Governed Decision Intelligence

The Decision Architecture for Governed AI

Open Specification | v3.0

By Mark Julius Banasihan

March 2026

"Every governance framework tells organizations what to govern. Every platform gives tools to manage governance processes. No published specification defines the governed decision record itself."

The GDI Thesis

Table of Contents

Section 01

The Governance Implementation Gap

The decision layer of AI governance is architecturally absent. Existing frameworks govern organizations (roles, policies, risk registers) and systems (model cards, bias audits, data lineage). The individual AI-informed decision, where a specific prediction is evaluated, contextualized, and acted upon by a specific human, remains ungoverned infrastructure.

A consistent pattern emerges across regulated industries adopting AI. Organizations invest in governance frameworks, hire compliance teams, map regulatory requirements, and produce policy documents. Then a system makes a prediction. An analyst acts on it. A decision is made. And the governance layer has no structured record of what happened at the decision point.

This gap is structural, and it has specific intellectual and organizational roots that explain why it persists.

The Decision Science Foundation

The theoretical foundations trace an unbroken line from Herbert Simon's bounded rationality (1955) through modern computational rationality. Simon established that agents cannot evaluate all alternatives due to cognitive limits, time constraints, and incomplete knowledge. They satisfice rather than optimize. His distinction between procedural rationality (how decisions are made) and substantive rationality (whether outcomes are optimal) maps directly to the governance question: should we govern the decision process or the decision outcome? GDI's answer is both, captured in a single artifact.

Gary Klein's Recognition-Primed Decision model, developed from fieldwork with firefighters and military commanders, showed that experts under pressure do not weigh options analytically. They recognize situations as familiar and simulate the first workable solution. This predicts that AI systems functioning as decision aids should amplify human pattern recognition rather than replace expert judgment. Daniel Kahneman's dual-process theory (System 1 fast inference, System 2 deliberative reasoning) maps directly to AI architecture: neural networks provide System-1-like pattern evaluation while search procedures provide System-2-like deliberation. Decisions produced by fast inference and those produced by deliberative chains may require fundamentally different governance structures.

The convergence point is computational rationality, synthesized by Gershman, Horvitz, and Tenenbaum (Science, 2015): identifying decisions with highest expected utility while accounting for computational costs. Thomas Icard's Resource Rationality (2023) formalizes this further, proposing that decision quality must trade off against the costs of computation. This framework treats bounded rationality as an optimization problem rather than a deficiency, and it is the intellectual foundation for GDI's confidence threshold model.

Three Altitudes of AI Governance

Altitude	What It Governs	Who Has Built This	Status
Organizational	Roles, policies, risk registers, board oversight, maturity assessment	NIST AI RMF, ISO 42001, AIGN OS, EU AI Act, OECD Principles	Mature
System	Model cards, bias audits, data lineage, model lifecycle, observability	OneTrust, IBM OpenPages, Credo AI, Holistic AI, model registries	Emerging
Decision	Individual AI-informed decisions with provenance, accountability, and audit trail	No published specification exists	ABSENT

Five Structural Reasons the Decision Layer Is Missing

1

IT Governance Inheritance

Frameworks evolved from IT and data governance traditions that focus on systems, data quality, and model lifecycles. Individual decisions were never the governed object.

2

Static-Dynamic Mismatch

Traditional compliance operates on quarterly or annual audit cycles. AI decisions occur at millisecond speed. Per-decision governance seemed impractical until policy-as-code made inline enforcement feasible.

3

Wrong Abstraction Layer

The EU AI Act classifies AI systems by risk level, pushing governance to the system layer. Individual decisions within a high-risk system receive no independent governance treatment.

4

Policy-Proof Gap

An AI policy that says 'humans remain accountable' is meaningless without an operational answer to: who had the right to delegate this action to the system, and can you demonstrate what happened?

5

Governance as Performance

Organizations perform oversight symbolically while real decisions occur outside formal channels. Only 24% of organizations have operationalized AI governance across their enterprise, even where policies exist.

Section 02

Where GDI Sits

The Infrastructure Reality

The accelerated computing era is converging two data ecosystems that have historically been governed separately. Structured data, the $120B+ ecosystem of SQL engines, data frames, and enterprise databases, represents the ground truth of business. Every financial transaction, patient record, supply chain event, and customer interaction lives in structured form. Unstructured data, growing at hundreds of zettabytes per year, represents the context of AI: PDFs, videos, images, natural language, sensor streams. Until recently, this data was computationally inaccessible. Multi-modal perception and vector search have changed that.

GPU-accelerated libraries now process both structured data frames and vector stores at speeds that compress the time between 'data exists' and 'decision happens' from hours to milliseconds. These libraries are being integrated across every major cloud platform, every major enterprise database, and every major open source processing engine. The result is a unified data fabric where AI systems consume both structured ground truth and unstructured context simultaneously to produce predictions, recommendations, and autonomous actions.

This convergence creates a governance requirement that no existing framework addresses: decisions drawing from both structured and unstructured sources need provenance records that span both data types. The governed decision record must capture what structured data informed the decision (which tables, which queries, which freshness) and what unstructured context was retrieved (which documents, which embeddings, which confidence).

The Five-Layer Architecture

GOVERNANCE FRAMEWORKS

NIST AI RMF | ISO 42001 | EU AI Act | AIGN OS | ARAF

GDI: GOVERNED DECISION LAYER

Decision Records | Confidence Thresholds | Gate Taxonomy | Accountability Chains | Audit Trail

AI INFERENCE + AGENTS

Predictions | Recommendations | Classifications | Generations | Autonomous Actions

DATA INFRASTRUCTURE

Structured (SQL, Spark, Data Frames) | Unstructured (Vectors, RAG, Knowledge Graphs)

ACCELERATED COMPUTE

GPU Clusters | Cloud | Edge | On-Prem | Confidential Computing | Inference Engines

GDI sits between the data/inference layer and the governance framework layer. It captures the transition point where AI-processed data becomes a recommendation, and a recommendation becomes a decision with consequences. Below GDI, data platforms and AI models produce outputs. Above GDI, governance frameworks define what should be governed and how maturity is assessed. GDI is the implementation layer that connects 'we have a policy' to 'we can prove what happened.'

1

What Service Does GDI Provide?

Decision-level accountability infrastructure. GDI makes the governed decision a first-class artifact with provenance, model metadata, human oversight records, confidence assessment, and audit-ready structure. It turns the question 'who is responsible for this outcome?' from a political question into a technical one with a verifiable answer.

2

Why Does GDI Matter Now?

Five converging pressures: (1) accelerated computing compresses decision cycles to machine speed; (2) agentic AI spawns sub-decisions no human directly authorized; (3) structured and unstructured data convergence means decisions draw from data types never governed together; (4) regulations mandate individual decision explanations organizations cannot currently produce; (5) automation bias degrades human oversight quality unless the oversight loop is architecturally enforced.

3

Where Does GDI Land?

In organizations that have governance frameworks and are discovering the implementation gap. They have policies, roles, and risk registers. They cannot produce a governed record of any individual AI-informed decision. GDI fills the space between governance policy and governance proof.

Section 03

The GDI Core Specification

The GDI Core Specification defines a domain-agnostic architecture for governing AI-informed decisions. It consists of five components that together form a complete decision governance layer: the Governed Decision Record, the Confidence Threshold Model, the Gate Taxonomy, the AI Assistance Disclosure Protocol, and the CI/CD Governance Integration pattern. Each component is designed to operate independently or as part of the integrated specification.

The Governed Decision Record (GDR)

The fundamental unit of GDI is the Governed Decision Record. A GDR is a schema-validated artifact that captures a single consequential decision at the point where an AI prediction becomes a human-authorized action. It is the artifact that makes ARAF's reconstructability principle mechanically achievable and the EU AI Act's Article 14 human oversight requirement demonstrably satisfied.

Element	Description	Enforcement
Decision Question	The specific question being decided, with deadline and scope	Schema-required. Validates non-empty.
AI Prediction	Model output, recommendation, or analysis that informed the decision	Structured disclosure: model identity, confidence, limitations
Options Considered	At least two options evaluated, including inaction	Schema-enforced minimum of two options
Evidence Base	Supporting evidence with completeness classification	Each item: complete, partial, or placeholder
Risk Posture	Accepted risks, residual risks, tolerance rationale	Residual risk field explicitly required
Decision Outcome	Go, conditional go, defer, no-go, or escalate	Enumerated outcome types with conditional logic
Accountability Chain	Named decision owner, reviewers, approvers	Human identity required. No delegation to systems.
Data Provenance	Structured and unstructured sources that informed the decision	Source type, freshness, governance status per source
Escalation Record	Triggers, paths, and outcomes of any escalation	Triggered by confidence thresholds or policy boundaries
Downstream Propagation	What must change if this decision changes	Dependency map with named owners

The GDR schema is machine-readable and validatable. Every field has a defined type, cardinality, and validation rule. This means governance checks can be automated: a CI/CD pipeline can verify that every decision record in the system conforms to the schema, that required fields are populated, that accountability chains are complete, and that confidence thresholds have been properly evaluated.

Relationship to ARAF's Decision Supply Chain

ARAF's foundational principle states that institutions can rely on autonomous systems only to the extent that decisions can be reconstructed from contemporaneous governance records. ARAF specifies that evidence must be chain-complete across the entire Decision Supply Chain: data inputs, model processing, human review, and execution infrastructure. The GDR is the artifact that satisfies this requirement at each link in the chain. Where ARAF asks 'can you reconstruct this decision?', GDI answers 'here is the schema-validated record that makes reconstruction mechanically possible.' Organizations adopting ARAF's assessment framework can use GDI's GDR as the evidence artifact that demonstrates reconstructability.

Section 04

The Governed Decision Record

The fundamental unit of GDI is the Governed Decision Record. A GDR is a schema-validated artifact that captures a single consequential decision at the point where an AI prediction becomes a human-authorized action. It is the artifact that makes ARAF's reconstructability principle mechanically achievable and the EU AI Act's Article 14 human oversight requirement demonstrably satisfied.

Element	Description	Enforcement
Decision Question	The specific question being decided, with deadline and scope	Schema-required. Validates non-empty.
AI Prediction	Model output, recommendation, or analysis that informed the decision	Structured disclosure: model identity, confidence, limitations
Options Considered	At least two options evaluated, including inaction	Schema-enforced minimum of two options
Evidence Base	Supporting evidence with completeness classification	Each item: complete, partial, or placeholder
Risk Posture	Accepted risks, residual risks, tolerance rationale	Residual risk field explicitly required
Decision Outcome	Go, conditional go, defer, no-go, or escalate	Enumerated outcome types with conditional logic
Accountability Chain	Named decision owner, reviewers, approvers	Human identity required. No delegation to systems.
Data Provenance	Structured and unstructured sources that informed the decision	Source type, freshness, governance status per source
Escalation Record	Triggers, paths, and outcomes of any escalation	Triggered by confidence thresholds or policy boundaries
Downstream Propagation	What must change if this decision changes	Dependency map with named owners

The GDR schema is machine-readable and validatable. Every field has a defined type, cardinality, and validation rule. This means governance checks can be automated: a CI/CD pipeline can verify that every decision record in the system conforms to the schema, that required fields are populated, that accountability chains are complete, and that confidence thresholds have been properly evaluated.

Relationship to ARAF's Decision Supply Chain

ARAF's foundational principle states that institutions can rely on autonomous systems only to the extent that decisions can be reconstructed from contemporaneous governance records. ARAF specifies that evidence must be chain-complete across the entire Decision Supply Chain: data inputs, model processing, human review, and execution infrastructure. The GDR is the artifact that satisfies this requirement at each link in the chain. Where ARAF asks 'can you reconstruct this decision?', GDI answers 'here is the schema-validated record that makes reconstruction mechanically possible.' Organizations adopting ARAF's assessment framework can use GDI's GDR as the evidence artifact that demonstrates reconstructability.

Section 05

Confidence Thresholds and Gate Taxonomy

The Confidence Threshold Model

GDI defines three confidence zones that determine how a decision flows through governance. These zones address the core tension between governance overhead and operational speed by applying proportional governance: low-risk decisions receive lightweight oversight, while high-risk or uncertain decisions receive intensive human review. This approach draws from IEC 62304's risk-based classification in medical device software and the FDA's Predetermined Change Control Plans, which govern the conditions under which AI systems can evolve rather than validating each output individually.

Zone	Behavior	Governance Requirement
GREEN (High Confidence)	AI prediction aligns with established patterns and policy boundaries. Decision proceeds within pre-authorized parameters.	Standard GDR logged. Human review documented, may be asynchronous. Periodic batch audit.
AMBER (Moderate Confidence)	AI prediction falls in uncertain range, involves novel conditions, or approaches policy boundaries.	Synchronous human review required before action. Additional evidence gathering triggered. Full GDR with justification.
RED (Low / Out of Bounds)	AI prediction conflicts with policy, exceeds trained domain, or addresses conditions outside governed parameters.	Mandatory escalation. No action without explicit senior human authorization. Complete GDR with escalation record.

Threshold boundaries are defined per vertical module. The core spec requires that thresholds exist and are enforced; the specific values are domain-dependent. A pharmaceutical company's amber zone for an IND submission decision differs fundamentally from a telecom network's amber zone for an automated routing decision. What remains constant is the architectural pattern: every AI-informed decision is assessed against defined thresholds, and the governance response is proportional to the assessed risk.

This model directly counters two well-documented failure modes. Automation bias (Parasuraman and Manzey, 2010) causes humans to over-rely on AI recommendations, particularly in high-throughput environments where decision fatigue degrades oversight quality. The threshold model forces the system to classify its own confidence and escalate uncertain decisions before they reach a fatigued human reviewer. Algorithm aversion (Dietvorst, Simmons, and Massey, 2015) causes humans to abandon algorithms after observing even small errors. The threshold model preserves human agency by making the AI's uncertainty visible, allowing humans to calibrate their trust appropriately rather than oscillating between blind trust and blanket rejection.

Gate Taxonomy

GDI recognizes five decision outcomes. This taxonomy is informed by phase-gate decision methodology from regulated industries, where decisions are discrete, documented events with defined governance consequences.

Outcome	Definition	Governance Implication
Go	Decision proceeds with full confidence and complete evidence	Standard GDR logged. Downstream propagation tracked.
Conditional Go	Decision proceeds with explicit conditions, owned follow-ups, and verification checkpoints	Conditions schema-validated. Each condition has named owner and deadline.
Defer	Decision paused due to insufficient evidence or unresolved dependencies	Re-entry criteria defined. Evidence gaps documented. Timeline set.
No-Go	Decision explicitly rejected with documented rationale	Re-entry logic captured. Stopping treated as risk reduction.
Escalate	Decision exceeds current authority scope or policy boundaries	Escalation path, receiving authority, and outcome all logged.

A critical design principle: defer and no-go are first-class outcomes with the same governance rigor as go decisions. In many organizations, the decision to stop or pause is treated as an absence of decision. In GDI, it is a governed artifact. Aviation's Crew Resource Management philosophy informs this design: the decision to go around on a landing approach receives the same procedural treatment as the decision to land. Stopping early is risk reduction, and risk reduction deserves a governed record.

Section 06

CI/CD Governance Integration

GDI decision records are schema-validated artifacts. This means governance checks can be embedded directly in deployment pipelines, transforming governance from periodic manual review into continuous automated enforcement. This pattern represents the shift from governance-as-oversight to governance-as-infrastructure.

The Policy-as-Code Paradigm: Financial institutions using policy-as-code have achieved reductions of 90% or more in compliance violations from human error. The principle is that governance rules translated into machine-executable controls enforce at runtime, making violations architecturally impossible rather than procedurally discouraged.

What CI/CD Governance Validates

Validation	What It Checks	Failure Response
Schema Conformance	Every GDR conforms to the defined schema with all required fields	Build fails. Decision record cannot be committed.
Accountability Completeness	Named human owner, reviewers, and approvers present	Build fails. Anonymous decisions are architecturally blocked.
Confidence Assessment	Threshold evaluation recorded with zone classification	Warning or failure depending on risk tier.
Evidence Linkage	Evidence items linked with completeness classification	Warning. Placeholder evidence flagged for follow-up.
AI Disclosure	When AI was used, structured disclosure is complete	Build fails. Undisclosed AI usage is blocked.
Propagation Map	Downstream dependencies identified and owned	Warning. Unowned dependencies flagged.

This pattern was first implemented in the RGDS reference repository (github.com/mj3b/rgds), where GitHub Actions validate every decision log against the schema on every push. The 170+ commit history demonstrates that schema-validated decision governance is operationally feasible in a CI/CD environment. The same pattern extends to any pipeline: Jenkins, GitLab CI, Azure DevOps, or custom orchestration.

Governance at Machine Speed

When AI inference operates at sub-millisecond timescales and agentic systems spawn autonomous sub-decisions, post-hoc governance becomes physically impossible. The CI/CD integration pattern addresses this by embedding governance in the pipeline itself. Decision records are produced inline with the decision event, validated automatically, and logged to an immutable audit store. Human review is triggered by threshold violations, not by calendar cycles. This creates a governance architecture that scales with the system rather than bottlenecking it.

For agentic AI systems, where agents spawn sub-agents and interact with external tools autonomously, the CI/CD pattern extends to runtime governance. Each consequential action by an agent produces a GDR. The agent's governance covenant (what it is authorized to decide, what must be escalated) is enforced as a runtime constraint. Chain-of-thought logging across inter-agent communications provides the decision provenance that ARAF's reconstructability principle requires.

Section 07

Why This Is Hard

Decision Science, Failure Modes, and Honest Counterarguments

Building a decision architecture for AI governance is genuinely difficult. The challenges are technical, philosophical, organizational, and economic. Any specification that does not acknowledge these challenges lacks the intellectual honesty to earn adoption. This section maps the failure landscape.

Automation Bias and Decision Fatigue

Automation bias, the tendency to over-rely on automated recommendations, produces two documented error types: omission errors (missing events the AI did not flag) and commission errors (accepting incorrect AI outputs even when contradictory evidence is available). Parasuraman and Manzey's 2010 integrative review found these are manifestations of the same attentional mechanisms. In clinical settings, physicians overrode their own correct decisions in 6% of cases in favor of erroneous decision support advice (Goddard, Roudsari, and Wyatt, 2012). In high-throughput environments, decision fatigue compounds the problem: when you review hundreds of identical transactions and they are all correct, you stop looking carefully. Oversight becomes a rubber stamp.

GDI's confidence threshold model is an architectural countermeasure. By forcing the system to classify its own confidence and escalate uncertain decisions, the model ensures that human attention is directed to the decisions that need it most, rather than spread uniformly across all decisions.

The Responsibility Gap

Andreas Matthias's foundational 2004 paper framed the problem: learning automata create situations where manufacturers and operators are in principle unable to predict future machine behavior, making traditional responsibility ascription impossible. Santoni de Sio and van den Hoven (2018) proposed that meaningful human control requires two conditions: a tracking condition (the system responds to relevant human moral reasons) and a tracing condition (behavior is traceable to proper human understanding). GDI's accountability chain, which requires named humans at every decision point and captures what each person reviewed, overrode, and approved, is the architectural mechanism that satisfies the tracing condition.

The Strongest Counterarguments

1

Governance Overhead Is Economically Significant

The US economy loses approximately $3 trillion annually to excessive bureaucracy. Enterprise system implementations fail at rates of 50-75%. Decision management systems have historically suffered from over-engineering, organizational resistance, and maintenance burden. GDI addresses this through risk-tiered governance: low-risk decisions receive lightweight, automated oversight; full GDR rigor is reserved for consequential decisions where the cost of ungoverned failure exceeds the cost of governance.

2

Decisions May Be Emergent, Not Discrete

The philosophical objection: if decisions are emergent properties of complex systems, they cannot be cleanly captured as schema-validatable events. Deep learning systems exhibit behaviors not explicitly programmed. Philip Anderson argued that reductionism does not imply reconstructionism. GDI acknowledges this by governing the conditions under which decisions emerge (the governed process) alongside the decision artifact itself (the governed record). The schema captures what was decided, by whom, on what basis, with what confidence. It does not claim to capture the full emergent complexity of the decision process.

3

Speed and Governance Are in Genuine Tension

Real-time AI systems respond in milliseconds. A single fraud model evaluates millions of transactions per hour. Schema-validating each decision introduces latency. GDI resolves this through tiered governance: high-frequency, low-risk decisions produce lightweight log records validated asynchronously; consequential decisions produce full GDRs validated synchronously. The CI/CD pattern ensures validation is automated and operates at pipeline speed rather than human speed.

Section 08

Lessons from Regulated Industries

Decision-level governance is operationally achievable. Several regulated industries have built it successfully, and their patterns directly inform GDI's architecture.

Industry	Governance Pattern	What GDI Learns
Aviation	Crew Resource Management governs decision process: situational awareness, communication structures, challenge authority. Flight data recorders provide continuous behavioral monitoring.	Decision quality depends on process quality. Governance should govern decision conditions, not just outcomes. CRM's go-around philosophy informs GDI's treatment of defer/no-go as first-class outcomes.
Nuclear	Probabilistic Risk Assessment (Kaplan-Garrick triplets: scenarios x likelihoods x consequences). Defense-in-depth philosophy: multiple independent barriers.	Risk-tiered governance with redundant barriers. No single point of governance failure. GDI's confidence zones and escalation paths implement defense-in-depth at the decision level.
Medical Devices	FDA Total Product Lifecycle approach. Predetermined Change Control Plans require manufacturers to pre-specify what may change and how changes are managed. IEC 62304 risk classification.	Govern the evolution of AI decisions, not just frozen snapshots. Risk-based classification determines governance rigor. GDI's vertical modules adapt governance intensity to domain risk.
Financial Services	SR 11-7 (Federal Reserve/OCC) provides mature model risk management: development, validation, and use. ECOA requires specific reasons for adverse decisions regardless of technology.	Model governance is necessary but insufficient. Decision governance must extend beyond model validation to capture how model outputs become consequential actions. Explainability at the decision level.

The Oxford Martin School (Simpson, Ortega, and Trager, 2025) extracted eight lessons from aviation and nuclear governance for AI. Their central finding: well-defined, checkable components based on rigorous theories of composition produce governable systems. GDI's schema-validated decision records are precisely this: well-defined, checkable components that compose into auditable decision chains.

Section 09

Vertical Module: Biopharma

Reference Implementation

The biopharma vertical module was developed over approximately seven weeks as the RGDS (Regulated Gate Decision Support) framework, focused on IND and BLA submission workflows. It is the proof-of-concept that demonstrates GDI is buildable, validatable, and auditable as a working system.

Regulatory Context

Biopharma operates under some of the most stringent governance requirements in any industry. FDA 21 CFR Part 11 governs electronic records and signatures. ALCOA+ principles (Attributable, Legible, Contemporaneous, Original, Accurate, plus Complete, Consistent, Enduring, Available) define data integrity expectations. The GDR maps directly to ALCOA+: every decision is attributable to a named owner, recorded contemporaneously, maintained as an original schema-validated artifact, and verifiable for accuracy through linked evidence.

Domain-Specific Decision Types

Decision Type	GDI Mechanism	Why It Matters
Author-at-Risk Drafting	Evidence completeness: complete / partial / placeholder. Risk posture with explicit residual risk.	Work proceeds on incomplete evidence with explicit, governed risk acknowledgment.
Reviewer Triage	Review plan with named reviewers, routing rationale, and scope assignments.	Routing of review responsibilities is itself a governed decision.
Regulatory Interaction	Decision category + regulatory context fields. Pre-IND / FDA interaction as first-class decision.	Decision to engage a regulator carries strategic implications requiring governance.
Publishing Lock Point	Publishing plan with lock events, verification checkpoints, and named authorities.	Content freeze for submission is a governed event with explicit accountability.

Implementation Evidence

github.com/mj3b/rgds contains: 170+ commits across 7 releases (v1.0 through v2.0.0) | 6 canonical decision examples demonstrating each outcome type (go, conditional go, no-go, defer, regulatory interaction, AI-assisted) | Schema-enforced decision log with CI/CD validation via GitHub Actions | Evaluation rubrics for decision quality assessment | Separated AI governance covenant repository (github.com/mj3b/rgds-ai-governance) | Apache 2.0 licensed, publicly auditable.

Section 09 (continued)

Why Decision-Centric Governance Is Different

Biopharma organizations already operate enterprise platforms for quality management, document control, and regulatory submissions. These platforms are mature, validated, and deeply integrated into operational workflows. The question practitioners will ask is: why does decision governance require a separate architectural layer? The answer is visible in the design orientation of every major platform.

Platform Orientation Comparison

Platform Category	Design Orientation	How Decisions Are Captured
Document Management (e.g., Vault-class systems)	Document-centric. Documents are primary objects. Workflows route documents through review and approval.	No dedicated decision log. Decisions are captured via document approval metadata. Rationale lives in free-text fields or attached memos.
Quality Management (e.g., QMS platforms)	Process and document-centric. Quality events trigger workflows. Documents are governed through lifecycle states.	Decisions captured within workflows but siloed by process. No unified cross-process decision view.
EHS / Risk Platforms (e.g., compliance suites)	Record-centric. Incidents, audits, and risk assessments are structured records.	No global decision repository. Regulatory decisions tracked in audit modules but not linked to operational gate decisions.
GDI (RGDS Implementation)	Decision-centric. Each gate decision is a primary governed object. Documents and evidence support decisions, not the other way around.	Structured decision records with fields for options, rationale, risk context, evidence completeness, and named approvers. All decisions in one queryable register.

The design orientation difference is foundational. In a document-centric system, the question 'show me all conditional-go decisions made in Q3 with outstanding CMC risks' cannot be answered from the system of record. The information exists, scattered across document approval histories, meeting minutes, and email threads. In a decision-centric system, that query resolves directly because every decision is a structured, indexed, schema-validated record.

Three Dimensions Where the Difference Is Consequential

1

Handling Incomplete Data

IND submissions routinely involve author-at-risk drafting where work proceeds on incomplete evidence. Document-centric platforms handle this through placeholder document statuses or workflow pauses with attached notes. GDI requires explicit evidence completeness classification (complete / partial / placeholder) at the field level, with conditions that trigger follow-up verification. Incomplete data does not get papered over with a workflow status change. It gets governed with named ownership and a verification deadline.

2

Decision Reconstructability Across Modules

A typical IND submission spans multiple CTD modules (quality, nonclinical, clinical) with cross-module dependencies. When a nonclinical finding triggers a CMC specification change, the decision chain must be traceable: who decided to change the spec, based on what evidence, with what risk acceptance, and who approved. Document-centric systems track the document change. GDI tracks the decision that caused the document change, with full provenance linking the nonclinical signal to the CMC action.

3

AI Governance in Regulated Workflows

As AI enters biopharma (document summarization, cross-module consistency checking, regulatory intelligence, gap analysis), the question of how AI assistance is governed becomes critical for 21 CFR Part 11 compliance. Most enterprise platforms either lack AI features or are introducing them with general explainability commitments. GDI has schema-enforced AI disclosure fields: tool identity, purpose, human review tiers, override logs, and confidence assessment. The AI assistance is governed at the same level of rigor as the decision itself. When an auditor asks 'was AI used in this submission decision?', the answer is in the decision record, structured and queryable.

The Complementarity Principle: GDI does not replace enterprise document management, quality management, or regulatory submission platforms. It provides the decision governance layer that these platforms structurally lack. Organizations continue using their validated platforms for document control and quality workflows. GDI captures the governed decisions that drive those workflows, creating the audit trail that connects 'what was decided' to 'what was done.'

Section 10

Vertical Expansion Map

GDI Core is domain-agnostic. Each vertical module adapts three domain-specific dimensions while preserving the core GDR structure, confidence threshold model, gate taxonomy, and CI/CD integration pattern. The infrastructure acceleration patterns emerging across industries create governance requirements in each vertical that existing frameworks cannot address at the decision level.

Vertical	Decision Taxonomy	Regulatory Mapping	Risk Thresholds
Financial Services ($35T+ industry)	Model-driven trading, credit decisioning, risk exposure, sanctions screening, fraud adjudication	SEC/FINRA, Basel III/IV, MiFID II, SR 11-7, ECOA adverse action requirements	Position limits, exposure caps, confidence floors for automated execution
Telecom ($2T infrastructure)	Network routing, spectrum allocation, edge AI deployment, predictive maintenance, base station intelligence	FCC, ETSI, 3GPP, data sovereignty, critical infrastructure protection	Service level thresholds, latency boundaries, coverage minimums for autonomous action
Supply Chain ($35T retail/CPG)	Sourcing adjustments, routing optimization, inventory rebalancing, supplier qualification, demand forecasting	Cross-jurisdictional trade compliance, customs, ESG reporting, FCPA	Cost variance limits, delivery time boundaries, quality thresholds for automated reorder
Healthcare (drug discovery + clinical)	Diagnostic assistance, treatment recommendation, resource allocation, clinical trial design, drug interaction	HIPAA, FDA SaMD, clinical decision support guidance, 21 CFR Part 11	Sensitivity/specificity floors, mandatory human review for all patient-affecting decisions
Robotics / Industrial ($50T manufacturing)	Motion planning, quality inspection, predictive maintenance, safety boundary enforcement	ISO 13849, IEC 62443, Machinery Directive, OSHA	Safety envelope boundaries, force limits, mandatory stop conditions

Each vertical module inherits the complete GDI Core Specification. The module adds domain-specific decision types (extending the GDR schema), regulatory mappings (connecting GDR fields to specific compliance requirements), and threshold definitions (calibrating green/amber/red zones to domain risk tolerances). This architecture allows organizations to adopt GDI Core and activate vertical modules as their domain requires.

Section 11

Framework Compatibility Mapping

GDI does not compete with existing governance frameworks. It completes them. The following mappings show how GDI's Governed Decision Records satisfy specific requirements of the major frameworks in production today. Organizations can adopt GDI alongside any combination of governance frameworks without conflict.

Framework	Relevant Requirement	GDI Implementation
NIST AI RMF	GOVERN function: establish accountability structures. MAP function: contextualize AI systems.	GDR accountability chain satisfies GOVERN. Decision context, evidence base, and data provenance satisfy MAP.
ISO/IEC 42001	Risk management for AI systems. Documented evidence of governance processes. Control A.6.2.8 event logs.	GDR risk posture and residual risk fields. Schema validation provides process evidence. GDRs are the event logs.
EU AI Act	Art. 12: automatic event-logging. Art. 14: human oversight. Art. 86: right to explanation of individual decisions.	GDRs satisfy Art. 12 logging. Accountability chains satisfy Art. 14. Decision rationale satisfies Art. 86.
AIGN OS	Layer 4 (Governance Implementation): translate duties into workflows. Layer 1 (Trust Infrastructure): audit-ready evidence.	GDI provides decision-level implementation for Layer 4. Schema-validated GDRs provide Layer 1 evidence artifacts.
ARAF v3.0	Reconstructability principle: decisions must be reconstructable from contemporaneous governance records. Chain-complete evidence across the Decision Supply Chain.	GDRs are the contemporaneous governance records ARAF requires. GDR provenance fields produce chain-complete evidence.
OECD Principles	Transparency, accountability, and human-centered values in AI systems.	AI disclosure protocol provides transparency. Accountability chain provides named ownership. Thresholds enforce human review.

The Compatibility Principle: GDI is the decision-architecture layer that sits inside any governance framework. It is the implementation artifact that turns framework requirements into governed system behavior. Organizations adopting GDI do not need to choose between frameworks. They need to implement the decision layer that makes their chosen framework enforceable.

Section 12

The Structured and Unstructured Data Convergence

Data governance and AI governance are converging because the data infrastructure feeding AI systems is converging. The implications for decision governance are profound and largely unaddressed.

Two Ecosystems Becoming One

The structured data ecosystem processes data frames: SQL queries, Spark jobs, pandas operations, enterprise data warehouse workloads. This is the ground truth of business. Every enterprise platform, from cloud-native data warehouses to on-premise proprietary engines, processes structured data as its core function. Simultaneously, the unstructured data ecosystem processes the 90% of data generated each year that historically could not be queried: PDFs, videos, natural language, sensor streams, images. Vector databases and multi-modal AI have made this data computationally accessible for the first time.

GPU-accelerated libraries are now integrating both ecosystems through a common acceleration layer. Accelerated data frame processing and accelerated vector search operate on the same compute infrastructure, are integrated into the same cloud platforms, and feed the same AI inference pipelines. The result: a single AI-informed decision may draw simultaneously from a structured SQL query (customer transaction history), a vector search result (similar past cases retrieved from unstructured documents), and a real-time sensor stream (current environmental conditions).

What This Means for Decision Governance

When a decision draws from both structured and unstructured sources, its provenance record must span both. The GDR's data provenance fields capture: which structured sources were queried (table, query, freshness, governance status), which unstructured sources were retrieved (document, embedding model, retrieval confidence, relevance score), and how the AI system synthesized both into the prediction that informed the decision. No existing governance framework specifies this dual-provenance requirement. Data governance standards address data quality and lineage for structured data. Vector database governance is in its infancy. GDI bridges both by treating data provenance as a required field in the governed decision record, ensuring that the convergence of data types is governed at the point where it matters most: the decision.

Section 13

Design Principles

AI generates predictions. Humans exercise judgment.

This is the foundational principle of GDI. No component in the architecture is permitted to silently decide, approve, or accept risk. AI systems produce outputs. Humans own decisions. The architecture enforces this boundary. When agentic AI systems require autonomous action within governed parameters, the governance covenant defines what the agent is authorized to decide. Actions outside the covenant trigger escalation. The human remains the governing authority.

The decision is the primary artifact.

Traditional approaches emphasize inputs: documents, analyses, reports, model cards. GDI treats the decision itself as the governed artifact. Everything else (evidence, analysis, models, data) serves the decision record. This inversion of emphasis reflects a fundamental insight from regulated delivery: programs fail because decisions cannot be reconstructed, not because analyses were missing.

Governance is infrastructure, not oversight.

Governance is embedded in the system architecture through schemas, validation pipelines, and enforcement logic. It operates continuously, automatically, and independently of individual awareness or discipline. The shift from governance-as-oversight to governance-as-infrastructure is the key architectural innovation that makes decision governance feasible at machine speed.

Risk determines rigor.

Drawing from IEC 62304 (medical device software classification), FDA's PCCP model, and nuclear defense-in-depth philosophy: governance overhead must be proportional to decision risk. Low-risk decisions receive lightweight automated governance. High-risk decisions receive full human review with complete GDR documentation. The confidence threshold model enforces this proportionality.

Stopping early is risk reduction.

Defer and no-go outcomes are first-class results in GDI. The architecture treats the decision to pause or stop as a governed outcome with the same structural rigor as a decision to proceed. Aviation's CRM philosophy informs this: go-arounds are not failures. They are governed decisions that prevent catastrophic outcomes.

AI assistance is optional and removable.

Every governed decision must remain defensible if all AI outputs are removed from the record. The governance architecture is durable regardless of whether AI assistance is present. This principle ensures that GDI is valid in environments where AI is restricted, untrusted, or absent, and that AI adoption does not create governance dependencies.

Section 14

Open Source Strategy and Prior Art

GDI is published under the Apache 2.0 license. The specification, schemas, reference implementations, and vertical modules are freely available for adoption, adaptation, and extension.

The open source strategy serves a specific purpose: establishing prior art for decision-level AI governance architecture. The Git commit history across the RGDS and GDI repositories provides date-stamped evidence of the design evolution, from the initial pharma-focused reference implementation through the domain-agnostic specification published here. The timestamps are the prior art.

Repository Structure

Repository	Purpose	Status
mj3b/governed-decision-intelligence	GDI Core Specification, schemas, vertical modules, compatibility mappings	New (canonical home)
mj3b/rgds	Biopharma reference implementation. 170+ commits, 6 canonical examples, CI validation.	v2.0.0 (active)
mj3b/rgds-ai-governance	AI governance covenants. Non-agentic boundaries, human ownership requirements.	Reference governance

Organizations are encouraged to adopt GDI, contribute vertical modules, propose schema extensions, and build tooling on top of the specification. The architecture is designed to be integrated, adapted, and extended rather than adopted wholesale. The goal is not a single authoritative framework but a shared implementation layer that makes every governance framework more enforceable at the decision point.

The AI governance market is building frameworks, platforms, certifications, and readiness assessments. What it has not built is the decision-level architecture that captures the governed decision itself. That gap is where consequential errors enter systems, where accountability becomes diffuse, and where audit trails break down.

GDI addresses this gap with a concrete, schema-validated, CI/CD-integrated architecture that is domain-agnostic at the core and domain-specific at the extension layer. It is compatible with every major governance framework in production. It is open source, freely available, and designed for integration.

The commit history tells the story. The architecture speaks for itself.

Mark Julius Banasihan | github.com/mj3b | Apache 2.0 Licensed | March 2026

Section 15

References and Bibliography

Decision Science and Cognitive Foundations

Anderson, P. W. (1972). More is different: Broken symmetry and the nature of the hierarchical structure of science. Science, 177(4047), 393-396.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2015). Algorithm aversion: People erroneously avoid algorithms after seeing them err. Journal of Experimental Psychology: General, 144(1), 114-126.

Dietvorst, B. J., Simmons, J. P., & Massey, C. (2018). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. Management Science, 64(3), 1155-1170.

Gershman, S. J., Horvitz, E. J., & Tenenbaum, J. B. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. Science, 349(6245), 273-278.

Goddard, K., Roudsari, A., & Wyatt, J. C. (2012). Automation bias: A systematic review of frequency, effect mediators, and mitigators. Journal of the American Medical Informatics Association, 19(1), 121-127.

Hamel, G., & Zanini, M. (2016). Humanocracy: Creating organizations as amazing as the people inside them. Harvard Business Review Press.

Icard, T. (2023). Resource rationality. Cambridge University Press (forthcoming).

Kahneman, D. (2011). Thinking, fast and slow. Farrar, Straus and Giroux.

Klein, G. (1998). Sources of power: How people make decisions. MIT Press.

Parasuraman, R., & Manzey, D. H. (2010). Complacency and bias in human use of automation: An attentional integration. Human Factors, 52(3), 381-410.

Sawyer, R. K. (2005). Social emergence: Societies as complex systems. Cambridge University Press.

Simon, H. A. (1955). A behavioral model of rational choice. The Quarterly Journal of Economics, 69(1), 99-118.

Simon, H. A. (1978). Rational decision-making in business organizations. Nobel Memorial Lecture.

AI Governance, Ethics, and Responsibility

Braun, M., Hummel, P., Beck, S., & Dabrock, P. (2022). Responsible AI: Bridging the gap between ethics and practice. AI and Ethics, 2, 437-445.

Matthias, A. (2004). The responsibility gap: Ascribing responsibility for the actions of learning automata. Ethics and Information Technology, 6(3), 175-183.

Mosier, K. L., & Skitka, L. J. (1996). Human decision makers and automated decision aids: Made for each other? In Parasuraman & Mouloua (Eds.), Automation and human performance (pp. 201-220). Lawrence Erlbaum.

Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. Frontiers in Robotics and AI, 5, Article 15.

Santoni de Sio, F., & Mecacci, G. (2021). Four responsibility gaps with artificial intelligence: Why they matter and how to address them. Philosophy and Technology, 34, 1057-1084.

Simpson, T., Ortega, A., & Trager, R. (2025). Voluntary industry initiatives in frontier AI governance: Lessons from aviation and nuclear power. Oxford Martin School Working Paper.

Regulated Industry Governance Patterns

Federal Reserve Board & Office of the Comptroller of the Currency. (2011). Supervisory guidance on model risk management (SR 11-7/OCC 2011-12).

Helmreich, R. L., Merritt, A. C., & Wilhelm, J. A. (1999). The evolution of Crew Resource Management training in commercial aviation. International Journal of Aviation Psychology, 9(1), 19-32.

International Electrotechnical Commission. (2015). IEC 62304:2006/AMD1:2015, Medical device software - Software life cycle processes.

Kaplan, S., & Garrick, B. J. (1981). On the quantitative definition of risk. Risk Analysis, 1(1), 11-27.

U.S. Food and Drug Administration. (2003). 21 CFR Part 11, Electronic records; electronic signatures. Code of Federal Regulations.

U.S. Food and Drug Administration. (2023). Marketing submission recommendations for a predetermined change control plan for artificial intelligence/machine learning-enabled device software functions. FDA Guidance Document.

World Health Organization. (2021). WHO Technical Report Series No. 1033: Guidance on good data and record management practices (ALCOA+).

Standards, Frameworks, and Regulatory Instruments

ARAF Standard. (2025). Agentic Risk Architecture Framework, v3.0. Institute for Autonomous Governance. https://araf-standard.org

European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). Official Journal of the European Union.

International Organization for Standardization. (2023). ISO/IEC 42001:2023, Information technology - Artificial intelligence - Management system.

National Institute of Standards and Technology. (2023). Artificial Intelligence Risk Management Framework (AI RMF 1.0). NIST AI 100-1.

Nygard, M. (2011). Documenting architecture decisions. Cognitect Blog. https://cognitect.com/blog/2011/11/15/documenting-architecture-decisions

Organisation for Economic Co-operation and Development. (2019). Recommendation of the Council on Artificial Intelligence. OECD/LEGAL/0449.

Singapore Infocomm Media Development Authority (IMDA). (2026). Agentic AI governance framework, v1.0.

Upmann, P. (2025). AIGN OS: The operating system for responsible AI governance. AI Governance Network. https://aign.global

Industry and Market Research

IDC. (2025). Global DataSphere forecast, 2024-2028. International Data Corporation.

NexaStack. (2025). Policy-as-code in financial services: Compliance automation outcomes report.

Responsible AI Governance Network (RAGN). (2025). State of AI governance operationalization: Global survey findings.