



Trisotech

Denis Gagne,
CEO & CTO,
Trisotech

Who Made That Decision?

Governing Decisions in the Age of AI Agents



The Age of AI Agents


The AI Exposure No One Intentionally Underwrote

The **risk profile** changes when AI moves from answering questions to taking actions, to making transactions.



The Rapid Evolution of AI Engineering


From Intern to
Mismanaged Genius



- Each wave adds capability and autonomy faster than it adds governance
- Ungoverned, that is a dangerous mismanaged genius: capable, fast, and accountable to no one

Capability Is Outpacing Absorption

Autonomy does not
equal **trust**



- AI capability is moving faster than deployment
- AI projects do not convert directly into value
- AI Models can do more than organizations can absorb
- Software accelerates faster than companies can reorganize themselves

Frontier Lab Companies are Playing Chicken


Nobody dares brake alone,
so they ask the **government**
to do it for all of them



- Triggered by an OpenAI model that broke its sandbox and hacked Hugging Face's production systems
- CEOs, Chief Scientists, VP of AI from most labs from OpenAI, Anthropic, Google, and Meta signed "Pacing the Frontier"
- Their ask: technical and governance tools to deliberately pace frontier-wide progress.
- We need the option to buy time to address emerging risks, develop security measures, and strengthen oversight

The Question We Can No Longer Answer

A fundamental question is becoming increasingly difficult to answer:
Who made that decision?



- AI assistants are evolving into AI performers, autonomous agents, and AI-driven orchestrators
- AI is now embedded inside operational workflows and decision-making
- When humans and AI act together, accountability blurs
- The question then is: Who is the Ultimate Decision Owner (UDO)? Who is responsible?



Responsible Decision Making


The New Governance Problem

Ensuring **Responsible**
Decision-Making in the age of AI




From AI Model-Centric to Decision-Centric Governance

Risk materializes in the **decision**,
not the model weights



What AI model is this? How accurate is it?
How does it work?

Where does AI participate? With what
authority? Under whose oversight?



Governance must move to where risk actually lives: **the decision**

Seven Words We Keep Confusing


Confusing them
leads to **governance blind spots**

Term	Definition
Explainability	Why the AI produced this output?
Traceability	What happened, step by step?
Provenance	Where data, models, and knowledge came from?
Accountability	Who answers for the outcome?
Auditability	Whether a third party can verify it after the fact?
Governance	Who decides what is allowed
Transparency	What is visible about AI participation?

- Related, overlapping, and routinely used interchangeably
- Each answers a different question about an AI-enabled decision

They Overlap, But They Are Not the Same


You cannot **audit**
what was never made visible



- Transparency is the foundation the others build on
- Traceability and provenance feed auditability
- Explainability supports accountability, but does not create it
- Governance orchestrates all of them
- Governed AI means being able to identify, constrain, observe, reconstruct and account for every consequential AI-enabled decision

Organizations Must Govern Decisions, Not Only AI Models

The natural unit of
governance is the **decision**



- Boards and committees do not approve AI model weight or gradients; they approve decisions
- Risk registers list decisions and their owners
- Regulators ask about outcomes, authority, and controls

Decisions: Where Risk, Authority, and Outcomes Converge

The natural unit of governance is the **decision**



- A decision binds inputs, knowledge, authority, and consequences
- Risk materializes at the decision point, not in the AI model weights
- Authority is granted, exercised, and lost at the decision point
- Decisions are countable, ownable, and auditable



AI Participation Taxonomy



AI Is Not One Thing:


We have learned to solve
1+1 in silicon years ago

Choose the appropriate tool for specific tasks



Natural Language

Designed for human communication




Silicon-Based Computation

Optimized for machine computation

AI Is Not One Thing: The Governed AI Taxonomy

AI Tool


- Invoked by a human, produces an output on demand
- No initiative, no authority, the human uses the result




AI Is Not One Thing: The Governed AI Taxonomy

AI Assistant

- Suggests, drafts, and recommends within a human task
- The human evaluates and remains the decision authority




AI Is Not One Thing: The Governed AI Taxonomy



AI Performer

- Executes a defined task within a process
- Operates inside human-defined boundaries
- Authority is shared: AI acts, humans set and monitor the limits
- Escalation and approval paths become essential


AI Is Not One Thing: The Governed AI Taxonomy



■ AI Agent

- Pursues goals with delegated authority
- Chooses actions, sequences, and tools autonomously
- Operates within explicit constraints (or should)
- Accountability must survive the delegation

AI Is Not One Thing: The Governed AI Taxonomy




■ AI Orchestrator

- Dynamically coordinates processes, decisions, humans, and other agents
- Allocates work and authority at runtime
- Unconstrained orchestration: logic becomes latent, policies implicit, and accountability hard to establish
- Governed orchestration: AI consumes explicit, deterministic operational contracts instead of inventing behavior
- The highest governance burden: authority itself becomes dynamic
- Oversight must be continuous, not periodic

AI Is Not One Thing: The Governed AI Taxonomy

Five distinct modes of **AI participation**
in work and decisions



- Each mode carries different authority
- Each mode demands different governance



Human-AI Agency



How Responsibility Shifts as Autonomy Increases


Responsibility never disappears, it changes form



- As autonomy moves from human-driven to AI-driven
 - The human role shifts from doing, to approving, to supervising, to governing
- As agency shifts toward AI, governance, transparency, and orchestration requirements all increase
 - Every shift in agency must be explicit and recorded

Decision Authority: Human, Shared, Delegated, Dynamic


Authority must always be traceable to a human grant



- Human: AI informs, the human decides
- Shared: AI acts inside human-set boundaries
- Delegated: AI decides within explicit constraints
- Dynamic: authority is allocated at runtime
 - Dynamic authority does not mean unbounded authority. Runtime allocation must remain within a human-approved delegation envelope.

More Autonomy Requires More Governance

Governance is not a flat requirement;
it is a **gradient**




- Transparency requirements rise with autonomy
- Oversight moves from periodic to continuous
- Controls move from implicit habits to explicit, machine-enforced mechanisms

Who Is Really in Charge?

The Human-AI Agency Scale (HAAS)

A scale for reasoning about responsibility, oversight, and decision authority



Level-1 Human-Driven: human decides and acts, AI is absent [or trivial]



Level-2 Human-Assisted: human decides, AI informs



Level-3 Human-AI Collaborative: shared agency, AI proposes or acts with human approval



Level-4 AI-Assisted: AI acts, human supervises and can intervene



Level-5 AI-Driven: AI decides and acts, human sets policy




Governance via Contracts



The Governance Gap

The gap is **architectural**,
not technological



- Most organizations can inventory their AI models
- Few can inventory where AI participates in decisions and workflows
- Model registries are not decision governance

Where Does Governance Live? In the Decision Architecture

If you can **model it**,
you can **govern it**




- Governance needs an address: a place where it can attach
- BPM+ standards make decisions, processes, cases, and knowledge explicit
- Explicit models create explicit governance points

BPM+ refers to a family of open Business Process Management international standards produced by a Standard Development Organization called OMG <https://www.trisotech.com/bpm-plus/>

BPMN and CMMN: Where AI Acts


Authority assignments
is **explicit**



- BPMN: structured processes
 - Tasks show exactly where AI performs work
- CMMN: dynamic cases
 - Discretionary actions AI may take
- Both make AI participation visible at the task level

DMN, SDMN, and KEM: Where AI Decides and Knows

Together: an **explicit** map of decision dependencies




- DMN: decision logic, decision requirements, and who or what decides
- SDMN: the situational data feeding decisions
- KEM: the knowledge and vocabulary decisions depend on

DMN is the Decision Model and notation <https://www.trisotech.com/dmn/>
SDMN is the Shared Data Model and Notation <https://www.trisotech.com/sdmn/>

Explicit Execution Contracts, Not Opaque Behavior

Deviation from the contract
becomes **detectable**



- BPM+ models are executable contracts: inputs, outputs, roles, boundaries
- AI participation becomes a declared capability, not an emergent side effect
- Contracts define what AI may do and what it may not



Transparency



Transparency by Design, Not Compliance After the Fact

Retrofitted transparency is
expensive, partial, and late



Transparency
patch

- Transparency must be built into the decision architecture itself
- Make AI participation explicit and machine-readable from day one
- Transparency becomes a capability, not a report



Transparency
woven

AI participation must never be hidden or implicit.

What should be Clearly Identified?

Toward AI
Transparency

AI Transparency Checklist



**Which parts
were AI-
generated?**



**Which AI
systems
participated?**



**Which
prompts or
context were
used?**



**Whether
humans
reviewed or
approved AI-
generated
content?**



**Which AI
models
influenced
execution
behavior?**



**Whether
outputs are
determinis-
tic, proba-
bilistic, or
externally
generated?**

AI Transparency Tags

Machine-readable governance metadata
attached to assets



- Not a single "AI = true" flag: transparency is categorized and layered
- Declare AI participation: where, in what role, with what authority, under whose oversight
- Attach to processes, decisions, agents, prompts, and knowledge assets
- Queryable: governance questions become simple queries, not investigations

Anatomy of a Transparency Tag

- Provenance: which model, version, vendor, prompt and context references
- Role: Tool, Assistant, Performer, Agent, or Orchestrator
- Authority: HAAS level and the decision authority granted
- Oversight: who supervises, human review state, approval and escalation paths
- Constraints: autonomy scope, boundaries, and policy source
- Classification: confidence, risk category, deterministic or probabilistic behavior




Bringing It Together



Bringing It Together: AI Taxonomy, Human-AI Agency, and Transparency

Three lenses,
one governance model



The Governance Matrix



AI Role	HAAS Level	Decision Authority	Transparency Requirement	Governance Requirement
Tool	Level 2	Human	Basic	Low
Assistant	Level 2/3	Human/Shared	Moderate	Moderate
Performer	Level 3/4	Shared/Delegated	High	High
Agent	Level 4/5	Delegated	Very High	Very High
Orchestrator	Level 5	Dynamic	Maximum	Maximum


Minimal AI Transparency Declaration Example for This Presentation

Declaration example using the **transparency framework** proposed in this presentation

- **AI involved:** True
- **Role:** Tool
- **Provider:** OpenAI
- **Model:** GPT5.5 High
- **Source:** Trisotech BPM+ Governed AI Taxonomy and Positioning
- **Source:** Trisotech AI Transparency Tagging PRD
- **HAAS:** Level2 Human-Assisted
- **Accountability:** Denis Gagne
- **Reviewed:** True
- **Reviewer:** Denis Gagne
- **Approval State:** Approved.

Standards Are Moving in This Direction

BPM+ provides the **architectural substrate** these efforts need



- HL7 AI Transparency Implementation Guide: layered declarations and provenance for AI in healthcare decision support https://build.fhir.org/ig/HL7/aitransparency-ig/en/general_guidance.html
- W3C PROV: entities, activities, and agents as the backbone of provenance modeling <https://www.w3.org/TR/prov-overview/>
- SPDX and SBOM practice: explicit, machine-readable declarations of what software contains <https://spdx.dev/>
- Audit and accountability requirements in emerging AI regulation (EU AI Act, NIST AI RMF) <https://artificialintelligenceact.eu/> <https://www.nist.gov/itl/ai-risk-management-framework>



AI In, On, and As



Trisotech: One Governed Framework

Decision-Centric
Orchestration



**AI In
Orchestration**



**AI On
Orchestration**



**AI As
Orchestration**

- The Tool through Orchestrator AI Taxonomy describes the AI's operational role
- In, On and As describes its relationship to the orchestration architecture

Trisotech: One Governed Framework

AI **In** Orchestration
AI **On** Orchestration
AI **As** Orchestration



AI In Orchestration:

tools, assistants, performers, and agents can act inside governed execution
BPM+ controls invocation, escalation, and accountability

E.g., Summarization, extraction, classification, recommendations, AI task performers



AI On Orchestration:

AI generates, optimizes, and analyzes BPM+ models and runtime behavior
AI identifies improvement recommendations, insight analysis

Generated does not mean trusted: generated assets enter the governed model lifecycle



AI As Orchestration:

AI coordinates work dynamically while consuming explicit deterministic BPM+ operational contracts via MCP (Model Context Protocol)

How Trisotech Implements This?

The result:
deterministic operational contracts
for probabilistic AI systems



- The Trisotech platform operationalizes these practices through BPM+ standards
- BPM+ models are explicit behavioral contracts: what is allowed, what is observable, what is governed, where autonomy is permitted
- DMN governs decisions, BPMN governs orchestration, CMMN governs adaptive work
- SDMN governs shared logical data, KEM governs semantics and enterprise meaning
- AI outputs, recommendations and proposed actions become governed inputs to orchestration and policy enforcement

The Result: Governed Agentic Orchestration Through BPM+ MCP

A **neuro-symbolic architecture**:
probabilistic AI reasoning plus
deterministic BPM+ governance



- Traditional BPMN-centric orchestration is governed but often less adaptive vs pure agentic AI is adaptive but difficult to govern
- Governed Agentic Orchestration is both: adaptive and governable
- BPM+ MCP exposes processes, decisions, cases, and semantics as callable, governed tools for AI agents
- The AI may reason probabilistically, but it does not invent enterprise rules, policies, authority or operational contracts. It consumes them explicitly.
- Decisions stay externalized in DMN, policies stay observable, boundaries stay enforceable

Transparency by Design: The Trisotech Direction

A **neuro-symbolic architecture**:
probabilistic AI reasoning plus
deterministic BPM+ governance



- AI transparency tags embedded in the XML serialization of every modeling asset: BPMN, DMN, CMMN, SDMN, KEM, forms, and tests
- Layered: lightweight declarations plus rich, extensible provenance
- Captures AI identity, role, human review state, confidence, and prompt and model references
- Covers both authoring-time AI and runtime AI participation
- Vendor neutral aligned with the HL7 AI Transparency IG and W3C PROV

Governed, Transparent, Explainable AI Participation

A Practical Adoption Roadmap

Expect surprises:
AI **already** participates in more
places than anyone thinks

What You Can Do Tomorrow




Made with  Napkin

Conclusion: Who Made that Decision?

Transparency by design is a foundational capability, not a compliance exercise

Governed AI means being able to answer that question every time



Any questions?



Denis Gagne,
dgagne@trisotech.com

THANKS!