Chapter 15

Direction: the Governance Blueprint

Picture a Governance Blueprint for a specific domain. Let us say the domain is software deployment at a mid-sized enterprise technology company. On the page, there is a central horizontal spine of six nodes, representing the sequence of actions an agent might take: review and merge a pull request, deploy a service, triage a security vulnerability, respond to an incident, apply a database migration, release technical data.

Each spine node has smaller nodes attached to it, representing the specific checks and conditions that apply to the action at that spine point. The deploy-service node might have attached nodes for build artifact integrity (type: check), test coverage (type: evidence), canary health (type: probe), SLO budget (type: boundary), on-call presence (type: authority), and schema change authorization (type: evidence). Each attached node has a scrutiny level and a signer.

The whole thing is a structured representation of what this organization expects to happen before and during the deployment of each change to production. Not a document. Not a policy. A topology that executes.

This is a Governance Blueprint. This chapter describes what it is, how it is constructed, and why it works.

What a Governance Blueprint is

A Governance Blueprint is a formal object with four kinds of content.

Nodes. Each node represents a specific aspect of the organization's institutional judgment relevant to a class of decision. Nodes are typed, which is to say that each node has a declared category that determines how verification engages with it. The node types are:

- *Check*: a condition that must be true for the decision to proceed (test passes, status is active, flag is not set).
- *Evidence*: data that must be present and conforming before the decision is adequately supported (fraud model score, policy coverage document, authorization record).
- *Probe*: an active inquiry the verification layer runs that may go beyond what

the agent brings (query a secondary source, compute a property not in the agent's inputs).

- *Competency*: a capability the agent must have demonstrated for this class of decision (model has been evaluated at the required level, agent has not been flagged for anomalous behavior).
- *Verification*: a redundant check of a decision whose risk profile warrants duplicate validation.
- *Authority*: a signature or sign-off from a specific role that must be present (lead of data engineering, compliance officer, product owner).
- *Failure*: a condition that, if true, blocks the decision regardless of other checks (known-bad state, sanctions match, explicit policy violation).
- *Boundary*: a global constraint that defines the perimeter of what the agent can do (jurisdictional limits, time-of-day limits, scope-of-authority limits).

The eight types are not arbitrary. They correspond to recognizable patterns in enterprise governance that have historically existed in scattered forms — checks are what tests are, authorities are what sign-offs are, boundaries are what jurisdictional policies are. The graph formalizes these patterns as a shared vocabulary.

Edges. Edges connect nodes to represent dependencies and sequencing. A deploy node's check for test coverage may depend on an evidence node for the test suite having completed; a release node may depend on an authority node for sign-off. The edges allow verification to understand the structure of what must be true for the decision to be well-supported.

Scrutiny level configuration. Each node has a specified scrutiny level — minimal, standard, rigorous, or adversarial. The scrutiny level determines the depth of verification applied at the node. Chapter 17 develops the four levels in detail.

Signatures. Each node has a signer — a specific person or role in the organization who has committed to the node's content and to its ongoing maintenance. The signature is literal: a formal commitment recorded in the graph's metadata, with a timestamp and a review cadence. Signatures are not perpetual; they expire if not renewed, which forces periodic revisitation of the graph's content against the organization's current state.

What a Governance Blueprint is not

Five things it is often mistaken for.

A Governance Blueprint is not a workflow definition. Workflow tools (BPMN, state machines, orchestration engines) describe the sequence of

actions an agent should take. The Governance Blueprint describes what should be true about each action. The two coexist; they are different objects.

A Governance Blueprint is not a knowledge base. Knowledge bases store information the agent may consult. The graph stores what the organization expects about decisions in a class. The two coexist; they are different objects.

A Governance Blueprint is not a set of instructions for the agent. The operating agent makes decisions using its own instructions, models, tools and context. The Governance Blueprint independently represents how the organization expects consequential decisions to be made. The Steward compares the agent's observable decision trajectory against that organizational reference. An agent can be well trained and well instructed and still reach a decision by a different path from the one a competent employee would take. The Blueprint makes the expected path explicit: what must be looked at, weighed and ruled out on the way to a decision in this class, and how those considerations depend on one another. This is the expected **decision trajectory** introduced in Chapter 2. The workflow tells us what the agent does; the decision trajectory shows how it arrives at a decision; the Blueprint makes visible the judgment the organization expects. Most organizations have a workflow for such decisions. Few have the trajectory.

Workflow	Decision trajectory
What actions or tasks happen, and in what sequence	How a particular decision develops toward an outcome
Which tools, systems and process steps are involved	What information is considered, which conditions matter, what is omitted, and when escalation is needed
Maps operational execution	Makes organizational judgment and decision dependencies visible
Can run entirely as designed	Can diverge from the expected decision basis even when execution succeeds

The expected trajectory is not one employee's reasoning copied step by step. Different people reach sound decisions by different paths; the Blueprint holds what any sound path must include: the necessary considerations, their dependencies, the boundaries, and the points where someone must step in. The Steward, for its part, can compare only what is observable or evidenced in the agent's trajectory, not every element of its internal reasoning.

Gigerenzer's work on heuristics is the strongest argument that expert judgment is not a rule set. Experienced people decide with fast-and-frugal heuristics fitted to their environment, and under real uncertainty, where the future does not resemble the past, those heuristics often beat the algorithm (Gigerenzer, 2022; Katsikopoulos & Gigerenzer, 2025). The Blueprint does not try to write the heuristic down. It records what the heuristic attends to: the household rather

than the claimant, the pattern rather than the count. That is what a heuristic can be interrogated for even when it cannot be stated as a procedure, and it is the honest limit of the reference. It holds the considerations experienced judgment uses, not the judgment itself. Chapter 24 returns to what that leaves out. Where the agent's trajectory and the expected one diverge, the Steward has something specific to surface, and an owner has something specific to judge.

A Governance Blueprint is not an org chart or an RACI matrix. Those tools describe who is responsible for what in a general organizational sense. The graph describes who has signed a specific node that governs a specific aspect of agent decisions. The specificity is important: a general "responsible for compliance" assignment does not survive translation into automated enforcement, but a signature on a specific node does.

A Governance Blueprint is not a one-shot artifact. The graph is a living object, updated continuously as the organization's institutional judgment evolves. New nodes are added when new patterns are recognized; existing nodes are modified when the responsible party updates their position; signatures are renewed on cadence or re-executed when content changes.

How a Governance Blueprint is constructed

Constructing a Governance Blueprint for a given decision class is a specific craft. The playbook developed in Chapter 26 goes through it in detail. The compressed version here establishes the shape of the work.

The construction has six phases.

Scope. The direction designer, working with domain stakeholders, defines the decision class the graph will govern. "Software deployment" is a scope. "Insurance claim adjudication for product line X" is a scope. "Supplier qualification for commodity Y" is a scope. The scope is the universe of decisions the graph will apply to.

Judgment extraction. The designer interviews the responsible parties who currently exercise judgment in the scope. The interviews surface what matters, what rules apply, what prior incidents have taught the organization, what distinctions the responsible party makes that are not currently documented. The output is a corpus of institutional judgment, often several times the size of what the designer initially expected.

Node proposal. The designer translates the extracted judgment into draft nodes, assigning types and scrutiny levels. The draft is a first approximation; it will be revised.

Signer identification. For each proposed node, the designer identifies who should be the signer. The signer is the person with domain authority over the judgment the node encodes, who can both commit to the content and maintain it over time. Signer identification often surfaces organizational gaps — nodes where no one person owns the judgment, requiring the organization to clarify ownership before the node can be deployed.

Signer review and signature. The designer reviews each proposed node with its intended signer. The signer may revise the content, challenge the scrutiny level, or push back on the node's necessity. The negotiation produces nodes that the signer is willing to formally commit to. The signatures are then captured.

Deployment and iteration. The signed graph is deployed into the verification architecture. Initial operation is typically in observation mode — the verification runs against each decision, but the intervention ladder is set to log-only. Observation reveals edge cases, missing nodes, and over-cautious nodes. The designer iterates based on observation, re-engaging signers as needed, before the graph is promoted to enforcement.

The whole cycle, for a well-scoped domain, takes between six weeks and six months, depending on organizational maturity and the number of nodes required. The first graph in an organization typically takes longer because the craft is new; subsequent graphs are faster because the pattern is established.

Why the signing matters

The specific property that makes the Governance Blueprint work — and that distinguishes it from previous attempts at codifying institutional judgment — is the signing.

Previous codification attempts have typically been authored by one party (often legal, compliance, or a dedicated policy team) and distributed to the organization for compliance. The authoring party owns the content; the rest of the organization is expected to follow it. This model has a well-known failure mode: the content drifts from the actual institutional judgment because the authors do not have the operational context to maintain it, and the operators do not own the content, so neither party keeps it current.

The Governance Blueprint inverts this. The signer is the person who owns the operational judgment in the domain. They are the authority on what should be checked and how. Their signature means they will maintain it, defend it, and update it as the organization learns. The content lives with its natural owner.

This matters for three reasons. The content is more accurate because it is authored by the person with the deepest context. The content is maintained because the signer is accountable for maintenance. The organization has a clear

party to escalate to when a verification outcome is contested, because the signer is identifiable and committed.

The signing is not just a formality. It is the architectural mechanism that solves the drift problem that has defeated every previous attempt at codifying institutional judgment in enterprise systems.

What the graph enables

Once a Governance Blueprint exists and is signed, the architecture can do things that were previously impossible.

It can verify every agent decision against the applicable subgraph in milliseconds, instead of batching for later review.

It can respond to verification outcomes proportionally through the intervention ladder, instead of binary pass-or-block.

It can escalate specific issues to specific signers when the architecture identifies that a node's judgment is implicated, instead of generic "please review" messages.

It can produce audit records of what was checked, what was found, and who was responsible, for every decision, retrievable on demand.

It can enable rapid update cycles — when the organization learns something new about a decision class, the relevant signer updates the relevant node and the update is live for all subsequent decisions within hours.

It can scale, because the architecture does not require more human attention as decision volume grows; only more capacity in the verification infrastructure.

Each of these properties is a specific capability that classical governance did not have and that shift-left alone cannot provide. The next chapter turns to the artifact that bounds how the agent operates while exercising the authority the Blueprint grants: the Service Integrity Agreement. Scrutiny levels, the tool that calibrates how deeply verification engages with each node, follow in Chapter 17.
