives.[8] The regulatory era rewards one posture: evidence as a byproduct of operations.

The audit scramble, and its cost

Most organizations meet challenges with what this paper calls the audit scramble: a weeks-long forensic project to reconstruct, from fragments across systems, what a decision record would have captured in minutes. The scramble has three costs. The direct cost is the project itself, senior people diverted to archaeology. The evidentiary cost is that reconstructions are arguments where records are facts; a trail assembled after the challenge is inherently impeachable. The strategic cost is the largest: an organization that fears the question begins declining the delegation. It slows its AI adoption not because the technology fails but because leadership, rationally, will not extend authority it cannot verify. This is the mechanism behind the audit dividend: organizations that regularly audit and assess their AI systems are more than three times as likely to report high value from generative AI.[3] Provability is not the tax on delegation. It is the license for it.

FIELD VIGNETTE · ILLUSTRATIVE COMPOSITE · TRADING & MARKETS

A gas marketing desk runs an agentic execution layer: models monitor flows and weather, propose hedges, and execute within delegated limits. It works, and the desk is proud of it. Then internal audit, prompted by a board risk committee question, asks for something simple: evidence that every machine-initiated execution in the last quarter stayed within delegated authority, with the authority chain for each. The evidence exists, in the sense that atoms exist: order logs in the EMS, limits in a risk system, model versions in a repository, approvals in email, none joinable without manual work. Three weeks and two consultants later, the desk produces a reconstruction with four unexplained residual items, which is four more than a control environment allows. Nothing was actually wrong. The desk still loses its expansion request for two quarters, because the committee correctly distinguishes between "probably fine" and "proven." The control plane that would have answered the question in an afternoon is approved, ironically, as part of the remediation.

What the pillar requires

Evidence by design: decision records, provenance capture, authority validations, and representation baselines are generated as operational byproducts, not compiled as projects. If producing evidence requires a project, the organization is at Level 3 at best. **Joinability:** the artifacts share identifiers so that a decision can be walked end to end: input to recommendation to review to authority to action. Evidence that cannot be joined is a scavenger hunt with better branding. **Testing:** controls are exercised, not assumed. The organization periodically pulls a decision thread on demand, times the retrieval, and treats failure to produce as a finding, exactly as it treats a failed pressure test. **Scope honesty:** the audit posture covers the actual decision surface, including the embedded and vendor AI where most influence lives, not just the flagship systems that demo well.

> **The mature organization does not prepare for audits. It operates in a way that makes audit a retrieval task.**

SECTION 10

The seven principles of decision governance

The pillars say what must be true. The principles say how a governed organization behaves. They are written to be quoted in a policy, tested in an audit, and invoked in a disagreement.

PRINCIPLE 01

The decision is the unit of governance.

Govern decisions, not models. Model soundness, data quality, and record retention are inputs to a governed decision, never substitutes for one. Any control, metric, or review that cannot be expressed in terms of a decision's defensibility is supporting infrastructure, not governance.

IN PRACTICE The governance program's core inventory is a decision inventory, and its core question in every review is: which decisions does this make more defensible?

PRINCIPLE 02

Every consequential decision has one named human owner.

Accountability is a person, not a system, a committee, or a vendor. Machine participation in a decision never dilutes human ownership of it; if anything it concentrates it, because the owner now answers for the delegation as well as the outcome. Where no owner can be named, the decision is not authorized to occur.

IN PRACTICE The decision record carries exactly one accountable name per decision. "The system decided" is recorded, when true, as "this owner delegated to this system under this authority."

PRINCIPLE 03

Provenance before inference.

No consequential decision rests on information whose origin, freshness, and trust status are unknown at the point of use. Imperfect information, known and recorded, is defensible; unknown information is not, regardless of how it turns out. Machine-generated content is marked as such wherever it travels.

IN PRACTICE Recommendations surface with their basis attached. A recommendation that cannot state its sources is treated as an unsourced opinion, whatever its fluency.

PRINCIPLE 04

The record is created at decision time.

Contemporaneous capture is the difference between evidence and argument. A record assembled after a challenge is a reconstruction, and a reconstruction is impeachable. Capture is therefore built into the decision workflow itself, proportionate to the decision's tier, so that recording is a byproduct of deciding rather than a demand on discipline.

IN PRACTICE If creating the record requires remembering to create the record, the design has failed for every decision that matters.

PRINCIPLE 05

Authority is explicit, hierarchical, and validated.

Decision rights are written down: who and what may recommend, approve, and execute, at what thresholds, with what escalation. Machine-initiated actions are validated against that hierarchy before execution, structurally and every time, the way a flight director polls the loop. Authority that exists only in culture is authority that fails silently.

IN PRACTICE The authority map covers machines as well as people, and no agentic system executes outside a delegation it can be checked against.

PRINCIPLE 06

Overrides are deterministic.

Hard limits on machine-initiated action are enforced by non-probabilistic mechanisms. A model supervising a model produces a likelihood; a governance limit must produce a wall. Position limits, spend ceilings, safety envelopes, and prohibited actions are enforced by controls that no quality of argument, human or machine, can talk past.

IN PRACTICE For every agentic system, the organization can point to the specific deterministic mechanism that stops it at its limit, and has watched that mechanism fire in a test.

PRINCIPLE 07

Defensibility is proven, not presumed.

Controls are exercised against live decisions on a fixed cadence, and failure to produce evidence is treated as a finding with the same seriousness as a failed integrity test. The organization's representation across external AI systems is part of the proven surface. What cannot be demonstrated is assumed absent.

IN PRACTICE Someone periodically pulls a decision thread without warning and times how long the full trail takes to produce. The number is reported to leadership like any other integrity metric.

The principles are deliberately few and deliberately testable. A policy manual can elaborate them; an auditor can verify them; an executive can recite them. An organization that honors all seven is, by construction, operating the four pillars. An organization that honors six is usually about to discover which one it skipped.

SECTION 11

The DIG Maturity Model

A five-level scale measuring one thing: when an AI-influenced decision is challenged, how readily can the organization defend it? Most AI maturity models rate the model, or the adoption. This one rates the decision.

The maturity model exists because "are we governed?" is not a yes-or-no question, and treating it as one produces either complacency or paralysis. The five levels describe observable postures, each defined by a single diagnostic: **when is the defensibility record created?** At Level 1 it is never created. At Level 3 it is assembled on request. At Level 5 it is captured at decision time and verified continuously. The climb is from reconstructing defensibility under pressure to producing it on demand, which is precisely the posture modern regulation rewards.[5, 7, 8]

L1 Ad hoc

AI shapes decisions, but nothing durable is recorded. When a decision is questioned, the organization reconstructs it from memory, if at all. Shadow AI is pervasive and invisible. No decision is defensible by design; some are defensible by luck.

MARKER: A CHALLENGED DECISION TRIGGERS INTERVIEWS, NOT RETRIEVALS.

L2 Aware

Policies exist and some systems log activity, but coverage is uneven and unowned. Defensibility depends on which individual made the decision and whether they happened to keep a record. The organization can describe its intentions and cannot demonstrate its practice.

MARKER: THE ANSWER TO "WHO DECIDED" VARIES BY WHO IS ASKED.

L3 Defined

The four pillars are standard practice for high-stakes decisions. The decision inventory exists and is tiered, provenance is tracked for Tier 1 decisions, decision records are captured, owners are named, and most consequential decisions can be reconstructed on request. This is the first level a regulator would call governed, and the appropriate first target for any operator starting the climb.

MARKER: A TIER 1 DECISION CAN BE RECONSTRUCTED ON REQUEST, IN DAYS.

L4

Managed

Controls are tested, not assumed. The organization is audit-ready on demand, decision coverage is measured rather than hoped for, authority validations run on machine-initiated actions, and the representation baseline is monitored on cadence. Evidence production is a retrieval task with a known clock time.

MARKER: AN UNANNOUNCED THREAD-PULL PRODUCES THE FULL TRAIL IN HOURS.

L5

Defensible by default

Decision integrity is captured automatically as each decision is made. Audit is continuous, representation is governed, deterministic overrides bound every agentic system, and defensibility is the organization's resting state rather than a scramble after a challenge. The governance layer is infrastructure, invisible in the same way custody transfer measurement is invisible: always on, rarely discussed, absolutely relied upon.

MARKER: THE AUDIT IS A QUERY. THE ANSWER IS ALREADY WRITTEN.

Reading the model honestly

Three usage notes keep the model useful. **Score by pillar, not in aggregate.** Organizations are routinely Level 3 on traceability for capital decisions and Level 1 on representation integrity, because nobody owns the latter. A single blended number hides exactly the gap that will surface first. **Level 3 is the regulatory floor, not the ceiling.** Operators making consequential AI-influenced decisions in safety-critical or market-facing contexts should target Level 4 or above, because "reconstructable in days" is adequate for an audit and inadequate for an incident. **Do not confuse tooling with level.** Buying a governance platform at Level 1 produces a Level 1 organization with a dashboard. The levels describe operating behavior: named owners, live records, tested controls. Tools accelerate the climb; they do not constitute it.

Placement is an assessment exercise: each pillar is scored against the level markers using live decisions pulled from the actual surface, not policy documents. The output that matters is not the score but the gap: the specific distance, in artifacts and behaviors, between the current level and the next one. Section 15 turns that gap into a sequence.

SECTION 12

The risk taxonomy

Four risk families, one per pillar, and two amplifiers that act on all four. The taxonomy exists so that risk registers, audit plans, and board reporting can name what they are looking at.

AI risk discussions fail most often by being either too abstract (existential hand-wringing) or too narrow (model bias metrics). The DIG taxonomy classifies risk by the pillar whose failure produces it, which keeps every entry connected to a control and an owner. Each family is stated with its mechanism, its characteristic energy-sector expression, and its signature: the tell that shows up in the post-mortem.

RISK FAMILY	PILLAR	MECHANISM AND ENERGY EXPRESSION
Contaminated basis	Information Provenance	Decisions consume unsourced, stale, or synthetic information without knowing it. Expression: valuations anchored on unvetted vendor curves; integrity assessments fed by mislabeled historian tags; model output re-ingested as source. Signature: "we cannot say what the analysis was based on."
Unreconstructable decision	Decision Traceability	The decision event leaves no usable trail: no recommendation as issued, no named reviewer, no authority chain. Expression: deferrals, dispatch choices, and screening decisions absorbed from copilots into work product; shadow AI on personal accounts. Signature: "the answer to who decided depends on who you ask."
Misrepresentation	Representation Integrity	External AI systems describe the company inaccurately and third parties decide on that description. Expression: entity conflation in counterparty diligence; resurrected enforcement history in insurance and credit review; stale operational footprint in prequalification. Signature: "we lost to a sentence we never saw."
Unprovable oversight	Audit Readiness	Controls existed but cannot be evidenced; the challenge triggers a reconstruction project instead of a retrieval. Expression: the audit scramble; residual unexplained items in agentic activity; regulatory responses built from fragments. Signature: "we are confident it was fine."

FIG. 04 · THE FOUR RISK FAMILIES, MAPPED TO PILLARS

The two amplifiers

Two forces act multiplicatively across all four families. They deserve their own line in any risk register because they change the rate at which the families compound.

Amplifier one: decision velocity. Agentic systems collapse the interval between recommendation and action, and multiply the number of decisions made per unit time. Every ungoverned property of the decision layer scales linearly with decision count: more unrecorded decisions, more unvalidated authority, more residual items per quarter. An organization with a manageable accountability gap at human decision speed acquires an unmanageable one at machine decision speed without changing anything else. This is why the taxonomy treats agency not as a new risk family but as a throttle on all of them, and why Section 14's authority architecture is the control that matters most over the planning horizon.

Amplifier two: epistemic drift. The information environment itself degrades as machine-generated content recirculates: summaries citing summaries, models consuming model output, synthetic consensus mistaken for corroboration. Drift acts slowly and directionally, like corrosion. Its danger is that it lowers the quality of the basis for every decision simultaneously, while making the basis look more consistent, because homogenized sources agree with each other. The control is provenance discipline applied continuously: synthetic marking, source tiering, and the deliberate preservation of primary sources in the loop.

Using the taxonomy

The taxonomy is designed to be operationalized three ways. In the **risk register**, the four families become standing entries with pillar controls as mitigations and the amplifiers as velocity factors, which gives AI risk a stable structure instead of a rotating list of headlines. In the **audit plan**, each family implies a test: pull a decision and demand its basis; pull a decision and demand its owner; pull the representation baseline and check its age; pull a thread and time the retrieval. In **board reporting**, the families give the risk committee a fixed vocabulary, so the conversation matures from "what is our AI risk" to "contaminated-basis exposure is down, decision velocity is up 4x, and here is the coverage number." Named risks get managed. Unnamed risks get discovered.

SECTION 13

The regulatory environment

This is not a compliance paper, and DIG is not a compliance framework. But the statutes now arriving all converge on the same demand, and an operator should understand why practicing DIG satisfies obligations it was not written from.

Every major AI governance regime, read past its definitions and schedules, asks the operator for the same four things: know what fed your systems, keep records of what they did and who oversaw them, do not deceive, and be able to prove the foregoing. Those are the four pillars, discovered independently by legislators. This convergence is why DIG is practiced once and mapped many times: an organization at Level 3 or above finds each new regime a documentation exercise rather than a program.

Texas: TRAIGA

For energy, the statute that matters first is at home. The Texas Responsible Artificial Intelligence Governance Act (HB 149, 89th Legislature), effective January 1, 2026, made Texas one of the first US states with a comprehensive AI law.[8] Three features shape operator exposure. **Enforcement is real:** the Attorney General enforces, with civil penalties up to $200,000 per uncurable violation and up to $40,000 per day for continuing violations. **The obligations are conduct-shaped:** the Act's prohibitions and duties turn on how AI systems are used and represented, which means exposure runs through the decision layer, exactly where most companies have no records. **Cure depends on evidence:** the Act's structure rewards operators who can promptly demonstrate what their systems did and what oversight existed. An operator that cannot reconstruct its AI-influenced conduct cannot effectively use a cure period, because it cannot establish what happened, when, or that it stopped. For a Texas-headquartered industry, TRAIGA converts the accountability gap from a governance concern into a quantified, per-violation, per-day exposure with the state's chief law enforcement officer as counterparty.

$200K

TRAIGA CIVIL PENALTY, PER UNCURABLE VIOLATION

$40K

PER DAY, CONTINUING VIOLATIONS

€35M / 7%

EU AI ACT MAXIMUM: FINE OR SHARE OF GLOBAL TURNOVER

FIG. 05 · THE PRICE OF THE GAP, AS LEGISLATED · SOURCES: SEC. 17

The wider map

REGIME	STATUS	WHAT IT DEMANDS, IN DIG TERMS

TRAIGA (Texas HB 149)	Effective Jan 1, 2026; AG-enforced	Conduct-level accountability for AI use; evidence sufficient to establish and cure. Leans on Traceability and Audit Readiness.[8]
EU AI Act (2024/1689)	In force; obligations phasing through 2027; prohibited practices since Feb 2025	For high-risk systems: automatic event logging supporting traceability (Art. 12), technical documentation (Art. 11), human oversight (Art. 14); penalties to €35M or 7% of global turnover (Art. 99). Effectively mandates decision records for covered systems.[5, 6]
NIST AI RMF 1.0	Voluntary; de facto US reference	Govern, Map, Measure, Manage. States that trustworthy AI depends on accountability, and accountability presupposes transparency: the principle behind Decision Traceability. DIG operationalizes the Govern and Manage functions at the decision layer.[4]
ISO/IEC 42001:2023	Certifiable standard; first AI management system standard	An auditable AI management system: documented processes, records, continual improvement. DIG artifacts (inventory, records, test logs) are the operational evidence an audit consumes.[7]
US state laws (Colorado and successors)	Phasing in from 2026	Duties around consequential automated decisions and algorithmic discrimination; documentation and impact assessment requirements that presume reconstructable decisions.

FIG. 06 · REGULATORY MAP · ONE DISCIPLINE, MAPPED MANY TIMES

The strategic reading of this table is not "compliance burden rising," although it is. It is that **the record-keeping posture regulators reward and the operating posture that captures the audit dividend are the same posture.** The organization that builds the decision layer once, for its own reasons, gets the statutes as a discount. The organization that builds regime by regime buys the same infrastructure several times, each time under deadline, and never gets the operating benefit.

SECTION 14

The operating model

Frameworks fail as posters and survive as operating models. This section specifies the roles, the artifacts, the cadence, and the control architecture: what a leadership team actually stands up, staffs, and runs.

Roles

DIG runs on five roles. Titles can flex to the organization; the accountabilities cannot.

ROLE	ACCOUNTABILITY
Accountable Executive	Owns the program at the leadership table. Approves the authority map, receives the coverage and thread-pull metrics, and answers to the board for the decision layer. Typically the COO or an operations-facing C-suite officer; deliberately not delegated to a working group.
Decision Owners	The named humans accountable for each decision class in the inventory (well intervention deferrals, hedge execution, vendor screening). They approve within delegated authority and their names appear on the records. This role is not new headcount; it is existing accountability made explicit.
Governance Lead	Runs the program: maintains the decision inventory and authority map, operates the cadence, coordinates assessments, reports the metrics. The one substantially new seat in most organizations.
Governance Engineers	Build capture into workflows: decision-record instrumentation, provenance capture at retrieval, authority validation, deterministic overrides. Sit with the automation and platform teams, not in a policy function, because the controls are built, not written.
Assurance Partner	Internal audit or an external assessor. Runs the unannounced thread-pulls, tests the overrides, scores maturity by pillar, and reports findings outside the program's own chain.

Artifacts

Six artifacts constitute the program's physical existence. If they are current, the program is real; if they are stale, the program is a slide.

ARTIFACT	CONTENT AND TEST OF HEALTH

Decision Inventory	The mapped decision surface, tiered by consequence (Tier 1: safety-critical, market-facing, capital-committing; Tier 2: significant operational; Tier 3: routine). Healthy when it includes embedded and vendor AI, not just flagship systems.
Authority Map	Decision rights across humans and machines: recommend, approve, execute, thresholds, escalation. Healthy when every agentic system appears on it with explicit limits.
Decision Records	The contemporaneous per-decision artifacts, specified below. Healthy when produced by workflows rather than willpower, and joinable end to end.
Provenance Register	Sources feeding Tier 1 decisions, with owner, freshness expectation, and trust tier. Healthy when a recommendation can be traced to registered sources.
Representation Baseline	The documented external representation survey and its correction log. Healthy when younger than one quarter and refreshed after material events.
Control Test Log	Results of thread-pulls, override tests, and coverage measurements, with findings and remediations. Healthy when it contains failures, because a test log with no failures records a testing program that is not testing.

The Decision Record, specified

The atomic artifact of the discipline. Eleven fields, captured at decision time, proportionate in depth to tier.

ARTIFACT SPEC · DECISION RECORD · DIG V1.0

01 · DECISION ID	Unique identifier; joins all related artifacts
02 · TIMESTAMP	Date and time of decision, captured contemporaneously
03 · CLASS / TIER	Decision class per inventory; consequence tier
04 · DECISION OWNER	Named accountable human
05 · SYSTEMS	AI systems involved, with version and configuration reference
06 · RECOMMENDATION	Machine output as issued, verbatim or by durable reference
07 · BASIS	Provenance references: what was actually read or retrieved
08 · HUMAN REVIEW	Reviewer(s); disposition: accepted / modified / rejected, with deltas
09 · AUTHORITY	Delegation invoked; threshold checked; validation result
10 · OVERRIDE STATUS	Deterministic limits evaluated; any override or escalation fired
11 · ACTION & FOLLOW-UP	What was executed; linked outcomes and review triggers

The control architecture for agentic AI

Recommendation-era controls assumed a human between the machine and the world. Agentic systems remove that assumption, so the control must move into the execution path. The architecture DIG specifies is a **governance control plane**: a layer that sits between AI systems and the systems of action they command (the trading gateway, the maintenance management system, the procurement platform), through which machine-initiated actions must pass.

The control plane does three things, and only three, which is why it can be trusted. **It validates authority hierarchically:** every machine-initiated action is checked against the authority map before execution: is this system delegated to take this action, at this magnitude, in this context, and does the delegation chain terminate in a named human owner? The check is structural and unconditional, the GO/NO GO poll performed in milliseconds. **It enforces deterministic overrides:** hard limits (position ceilings, spend caps, safety envelopes, prohibited action classes) are evaluated by non-probabilistic logic that no model output can argue past. Probabilistic systems propose; deterministic systems dispose. **It emits decision records as exhaust:** every validation, every action, every override event writes the record automatically, which is how an organization reaches Level 5 without asking anyone to fill in a form. Methods within this architecture, including authority-hierarchical validation and non-probabilistic governance override, are the subject of pending patent applications; the architectural pattern is stated here because the discipline requires it regardless of whose implementation an operator runs.

Cadence

RHYTHM	WHAT HAPPENS
At decision time	Records captured; authority validated; overrides evaluated. Continuous and automatic for instrumented workflows.
Monthly, operational	Governance Lead reviews coverage (share of Tier 1 decisions with complete records), new decision classes, shadow AI findings, and register freshness.
Quarterly, executive	Accountable Executive reviews the metrics that matter: coverage, thread-pull retrieval time, override test results, representation baseline age, maturity movement by pillar. One page, five numbers, no theater.
Annually, assurance	Assurance Partner conducts the full assessment: maturity scored by pillar against live decisions, findings reported outside the program chain, targets set for the year.

A closing note on scale. The operating model reads heavier than it runs. For a mid-cap operator, the standing cost is one Governance Lead, a fraction of existing engineering capacity, and the instrumentation of a few dozen workflows that matter. The design principle throughout is minimal effective dose: govern the top of the decision inventory completely rather than the whole surface thinly. Twenty fully

governed Tier 1 decision classes are worth more, in defensibility and in regulatory posture, than two hundred half-governed ones.

SECTION 15

The adoption path

From wherever the organization stands to Level 3 in a year, with the first proof inside ninety days. The sequence is designed around a rule that survives contact with real organizations: prove the discipline on the decisions that matter before attempting the decisions that don't.

Days 1 to 90: make the gap visible, then close it where it counts

Weeks 1 to 3: map the surface. Build the first decision inventory through structured interviews and tooling review, deliberately hunting the embedded and shadow AI where influence actually lives. Tier the inventory by consequence. The output is one artifact and one number: the inventory, and the count of Tier 1 decision classes currently leaving no usable trail. That number is the program's founding metric, and in most organizations it lands between "uncomfortable" and "unpresentable."

Weeks 3 to 6: name the owners, run the baseline. Assign a named Decision Owner to every Tier 1 class; where naming proves hard, the difficulty is itself a finding about pre-existing accountability. In parallel, commission the first representation baseline and the TRAIGA exposure screen: which Tier 1 decisions touch conduct the Act reaches, and what evidence exists today. These two workstreams produce the early discoveries that convert executive sponsorship from polite to genuine.

Weeks 6 to 13: instrument the top twenty. Stand up decision records for the top Tier 1 classes, built into the workflows (forms where necessary, capture where possible), with provenance references for the sources they consume. Draft the authority map for these classes, including any agentic system that executes. Close the quarter with the first thread-pull: the Accountable Executive picks a live decision, and the team produces the full trail with a stopwatch running. The retrieval time, compared against the founding metric, is the ninety-day proof.

Months 4 to 12: from proof to posture

Months 4 to 6: extend records across the full Tier 1 inventory; complete the authority map; deploy deterministic overrides for every agentic execution path, and test-fire each one. Establish the monthly operational review and the provenance register.

Months 7 to 9: begin unannounced thread-pulls on cadence; refresh the representation baseline and remediate at the source; extend instrumentation into Tier 2 where the engineering is cheap. Publish the first quarterly executive page: coverage, retrieval time, override tests, baseline age, maturity by pillar.

Months 10 to 12: the Assurance Partner runs the first full assessment against live decisions. Target posture at the year: Level 3 across all four pillars, Level 4 on traceability for the instrumented Tier 1

surface, and a standing answer to the only question that matters: pull any consequential decision from the last two quarters, and the trail arrives in hours.

The three failure modes, and their prevention

The policy-first failure: twelve months of framework documents, zero instrumented decisions. Prevention: no policy artifact is approved until the first twenty decision classes are producing records; the program's currency is trails, not documents. **The boil-the-ocean failure:** attempting the full surface at once, collapsing under form fatigue, and getting bypassed by the operators it was meant to protect. Prevention: minimal effective dose, enforced by the tiering; depth before breadth, always. **The tooling-first failure:** buying a platform and mistaking deployment for governance. Prevention: the maturity model scores behavior, not licenses; the unannounced thread-pull is unforgiving of dashboards with nothing behind them.

> **The organization does not need to govern everything. It needs to be unable to be surprised by anything it would have to defend.**

SECTION 16

Glossary

The canonical vocabulary of the DIG framework, v1.0. Terms defined here control over informal usage elsewhere.

AI-influenced decision

Any decision in which an AI system materially shaped the inputs, analysis, recommendation, or execution, regardless of whether a human made the final call or how the influence arrived.

Accountability gap

The condition, named in the research literature, in which harms from deployed AI systems cannot be traced back to their source and therefore cannot be assigned, remedied, or learned from.

Agentic AI

AI systems that take actions in the world (executing, transacting, dispatching) rather than only producing recommendations for humans to act on.

Audit dividend

The measured association between regular AI auditing and value capture: organizations that regularly audit and assess their AI are over three times as likely to report high value from generative AI.

Audit Readiness

Pillar four. The ability to prove, on demand, that AI-influenced decisions met their obligations: evidence as an operational byproduct rather than a project.

Audit scramble

The reconstruction project triggered by a challenge in an organization without decision records: weeks of forensic assembly producing an argument where a record would have produced evidence.

Authority map

The explicit assignment of decision rights across humans and machines: who or what may recommend, approve, and execute, at what thresholds, with what escalation.

Authority-hierarchical validation

The control pattern in which every machine-initiated action is validated against the authority map before execution, with the delegation chain terminating in a named human owner.

Decision inventory

The mapped and tiered catalog of the organization's decision surface: the foundational artifact of a DIG program.

Decision integrity

The condition in which a decision can be reconstructed, explained, and defended: known basis, recorded reasoning, named authority, contemporaneous record.

Decision Owner

The single named human accountable for a decision class, whose name appears on its records.

Decision provenance

The chain of custody of the information feeding a decision: source, transformations, freshness, and trust status at the moment of use.

Decision record

The contemporaneous artifact capturing an AI-influenced decision across eleven specified fields; the atomic unit of the DIG framework. See Section 14.

Decision surface

The full set of points where AI can influence a decision in the organization; always larger than the model inventory.

Defensible by default

Maturity Level 5: decision integrity captured automatically at decision time, with continuous audit, so defensibility is the organization's resting state.

Deterministic override

A governance limit enforced by non-probabilistic mechanisms that no model output can argue past: a wall, not a likelihood.

Digital Information Governance (DIG®)

The discipline of keeping AI-influenced decisions defensible and auditable. See the canonical definition, Section 5.

Epistemic drift

The gradual shift of an organization's working beliefs toward machine-generated content as synthesis recirculates and primary sources fall out of the loop.

Governance control plane

The architectural layer between AI systems and systems of action that validates authority, enforces deterministic overrides, and emits decision records automatically.

Information governance

The established records-management discipline governing the information lifecycle: creation, storage, retention, disposition. Records-centric; distinct from DIG, which is decision-centric.

Information Provenance

Pillar one. Where the information feeding a decision came from, and whether it can be trusted.

Representation baseline

The documented survey of how external AI systems and data environments currently describe the organization, maintained with a correction log.

Representation Integrity

Pillar three. Keeping the company accurately represented across AI systems, search engines, and data environments.

Shadow AI

AI use outside sanctioned tools and visibility: personal accounts, unapproved plug-ins, vendor features enabled by default.

Synthetic consensus

False corroboration produced when multiple AI outputs agree because they share sources or training, and the agreement is mistaken for independence.

Thread-pull

The assurance test in which a live decision is selected without warning and the complete trail is produced against a clock; the retrieval time is a standing integrity metric.

Tiering

Classification of decision classes by consequence (Tier 1: safety-critical, market-facing, capital-committing; Tier 2: significant operational; Tier 3: routine), setting proportionate record depth.

Decision Traceability

Pillar two. The contemporaneous record of what was decided, by what, on what basis, by whom reviewed, and on whose authority.

SECTION 17

References

1. Stanford University HAI. *The 2025 AI Index Report*, Responsible AI chapter (2024 data). hai.stanford.edu/ai-index/2025-ai-index-report/responsible-ai
2. Deloitte. *State of AI in the Enterprise*, 2026 edition (survey of 3,235 leaders across 24 countries). deloitte.com/us/en/insights/topics/emerging-technologies/ai-agents-scaling-faster.html
3. Gartner. Press release, November 4, 2025: regular AI system assessments and the likelihood of high GenAI value (survey of 360 organizations). gartner.com/en/newsroom/press-releases/2025-11-04
4. National Institute of Standards and Technology. *AI Risk Management Framework (AI RMF 1.0)*, NIST AI 100-1, January 26, 2023. DOI 10.6028/NIST.AI.100-1. nvlpubs.nist.gov/nistpubs/ai/nist.ai.100-1.pdf
5. European Union. *Regulation (EU) 2024/1689* (EU AI Act), Articles 11 (technical documentation), 12 (record-keeping and traceability), 14 (human oversight). eur-lex.europa.eu/eli/reg/2024/1689/oj
6. European Union. *Regulation (EU) 2024/1689*, Article 99 (penalties). artificialintelligenceact.eu/article/99
7. ISO/IEC 42001:2023. *Information technology, Artificial intelligence, Management system*. iso.org/standard/81230.html
8. State of Texas. *Texas Responsible Artificial Intelligence Governance Act* (TRAIGA), HB 149, 89th Legislature, enrolled. Effective January 1, 2026. capitol.texas.gov/tlodocs/89R/billtext/pdf/HB00149F.pdf
9. Raji, I.D., Smart, A., et al. "Closing the AI Accountability Gap: Defining an End-to-End Framework for Internal Algorithmic Auditing." ACM FAT* 2020. arxiv.org/abs/2001.00973
10. Mitchell, M., et al. "Model Cards for Model Reporting." ACM FAT* 2019. arxiv.org/abs/1810.03993
11. Stanford University HAI. *The 2026 AI Index Report*, Responsible AI chapter (2025 data; AI Index and McKinsey survey; AI Incidents Database). hai.stanford.edu/ai-index/2026-ai-index-report/responsible-ai
12. United States Patent and Trademark Office. Trademark Reg. No. 8147558, *Digital Information Governance / DIG*, owner Matthew Bertram. uspto.report/TM/99559923
13. Digital Information Governance reference site: canonical definition, framework, maturity model, and regulatory mapping, maintained in alignment with this paper. digitalinformationgovernance.com

Colophon

Digital Information Governance (DIG®): The Standard for Defensible AI-Influenced Decisions in Energy, v1.0. Written by Matthew Bertram and published by ModalPoint, Houston, Texas, 2026. Set in IBM Plex Serif, IBM Plex Sans, and IBM Plex Mono. Field vignettes are illustrative composites. Framework changelog and errata are maintained at digitalinformationgovernance.com/changelog.

ABOUT THE AUTHOR

Matthew Bertram

Matthew Bertram is the creator of the Digital Information Governance framework and President of ModalPoint, an AI governance advisory practice serving energy operators from Houston, Texas. His work sits at the intersection of decision quality, AI systems, and trust: two decades across energy-native ventures, including co-founding an oil and gas executive search firm acquired in 2015, and the leadership of EWR Digital, a Houston firm operating since 1999. He writes and speaks on decision integrity for energy audiences, including landman, legal, and operator associations across Texas.

PUTTING DIG TO WORK

ModalPoint operates the DIG framework as a practice. Engagements follow the sequence this paper describes, smallest effective step first.

---

01 **AI Visibility Audit**

Your representation baseline: how AI systems currently describe your company to the counterparties deciding about it.

02 **TRAIGA Readiness Assessment**

Your exposure under Texas HB 149, mapped to the decision surface, with the evidence gaps stated plainly.

03 **Full DIG® Assessment**

Maturity scored by pillar against live decisions, and the ninety-day instrumentation plan from Section 15.

MODALPOINT.COM DIGITALINFORMATIONGOVERNANCE.COM MATTHEWBERTRAM.COM