FIRST SAN FRANCISCO PARTNERS



AI GOVERNANCE FRAMEWORK

Decision Stewardship: Foundations & Operating Model

What Is Decision Stewardship?

Decision stewardship is the ongoing, accountable management of AI-related decisions on behalf of the individuals and functions that bear their consequences. It is not a single approval, and it is not an audit of model outputs on a per interaction basis.

It is the discipline of ensuring that the point at which organizational authority converts into AI's authority to reason, decide, and act is deliberately configured, continuously justified, and demonstrably trustworthy.

The **Decision Steward** secures answers to three categories of decisions across the lifecycle of an AI solution:

1. **QUALIFICATION**
   Is the use case worth pursuing?
2. **DELEGATION**
   What authority should AI hold?
3. **OPERATIONAL**
   Is the system still trustworthy?

1 Qualification

Qualification determines whether a proposed use case has enough demonstrated value relative to its risk to justify investment, before any technical build begins. These are business-facing decisions that define the role AI will play, whether as a tool or as a resource contributing directly to outcomes. Although these are ultimately business decisions, many questions will require consultation and input with data and technical communities. Findings are presented to the Data, Analytics, and AI Leadership Committee for review and approval, then packaged as business requirements for the technical team. Qualification is expected to answer key questions such as:

BUSINESS VALUE / What business value or outcome will the use case produce, and how will it be measured?

RISK TIER / What is the use case's approximate risk tier, prior to the application of any controls?

BASELINE SAFEGUARDS / What baseline technical safeguards (e.g., data-source restrictions, logging and traceability, human-review thresholds) does that risk tier presumptively require, and are those capabilities already in place as the starting requirements for engineering?

DATA READINESS / Does the data and metadata the use case requires exist, and is it acceptable quality, accessible, and authorized for AI use, or does it require new access or governance work before a technical build can begin? (Addressed in partnership with the Data Steward, who identifies and qualifies data readiness.)

GOVERNED FLOW / Does the use case fit within an existing governed flow, or does it require establishing a new one?

info@firstsanfranciscopartners.com
888-499-DATA

2 Delegation

Delegation determines how much control the organization concedes to AI at two distinct points: **1.)** in formulating the outcome, by shaping the options, synthesis, or recommendation on which the decision will rest, and **2.)** in executing that outcome once a decision has been made. These decisions configure the manner in which AI will support, influence, and/or act on the organization, and their answers become functional requirements for the technical team. Delegation is expected to answer key questions such as:

FORMULATION AUTHORITY / How much latitude does AI have to shape the option set, synthesis, or recommendation underlying the decision, versus humans defining the frame and options themselves?

EXECUTION AUTHORITY / Once an outcome is decided, where does AI sit on the human-in/on/over/out-of-loop spectrum for acting on it: recommending, deciding with review, or deciding and acting autonomously?

LIFECYCLE CONTROL / What controls exist, or must be established, to govern changes to either authority level, and what triggers them?

OWNERSHIP / Who is named accountable for the process across both formulation and execution, and what authority do they hold to intervene or override?

PRECONDITIONS & REVOCATION / What must be demonstrably true about the system's risk profile (e.g., it's grounded and traceability scores), before either form of authority is granted, and what signal would warrant its downgrade?

3 Operational

Operational decisions are those AI actually makes, or materially shapes, once deployed, occurring with each interaction in production. This category constitutes the evidentiary layer that validates the two preceding decision categories. The Decision Steward is not responsible for reviewing every individual output; rather, the role is responsible for continuous monitoring of the system, including periodic and formalized testing of outputs against original use-case intent, and for ad hoc support to the user community, fielding questions, resolving issues, and liaising with the technical team as needed. Operational stewardship is expected to answer key questions such as:

EXCEPTION RESOLUTION / When an output enters the exception queue, how should it be resolved, and by whom?

VALUE & DRIFT / Is the system still delivering the value for which it was authorized, or has drift set in?

ATTRIBUTION & OWNERSHIP / Has an attribution gap emerged, an output with no traceable source or no clear owner, and what is the appropriate response?

TRUSTWORTHINESS / Does audit of the outputs provide demonstrable evidence that the system is operating according to the established principles and characteristics of trustworthy AI?

RECONSIDERATION TRIGGERS / Does accumulated operational evidence warrant returning this system to Qualification or Delegation for reconsideration?

ACCOUNTABILITY ACROSS THE AI LIFECYCLE

Decision stewardship, therefore, is not reducible to a single approval, nor to the ongoing review of individual outputs. It is the responsible management of the point at which the organization's authority over a decision converts into AI's authority to reason, decide, and act is deliberately authority to act on it: deliberately configuring that conversion, sustaining its justification for as long as it stands, and monitoring it to demonstrate that the system continues to meet the organization's standards for trustworthy AI.

Why Is Decision Stewardship Needed?

Decision stewardship is necessary, first, because a single model can serve multiple use cases. Stewardship applied only at the model level neither comprehensively mitigates risk nor reliably delivers value: model stewardship demands a technical skill set, and because one model can support numerous use cases, it spans a scope of decision-making responsibility too broad for any one person to manage.

A large language model cleared for a Project Management Office to summarize project strategy, for example, bears little resemblance, in risk, in required domain expertise, or in appropriate controls, to the same model applied to screening vendor contracts.

Decision stewardship is necessary, second, because risk in AI-supported decisions is not solely a property of the outcome; it is a property of the decision process itself. The unique danger of AI-supported decisions lies upstream, in how the problem is framed, rather than downstream, in the outcome produced. Traditional governance, which watches outcomes, therefore systematically overlooks the actual point of failure: unintended consequences can carry an outsized impact when the right questions are not asked before a solution is built, regardless of how flawlessly it is executed. Risk also lives in the decision-making process, not only in its result.

Finally, decision stewardship cannot be a one-time exercise. Unlike data, AI is not static and does not represent a fixed point in time; it is dynamic, and it persists within the organization after deployment. Rather than being acted on, AI acts on other assets. The decision process that governs it must therefore be revisited on a defined cadence, to confirm that the solution continues to solve the problem it was built to address, especially as the business environment surrounding it continues to grow and evolve.

Decision Stewardship & the AI Governance Operating Model

The **Decision Steward** reports to the **Solution Owner**: the business owner of the use case, of the authority-conversion decision, and of the AI solution's continued validation and justification. The Solution Owner's counterpart is the **Product Owner**, who owns the technical instantiation of the use case.



COMMITTEE LEVEL

At the committee level, the Decision Steward reports to the **Decision Stewardship Committee**, and the **AI Steward** reports to the **AI Stewardship Committee**. Each committee maintains a distinct, recurring agenda because each is organized around a different question: “should we” and “how do we”, respectively.

Answering either question typically requires in-depth analysis with domain experts. On the business side, work such as business-process modeling and impact assessment is substantial and specialized enough to warrant a dedicated forum, with technical input brought in only as needed. On the technical side, the AI Stewardship Committee addresses the mechanics of the solution portfolio: system integration and interoperability, componentization, and system updates and versioning, as examples. Each committee is structured to support its own steady-state work; critically, however, formal joint sessions are required for delegation, cross-cutting decisions and escalations, and issue resolution.



WORKING-GROUP LEVEL

At the working-group level, pods form around individual AI solutions. These **AI Solution Pods** do not disband once a solution is deployed; they persist until the solution is fully decommissioned and should be convened periodically to revalidate the solution's risk, value, and performance.

The pod model answers a question most organizations struggle to resolve: who, specifically, is accountable for ensuring the right controls are in place, for checking the right things, and for ensuring the solution's sustainability and scalability over time? Pod members constitute the stewardship body itself and therefore do not report into either stewardship committee; the pod is simply a mechanism for organizing work and responsibility, one that is defined by business requirements rather than by technology.

AI Solution Pods

Each pod comprises six roles, by design rather than convention:

AI ENGINEER

Implements and builds the solution.

AI STEWARD

Validates that the solution works and is operationally sound.

DECISION STEWARD

Ensures the risk and value of the business decision that AI is meant to support is comprehensively understood and translates business requirements into functional requirements.

DATA STEWARD

Brought in as needed, ensures the underlying data and metadata foundation is governed and available.

BUSINESS SME

Contributes domain expertise and the ability to anticipate impact in context: the human judgment on which the framework depends in ambiguous situations.

AI COMMUNICATIONS

Ensures adoption and literacy, on the premise that accountability for a solution holds only if the people using it and affected by it, understand it.

Each pod member carries dual accountability, to the pod and to their governing committee, which is how lateral knowledge and cross-functional insight flow into the organization rather than remaining trapped within a single implementation. Viewed at scale, the pods operating simultaneously across the organization form the connective tissue that isolated, narrow implementations have historically lacked.

Integration with Data Governance

The **Data Steward** sits within the AI Solution Pod, functioning as a working-group member brought in as needed throughout the AI lifecycle to support the solution. Because the Data Steward is a role within the Data Governance operating model, the position does not formally belong to the AI Governance operating model; instead, it maintains a dotted-line report into the Solution Owner.

The **Data, Analytics, and AI Leadership Committee** sits above all three stewardship committees and exists to prioritize and resource their work as a single program rather than three competing ones, for example, by ensuring the program is organized around one integrated roadmap, timeline, and set of priorities rather than three separate ones.

This committee unites data governance and AI governance without merging them: data governance evolves as and where needed in service to AI rather than becoming it, so Data Stewardship, AI Stewardship, and Decision Stewardship remain distinct committees, each with its own standards, while this layer aligns their strategies to a single shared roadmap.

The value of this structure is that it avoids two failure modes simultaneously: collapsing the disciplines, which would forfeit their distinct expertise, and allowing them to operate as silos, which duplicates effort and leaves gaps, thereby achieving strategic alignment without redundancy. It is also the level at which the Solution Owner and Product Owner sit together, ensuring that business-value priorities and technical-capability priorities are reconciled at the one point in the organization positioned to see both.

info@firstsanfranciscopartners.com
888-499-DATA