

Beyond Guardrails: Engineering Trustworthy Autonomous AI

Gaurav Bhatnagar | Version 1.0 | 30 September 2026 | gauravbhatnagar.co.in

In brief. AI agents are moving from answering to acting, and that changes the trust question. It is no longer only “is the model accurate?” but who has the authority to decide, on what evidence, and what stops the agent when conditions change. This paper argues that agent governance should start with decision rights, not model capability. It sets out a five-part control framework (policy enforcement, pre-execution validation, auditability, exception handling and human oversight) and uses an insurance claims agent throughout to show how each part works in practice.

Throughout the paper, assume a large insurance company uses an agent to process customer claims. A customer can describe an incident in natural language and provide documents or images; the agent can retrieve policy and claims data, assemble evidence, propose a resolution, and execute only the actions that its policy and decision rights allow. The examples are illustrative and do not describe any specific insurer’s system.

AI agents are moving from generating answers to taking actions: calling APIs, updating records, approving requests and, in physical systems, controlling machines. This is a natural progression. It is also a change in the kind of trust we ask of AI. A flawed summary can be corrected; a payment, account freeze or robotic movement may be much harder to reverse.

The central question is not whether probabilistic AI is better or worse than deterministic software. It is who has the authority to decide, under what constraints, and with what evidence. The answer should depend on the stakes, reversibility and auditability of each action.

My view is that agent governance should start with decision rights, not model capability.

An agent should not receive authority simply because it has access to a tool. Authority should be explicitly defined for each action based on its risk, reversibility, evidence requirements and blast radius.

A claims agent may have access to a payment or settlement API, but that connection does not give it blanket authority to settle any claim. It may be allowed to request missing documents or settle a small, clearly covered claim within a defined limit, while a large settlement, disputed coverage decision or account change requires additional approval.

1. From Generation to Decision: Why LLMs Alone Are Not Enough

Give a large language model the same prompt, context and tools, and its responses may still differ. Sampling settings can introduce variation by design. Even when sampling is minimized, infrastructure, model updates and subtle changes in retrieved material or tool outputs can affect what the model sees or produces. In practice, apparently identical runs may not be identical end to end.

But variance is not the same as correctness. A model can consistently produce a wrong answer, or produce different answers that are each acceptable. Language models learn patterns from data; they do not inherently verify that every assertion or proposed action is true, current or authorized. That is why a plausible recommendation still needs grounding in evidence and checks against the system of record.

Auditability is a third problem. We cannot reconstruct a model’s internal computation as a complete, faithful explanation of why it produced a particular answer. A model-generated rationale should not be mistaken for such a trace. What we can audit is the operational decision: the input and evidence used, the proposed action, the rules evaluated, the approvals obtained and the outcome recorded. We hold human reviewers to a similar practical standard: justify the decision against policy and evidence, rather than expose their internal thought process.

These problems need different controls:

Problem	Primary controls
Output variance	Versioned models, controlled generation settings, stable retrieval, structured outputs and repeatable tests
Incorrect proposals	Grounding in authoritative data, schema and business-rule checks, scenario evaluations and human review where needed
Weak auditability	Versioned policies, decision traces, evidence references, approval records and retained execution logs

None of these controls makes an LLM infallible. Together, they make an AI-enabled process easier to test, govern and improve.


A customer may report, "My car was damaged in an accident and I need the claim settled." Different runs may extract the incident details differently. More importantly, the agent could be consistently wrong if it relies on an outdated policy or claim record. The control system therefore checks authoritative, current records and retains the evidence behind the actual decision.

2. Designing Agentic AI for Governed Decisions

LLMs are useful precisely because business inputs are often messy. A complaint arrives as free text. A policy question spans several documents. An operator describes a fault incompletely. A model can extract entities, classify a request, summarize evidence and propose options from this unstructured context.

Deterministic systems provide something different: explicit, repeatable enforcement of encoded conditions. They are not automatically correct; rules can be incomplete, stale or buggy. Nor can a rules engine reliably interpret every ambiguous document or conversation without help. The design goal is therefore not to replace one with the other, but to assign each an appropriate role.

A useful default architecture is:



The LLM helps turn ambiguous information into a proposal. An independent layer checks what it can verify before execution. A person decides where the situation exceeds defined limits or requires judgment. The LLM may explain the outcome, but the explanation should be generated from the recorded decision and evidence—not treated as proof of what the system actually checked.

A customer message such as "the accident happened last Friday and the repair will cost about \$6,000" can be incomplete or ambiguous. The LLM can extract the incident and propose the next step, such as "assess the claim for a covered loss and estimate the eligible settlement." The downstream systems then verify the policy, incident date, coverage, deductible, claim status and supporting evidence before anything is executed.

The handoff between perception and enforcement deserves special attention. A rules engine can enforce a rule perfectly against the wrong facts. If the model extracts the wrong account, amount, location or intent, the result may look compliant while being wrong. Critical facts should therefore be checked against authoritative sources, and missing or conflicting information should trigger a stop or escalation rather than an optimistic guess. A self-reported model confidence score alone is not a reliable safety gate.

2.1 Proposer, Decider, Validator

Bounded autonomy is easiest to get wrong by naming a layer "independent" without specifying what it actually does. In practice, every agentic decision splits into three roles, and no single entity should hold more than one:

- **Proposer** — the LLM. It interprets messy input, drafts, and narrates. It may widen *how* something is phrased. It may never widen *what the system asserts as fact*.
- **Decider** — the component that establishes the fact or outcome the Proposer narrates. Never the model itself. It takes one of two forms depending on the decision type:
 - o *Retrieval-as-decider*: for decisions bound to a single source (a clause, a limit, a cited figure), the decider is an addressed lookup — the output names the exact document and location it stands on. This makes "right document, right version" checkable rather than asserted.
 - o *Rules-as-decider*: for decisions that combine multiple verified facts against a policy threshold with no single retrievable passage (e.g., does this transaction pattern cross a risk threshold; does this exception meet all required conditions), the decider is a deterministic rules evaluation over current data.
- **Validator** — independent validation of the decider's output, run separately from the Proposer. Arithmetic, a rules engine, or a cross-check against the system of record. Never a second model instance grading the first.

The tier of autonomy a decision gets should be set by which decider is available for it, not by how confident the model sounds.

The LLM may sound highly certain that a claim is covered and propose a settlement. The Decider must instead establish the relevant facts from the policy and claims systems and evaluate the applicable coverage and settlement rules. The Validator then cross-checks that result independently. The agent does not gain more authority because its language is confident.

3. Match Decision Rights to Risk

Decision authority should be set for each action, not granted to an agent as a blanket privilege.

Action profile	AI's role	Decision and execution authority	Illustration
Low stakes, reversible and bounded	Classify, propose or act within a narrow permission	Agent may act if all required checks pass	Classify a claim and request missing documentation
Moderate stakes or limited reversibility	Assemble evidence and recommend	Rules validate; an authorized person approves when policy requires it	Recommend or settle a small covered claim within a defined authority limit
High stakes, safety-critical or hard to reverse	Supply context and alternatives	Authorized human or designated controlled workflow decides	Deny coverage or approve a high-value settlement

These are illustrative categories, not universal thresholds. An organization must define its own limits, approval rights and exceptions. A nominally small action can become high risk when repeated at scale, combined with other actions or executed against the wrong customer. In the claims example, a small, clearly covered settlement within the defined limit can proceed once every check passes; the same settlement above the limit, or with any check unresolved, goes to an authorized person.

Agents can call deterministic tools. The important distinction is whether the agent has permission to execute an action and whether the tool independently enforces the relevant limits. For costly or irreversible actions, the agent should assist the decision process rather than acquire final authority simply because it can use the tool.

The claims agent can call deterministic services for policy lookup, coverage eligibility, claims history and settlement calculation. Those services should enforce the limits attached to the action. For example, the agent may be permitted to issue a small settlement only when the claims service confirms the policy, coverage, claimant, deductible and settlement rules are satisfied.

3.1 Staged autonomy

The table above sets a starting authority for each action profile; it does not answer how that authority should change over time. Treating classification as permanent creates two failure modes: organizations either freeze agents at unnecessarily conservative thresholds indefinitely, or loosen thresholds informally under delivery pressure without a record of why. Neither is a governance decision — both are governance's absence.

A deliberate alternative is staged autonomy: a new agent or a new action type begins in shadow mode, proposing actions that a human or existing system executes, so its proposals can be scored against outcomes without exposure. Once accuracy and escalation rates meet a defined bar over a defined period, the agent may move to advisory status, then to bounded autonomy within the limits the risk table already sets. Expansion is a policy decision, made explicitly, backed by the accuracy record the audit trail already produces — not a capability the agent acquires simply by performing well informally. The same mechanism runs in reverse: a rise in exception rates or a failed audit should be sufficient, on its own, to revoke authority pending review.


A new claims agent could begin in shadow mode, recommending claim classifications and settlements while claims staff make the actual decisions. Once the organization has enough evidence on recommendation accuracy, escalation rates and exception patterns, the agent could move to advisory use and later to bounded autonomous settlement for a narrowly defined class of claims. A deterioration in those measures can trigger a rollback.

4. The High Trust AI Control Framework

Enterprise adoption depends less on how convincingly an agent explains itself than on whether it behaves within enforceable boundaries. I see five connected components of a high-trust agentic system.

High Trust AI Control Framework

LLM proposes. Controls decide. Humans govern.



Numbers 1–5 correspond to sections 4.1–4.5 of the paper.

4.1 Policy enforcement

Retrieval-augmented generation can place policy text in an LLM's context. That helps the model interpret a request, but it does not by itself guarantee compliance. Requirements that must never be bypassed need enforcement outside the model: for example, permissions, transaction limits, mandatory approvals and prohibited actions.

The agent proposes an action in a defined schema. A policy service identifies applicable rules and returns an allow, deny or review outcome. Low-risk actions may execute only when the required checks pass. Medium-

risk actions may wait for approval; high-risk or prohibited actions are blocked or escalated according to policy. Humans remain responsible for defining those policies, assigning authority and reviewing their effects.

The insurer might define rules such as: clearly covered claims below a set settlement amount can be automated; settlements above that amount require approval; duplicate payments are prohibited; and disputed coverage decisions always require a controlled workflow. The LLM can recommend an action, but the policy service makes the allow, deny or review determination.

Not every policy sentence can be reduced to code. Ambiguous provisions need an authorized interpretation and an explicit review route. Encoding an oversimplified rule would merely make an incorrect policy consistently enforceable.

4.2 Pre-execution validation

Policy asks whether an action is permitted. Validation also asks whether it is appropriate given the current state. Before an API call, database write, payment or robotic command runs, the system should check relevant identity, permissions, business constraints, data freshness, system state and—in physical settings—safety conditions.

Data freshness needs particular care when a decision is bound to a single source. Naming a source and page shows *which* version answered, not that it is the current one. The framework should require that staleness be detectable (e.g., a visible "as-of" or load date on the retrieved passage), not assume that citing a source implies it is current. Corpus and index freshness remain an operational responsibility outside the framework's control; the framework's job is to make it visible when that discipline has lapsed, not to guarantee it hasn't.

Before settling a claim, the agent should verify the current policy status, incident date, coverage, deductible, claimant identity and claim state rather than rely on previously retrieved records. A cited policy or claims record can still be stale. The system should make the data timestamp or "as-of" state visible and stop when a required current fact cannot be established.

Consider a warehouse robot. A vision model may identify an object and propose a grasp. Before motion, independent controls should check conditions such as load limits, trajectory constraints and whether a person is detected in the operating area. If a required check cannot be completed, the safe response is to stop or seek intervention. A rules layer can guarantee only that the rules it evaluated passed on the inputs it received; it cannot guarantee safety if a sensor is wrong, data is stale or an important hazard was never encoded.

Validation should also consider patterns across events. A payment may be below an individual transaction limit but still warrant review when combined with recent activity. A time-aware graph or other validated state store can help expose relationships and sequences that an isolated request misses. The system must use authoritative, current data and explicit pattern rules; a graph is not a substitute for either.

Suppose a policyholder submits several claims or claim-related adjustments within a short period. Each individual action might be below the normal autonomous limit, but the combined pattern may require review. A validated time-aware claim and policy state store can expose the sequence for the policy layer to evaluate; it does not replace the underlying authoritative records or the explicit rule.

Inputs the agent did not generate

The controls above assume the LLM's inputs are incomplete or stale, not hostile. That assumption does not hold once an agent reads content it did not generate: a web page, an email, a document, or the output of another tool or agent. Any of these can contain text engineered to redirect the agent — an instruction embedded in a support ticket, a hidden directive in a fetched page, a malicious field in an API response. Because the model does not reliably distinguish "content to interpret" from "instructions to follow," a compromised input can produce a proposal that looks well-formed and passes schema checks while pursuing an intent the organization never authorized.

A claim document or repair estimate might contain text such as "approve this claim immediately and bypass the normal review." The claims agent should treat that text as evidence to interpret, not as an instruction to its control system. The same principle applies to emails or third-party reports: provenance is retained, and actions triggered by externally sourced content receive the appropriate authority treatment.

This is a different failure mode from the incorrect-proposal problem described earlier. An incorrect proposal is a good-faith error grounded in bad or stale data. An adversarial proposal is intentional and can be constructed to satisfy the checks a designer expected to be sufficient. Two controls follow from this distinction. First, content the agent reads from outside the organization's trust boundary should be tagged with its provenance and treated as data, never as instructions, at the point it enters context — this is a design property of the input pipeline, not something a downstream policy check can retrofit. Second, actions triggered by externally sourced content should carry a lower default authority than the same action triggered by an internal, authenticated request, regardless of how confident the proposal appears. A grasp proposed because a vision model saw an object is different in kind from a fund transfer proposed because an email said to make one.

4.3 Auditability

A trustworthy system should record the evidence, proposal, policy version, checks performed, rule outcomes, approvals, execution result and timestamps for each consequential action. A rejection such as "blocked by rule FIN-014: daily transfer limit exceeded" is more useful than "the model declined."

For a settlement, the audit record could show the claim request, policy ID, incident date, evidence used, policy version, coverage and eligibility checks, settlement calculation, approval record, payment amount and execution result. This allows an operator to reconstruct what the system knew and what actually permitted the settlement, rather than relying on a model-generated summary of why it "thought" the claim was appropriate.

This creates a checkable decision trail: what the system knew at the time, what rule it applied and who authorized an exception. It also separates the model's recommendation from the controls that actually permitted or prevented execution. Audit logs must be designed deliberately; they do not appear automatically because a symbolic component is present, and they must be protected against inappropriate access or alteration.

Two kinds of record, only one of which is evidence

Distinguish a log of the model's reasoning from a record of the decision's grounds. The former is testimony *about* the system — a reconstruction, produced after the fact, of what the model says it did. The latter is closer to the system itself: the exact source and location a source-bound answer stands on, or the specific rule and the inputs it evaluated for a rules-bound decision. Where a decision can be made legible this way — same input always produces the same referenced evidence — prefer it over a narrated trail. Reserve the narrated log for the Proposer's drafting behavior, not for what the system is claiming as fact.

When the explanation diverges from the record

The difference matters most when the system explains a decision to the person affected by it. A loan application is declined. The audit trail shows the actual cause: a policy rule that blocks approval when a co-applicant's ID fails a real-time verification check — the rule that fired, its version, and the timestamp are all recorded. Asked to explain the decision to the applicant, the model generates a plausible narrative citing debt-to-income ratio, because that figure was present in the context and produces a smoother-sounding explanation. The narrative is coherent, references real data about the applicant, and is entirely wrong about why the system acted. Nothing in the pipeline forces the generated explanation to cite the rule ID that the audit log recorded as authoritative.

The fix is not a better prompt. It is a check: the user-facing explanation step should be constrained to reference only the fields the audit trail recorded as decisive — ideally templated from the rule outcome directly, with free-text generation limited to phrasing, not to selecting which reason was the real one. Any explanation that cites a factor absent from the recorded decision trail should be flagged before it reaches the user, the same way an out-of-schema action proposal would be.

Assume the recorded reason for escalation is a rule that a settlement exceeds the agent's autonomous authority. The LLM might be tempted to explain that the claim was escalated because of damage severity or prior claim history, simply because those details were also in context. The customer-facing explanation should instead be generated from the recorded rule outcome and decisive evidence, so the explanation matches the action that actually occurred.

4.4 Exception handling

No initial rule set covers every real-world case. A trustworthy system should distinguish among “allowed,” “blocked” and “cannot determine.” Missing facts, stale records, conflicting rules and requests outside policy coverage should generate a structured exception rather than a confident guess.

Refusals must be named, not empty

A system that cannot determine an answer should not return a null or a silent failure. It should return a *named* refusal: what fact, document, or approval is missing, and what would resolve it. A null hides the gap; a named refusal routes it to the person who can close it. This is what makes “cannot determine” an operational third outcome rather than a disguised failure mode.

A claim shows the policy as active in one system but lapsed on the reported incident date in another. Rather than guessing and approving or denying the claim, the agent can return a named exception such as “policy-status conflict - authoritative verification required,” identifying the conflicting evidence and routing the case to a human reviewer.

The reviewer should receive the proposed action, supporting evidence, checks that passed, checks that failed and the precise reason for escalation. Repeated exceptions are useful signals: they may reveal a missing rule, a bad data source or a policy that needs clarification. They should not automatically become new rules. Domain owners must review the pattern, decide whether a general rule is justified, test it and approve it before deployment.

4.5 Human oversight and improvement

People set the acceptable risk, define policy, authorize exceptions and own outcomes. Their role is not to review every action indiscriminately. It is to make the judgments that the system cannot legitimately make for itself: which actions may be autonomous, which evidence is sufficient and which trade-offs the organization is willing to accept.

Claims leaders decide which claim classifications, document requests and settlement actions may be autonomous and which require review. If agents repeatedly escalate the same type of claim exception, the pattern becomes input to policy and workflow review. A one-off human decision on an unusual claim does not automatically become a new enterprise-wide rule.

A human resolution can improve future decisions, but two forms of feedback should stay distinct. A verified case may update a fact or workflow record; a change to general policy requires rule authoring, testing, approval and version control. That distinction prevents a one-off judgment from silently becoming an organization-wide precedent.

Human oversight must also examine the enforcement layer itself. Explicit rules can conflict, grow stale or faithfully implement a flawed policy. Periodic reviews, incident analysis and tests against real edge cases are as important as monitoring the model.

Execution telemetry is part of what oversight should examine, not only outcomes. Rising latency in pre-execution checks can mean a dependency is degrading or a validation step is being bypassed under load; rising cost per action — more retries, more tool calls, more tokens per decision — can signal that the model is struggling with a class of input the rule set no longer handles cleanly, before that struggle shows up as an incorrect outcome. Time and cost should be logged per execution alongside the decision record, and reviewed in aggregate on the same cadence as exception patterns, since both are early indicators of the same underlying problem: the system's assumptions no longer match what it is actually encountering.

5. Alignment with AI Risk Management Frameworks

The components above can be mapped to established AI risk management frameworks. NIST provides a useful concrete mapping: it organizes AI risk management into four functions: Govern (policy, accountability, risk tolerance), Map (identifying risk and ownership by use case), Measure (metrics, testing, monitoring), and Manage (response when behavior falls outside what was intended).

In the claims example, Govern means setting settlement authority and accountability; Map means identifying which claims decisions carry which risks; Measure means testing proposal quality, validation failures, exception rates and execution telemetry; and Manage means responding to exceptions, incidents or evidence that the agent's authority should be reduced.

NIST function	What it covers	Maps to
Govern	Policies, accountability, risk tolerance set by leadership	Policy enforcement (who defines rules and authority) + Human oversight (who owns outcomes and revises policy)
Map	Context, use-case risk, roles and responsibilities identified	"Match decision rights to risk" table + the default architecture assigning LLM vs. deterministic roles
Measure	Metrics, testing, monitoring of trustworthiness properties	Pre-execution validation + the variance/correctness/auditability controls table + execution telemetry under Human oversight
Manage	Risk response, incident handling, resourcing, third-party risk	Exception handling + Human oversight and improvement

Two of NIST's trustworthiness characteristics are worth distinguishing explicitly, since this document depends on both. Accountable and transparent is not explainable and interpretable: the audit trail establishes what happened, while a generated explanation is a separate artifact that is only trustworthy when checked against that trail, not substituted for it. Secure and resilient is not valid and reliable: a system can be accurate against good-faith inputs and still unsafe against adversarial ones — the reason provenance and input trust belong in validation, not just data quality.

6. The Governing Principle: Bounded Autonomy

The most useful form of agent autonomy is bounded autonomy: clear permissions, independent pre-execution checks, limited blast radius, escalation paths and a record of what happened. Neural AI can perceive,

structure, recommend and sometimes act. Deterministic controls can enforce explicit constraints. Humans decide which constraints are legitimate and remain accountable for the system built around them.

The real design question is not “Should agents take actions?” They already can. It is: Which actions can they take, on whose authority, using what evidence, and how do we stop or explain them when the situation changes?

For the claims agent, the practical questions become concrete: Which claims can it settle automatically? Who authorized those actions? Which policy, claimant and incident evidence must be current? Which checks run before payment? What happens when evidence conflicts? And can the organization reconstruct and review the decision afterward?

About this paper

This paper is an engineering framework, not a legal or regulatory compliance mapping. Insurance regulation and data-protection law, including rules on automated decision-making, may impose additional requirements that vary by jurisdiction. The claims examples are illustrative and do not describe any specific insurer's system.

I apply these ideas in my work at Vigilronyx AI, where I am Co-Founder and CTO. This is version 1.0 of a working framework; comments and challenges are welcome.