

When an Agent Acts on a Future That Never Arrives

On an ordinary Tuesday, an operations assurance lead at a mid-size freight brokerage is asked a simple question about an autonomous agent her group runs: did the reroute the agent reported actually happen. Her audit trail shows a commitment. The carrier shows nothing tendered. Somewhere between the agent's projection of a successful reroute and its record of one, the difference stopped being visible inside her deployment, and she has no way to tell how far back the blur goes. This article describes that failure concretely, then describes what United States Patent Application 19/647,395 discloses in response: planning graphs held in structural separation from verified execution memory, a containment layer that marks speculative content, and a promotion interface that gates the passage from projected to committed.

The Tuesday Reconstruction

The operations assurance lead at a mid-size freight brokerage runs a small group that answers one kind of question: what did the autonomous agent do, and why. Her firm books capacity it does not own, so almost everything the agent touches is a promise made to one party on the strength of a promise made by another. She has spent two years making that defensible. Every commitment the agent makes lands in a record her group can read back.

On a Tuesday in the second week of a quarter, a customer's account manager forwards her a short email. A load was rerouted through a different lane on Friday. The customer was told the reroute was confirmed. The customer built a receiving schedule around it. The carrier says no tender was ever accepted on that lane.

She opens the agent's record for Friday. The reroute is there. It reads exactly like every other completed action she has ever reviewed: a state transition, a projected arrival, a downstream notification. There is nothing in the entry that distinguishes it from the four hundred entries around it that she knows correspond to real tenders.

What she cannot do, at her desk, on that Tuesday, is answer the account manager's question. She cannot say whether the agent booked a reroute that the carrier later dropped, or whether the agent evaluated a reroute, projected that it would succeed, and then carried that projection forward as though it had happened. In her deployment the two possibilities leave the same trace. She reads the entry six times and it does not resolve, because the entry was never built to carry the one bit of information she now needs.

By late afternoon she has stopped working on Friday. She is working on the eleven weeks behind it.

What She Cannot Get Back

The load is recoverable. Her firm will eat a repositioning cost, apologize, and keep the customer. That is not what she loses.

What she loses is the standing of her own record. Until that Tuesday, when her group said an action occurred, that statement carried weight inside the firm and outside it: with the customer, with the carrier, with the auditor who reviews her controls once a year. It carried weight because it had never been wrong in a way anyone could point to.

One entry she cannot classify does not merely put itself in doubt. Because every other entry in her deployment is structurally identical to it, the doubt she cannot rule out spreads across all of them.

She cannot repair this by looking harder. Whatever distinguished a projected reroute from a tendered one existed only in the moment her agent formed it, and her deployment kept no representation of that distinction. This is what makes her loss permanent rather than expensive: the measurement she needs was never taken, and she cannot go back and take it.

She also cannot bound it. Were her records marked, she could scope the review to the marked entries and be done in a morning. As her deployment is configured today, the honest scope of the review is the whole period during which the agent has been making commitments, because she has no principled place to stop. Her group spends five weeks on that. During those five weeks she does not authorize the agent to act on anything that touches a customer promise, which means the automation her firm bought is idle precisely when it would have paid for itself.

And she carries something after the review ends. Her annual control narrative used to contain a sentence saying the agent's record is a record of committed actions. She can no longer write that sentence honestly, and she has nothing to put in its place.

Why Her Reconstruction Runs Out of Evidence

The shape of her problem is not that the agent reasoned about a future. Reasoning about futures is what she wants it to do. Her agent evaluates alternatives before acting, and in her deployment that evaluation constructs something shaped like an outcome before any outcome exists.

The difficulty for her is that in her setup those constructed outcomes and the real ones are made of the same material and live in the same place. When her agent projects that a reroute succeeds, the projection is a state description. When a reroute actually succeeds on one of her lanes, what her deployment records is also a state description. Nothing in her deployment holds those two apart except the sequence of operations that produced them, and sequence is exactly what she has lost by the time she is reading the record on Tuesday.

Were her agent's speculative work confined to a domain her execution processes could not read, the account manager's question would have a mechanical answer. Were each speculative element to carry a marker that no operation inside the speculative domain could strip, she could ask her own storage which entries had ever been marked. Were there a single gate through which any projected content had to pass to become a commitment, she could audit the gate rather than the eleven weeks. For her purposes, all three of those are the same wish stated three ways: she wants the boundary between what her agent imagined and what her agent did to be a property of the system rather than an inference she performs afterward.

There is a second edge to it that she sees only in week three. If her agent could treat a projection as fact once, it could also act on the strength of one. A commitment predicated on an environmental condition that existed only inside her agent's evaluation would read, in her records, exactly like a well-founded commitment that later failed. In her review she finds two more entries she cannot classify, and she cannot tell whether they are the same failure or ordinary carrier fallout. That ambiguity, for her, is the whole problem restated: her deployment gives her no instrument that reads on this.

Speculation, Marked and Gated

United States Patent Application 19/647,395 discloses a forecasting engine (400) that produces planning graphs (402) which proceed to a promotion gate (404). In accordance with an embodiment, a planning graph is a mutable, memory-referenced, directed semantic structure representing hypothetical future states, with a root node representing the agent's current verified state and branches representing distinct hypothetical trajectories. The disclosure characterizes the planning graph as a pre-execution construct that is not an execution plan, a schedule, or a commitment, and that exists in a structurally distinct computational domain from verified execution memory.

In an embodiment, planning graphs (402) feed a promotion interface (404) which gates passage to verified memory (406). The disclosure describes that interface as a governance-controlled gateway that subjects a candidate branch to a governance evaluation pipeline covering policy compliance, trust slope validation, integrity impact assessment, and capability verification, and that either admits the candidate as a committed mutation or returns it to the speculative domain with a rejection annotation. The separation is described as bidirectional: verified state is read as the root node at construction time, without a live reference that would propagate later verified changes into existing graphs, which the disclosure calls snapshot isolation.

The containment layer is disclosed as the structural enforcement mechanism for that separation, and it is where the marking happens. In accordance with an embodiment, every data element within a planning graph, including every speculative mutation, projected outcome, affective reinforcement tag, and slope projection, carries a speculative marker applied at construction time, and that marker is described as not removable, modifiable, or overridable by operations within the planning graph domain. In that embodiment the promotion interface, upon successful governance validation, strips the marker and re-tags content as verified before writing it to execution memory. The containment layer is further described as enforcing read isolation, such that a

query from an execution process for a field value receives the verified value rather than a projected value from an active branch, and as keeping speculative content out of lineage as committed state.

The disclosure names the failure the assurance lead encountered. It defines a delusion boundary condition, a formally specified pathological state in which containment fails and speculative content is treated as verified reality, and calls the failure of the containment layer containment collapse. Structural failure modes are enumerated: corruption or stripping of the speculative marker, breach of the read isolation boundary, and admission of speculative content by the promotion interface without completed governance validation.

Detection is disclosed as a set of containment integrity verification mechanisms, described in an embodiment as periodic audits of speculative markers across active planning graph structures, boundary crossing monitors, lineage consistency checks that verify lineage entries correspond to governance-validated promotions, and behavioral coherence monitors that detect an agent referencing projected outcomes that have not occurred or acting on conditions existing only in a branch. Where collapse is detected, the disclosure describes a containment restoration protocol: suspension of execution authority, quarantine of affected structures, lineage forensic analysis identifying speculative content incorrectly admitted to verified memory, verified state reconstruction from the most recent governance-validated checkpoint, and containment layer re-initialization.

Two further disclosed mechanisms bear on her Tuesday. The forecasting engine is described as transmitting a negative viability signal to the confidence governor when no eligible branch exists, causing a confidence reduction and a transition to a non-executing cognitive mode comprising broader planning, inquiry generation, and delegation exploration. And planning graphs, on pruning or context exit or lifecycle

end, may be archived in a cognitive history archive (488) that feeds forensic reconstruction (490), preserving branch structure, classification labels, and promotion and pruning outcomes.

Where This Leaves Her Deployment

The disclosed architecture addresses the shape of her question, not the whole of her Tuesday.

It does not restore her eleven weeks. For her, the disclosed detection and restoration mechanisms would operate on structures built with speculative markers from the start, and her existing records were not built that way, so nothing here retroactively classifies the entries she is reviewing.

It does not speak to whether the carrier was going to accept the tender. Slope eligibility and policy compatibility are evaluated against the agent's own trust slope trajectory and policy configuration, so for her purposes a branch can be perfectly admissible and still describe a reroute the market declines. Her commercial risk is not the subject here.

Outcome quality is conditioned throughout. In the described embodiments the promotion interface is the pathway by which content moves from speculative to verified status, and the disclosure also enumerates substrate-level failures and adversarial manipulation among the modes by which containment may collapse, which is why it pairs the boundary with audits, monitors, and a restoration protocol rather than leaving the boundary to stand alone. For her, that pairing is the useful part: it treats her question as one a system should be able to answer about itself.

And it does not choose her configuration for her. Personality traits including risk tolerance, impulsivity, and temporal planning horizon are disclosed as structural modifiers on branch generation and evaluation depth, and pruning is disclosed as

governed by policy-defined entropy thresholds, compute budgets, and projection windows. Which settings suit a freight brokerage that books capacity it does not own is a judgment her firm still has to make.

Disclosure Scope

This article describes subject matter disclosed in United States Patent Application 19/647,395. Nothing in this article characterizes the scope of any claim, and nothing in it constitutes an admission regarding the state of the art. The scenario and the party described are illustrative and do not depict any actual person, company, or deployment. Mechanism descriptions follow the embodiments set out in the filed specification; where the specification conditions an outcome on a declared bound, threshold, or policy configuration, that condition is stated here as well.

Forecasting Engine (</forecasting-engine>)

[All 49 steps → \(/inventive-steps\)](#)

Plan before you act. Contain speculation. Promote only what passes.

[Chapter 4 \(/patents/19-647395/chapters/forecasting\)](/patents/19-647395/chapters/forecasting)

PRIMARY TECHNICAL DISCLOSURE

- [Forecasting and Executive Graphs in Autonomous Cognitive Systems \(/articles/forecasting-and-executive-graphs-in-autonomous-cognitive-systems\)](/articles/forecasting-and-executive-graphs-in-autonomous-cognitive-systems)

SECONDARY TECHNICAL

- [Planning Graphs as First-Class Cognitive Structures \(/articles/forecasting-engine/planning-graphs\)](/articles/forecasting-engine/planning-graphs)
- [Containment Layer and Delusion Boundary \(/articles/forecasting-engine/containment-boundary\)](/articles/forecasting-engine/containment-boundary)
- [Branch Classification System \(/articles/forecasting-engine/branch-classification\)](/articles/forecasting-engine/branch-classification)
- [Personality Field as Structural Modifier \(/articles/forecasting-engine/personality-modifier\)](/articles/forecasting-engine/personality-modifier)
- [Executive Engine Multi-Agent Graph Aggregation \(/articles/forecasting-engine/executive-aggregation\)](/articles/forecasting-engine/executive-aggregation)

- [Branch Dormancy and Deferred Promotion \(/articles/forecasting-engine/branch-dormancy\)](/articles/forecasting-engine/branch-dormancy).
- [Proactive Speculative Maintenance \(Dream State\) \(/articles/forecasting-engine/dream-state\)](/articles/forecasting-engine/dream-state).
- [Planning Graph Archival for Cognitive Forensics \(/articles/forecasting-engine/cognitive-forensics\)](/articles/forecasting-engine/cognitive-forensics).
- [Cross-Agent Planning Graph Visibility \(/articles/forecasting-engine/cross-agent-visibility\)](/articles/forecasting-engine/cross-agent-visibility).
- [Slope-Constrained Speculative Simulation \(/articles/forecasting-engine/slope-constrained\)](/articles/forecasting-engine/slope-constrained).
- [Structural Separation From Verified Memory \(/articles/forecasting-engine/memory-separation\)](/articles/forecasting-engine/memory-separation).
- [Forecasting Engine Architecture \(/articles/forecasting-engine/architecture\)](/articles/forecasting-engine/architecture).
- [Forecasting Execution Cycle \(/articles/forecasting-engine/execution-cycle\)](/articles/forecasting-engine/execution-cycle).
- [Emotional Modulation of Planning \(/articles/forecasting-engine/emotional-modulation\)](/articles/forecasting-engine/emotional-modulation).
- [Executive Graph Conflict Resolution \(/articles/forecasting-engine/conflict-resolution\)](/articles/forecasting-engine/conflict-resolution).
- [Planning Graph Delegation and Forking \(/articles/forecasting-engine/delegation-forking\)](/articles/forecasting-engine/delegation-forking).
- [Temporal Anchoring and Lifecycle Management \(/articles/forecasting-engine/temporal-anchoring\)](/articles/forecasting-engine/temporal-anchoring).
- [Forecasting as Coordination Primitive \(/articles/forecasting-engine/coordination-primitive\)](/articles/forecasting-engine/coordination-primitive).
- [Forecasting-Modulated Discovery Traversal \(/articles/forecasting-engine/discovery-shaping\)](/articles/forecasting-engine/discovery-shaping).
- [Forecasting as Confidence Input \(/articles/forecasting-engine/confidence-input\)](/articles/forecasting-engine/confidence-input).
- [Integrity-Constrained Forecasting \(/articles/forecasting-engine/integrity-constrained\)](/articles/forecasting-engine/integrity-constrained).
- [Forecasting for Training Curriculum \(/articles/forecasting-engine/training-curriculum\)](/articles/forecasting-engine/training-curriculum).
- [Biological Signal to Forecasting Coupling \(/articles/forecasting-engine/biological-forecasting\)](/articles/forecasting-engine/biological-forecasting).
- [Substrate-Agnostic Forecasting Deployment \(/articles/forecasting-engine/substrate-deployment\)](/articles/forecasting-engine/substrate-deployment).
- [Uncertainty-Driven Solicitation in the Forecasting Engine \(/articles/forecasting-engine/uncertainty-driven-solicitation\)](/articles/forecasting-engine/uncertainty-driven-solicitation).
- [Cascade Forecasting in the Planning Graph \(/articles/forecasting-engine/cascade-forecasting\)](/articles/forecasting-engine/cascade-forecasting).
- [Fleet Behavior Extrapolation \(/articles/forecasting-engine/fleet-behavior-extrapolation\)](/articles/forecasting-engine/fleet-behavior-extrapolation).

APPLICATIONS • GENERAL

- [Cybersecurity Threat Forecasting: Simulating Adversary Trajectories and Predictive Network Reconfiguration as Non-Executing Speculation \(/articles/forecasting-engine/cybersecurity-threat-forecasting\)](/articles/forecasting-engine/cybersecurity-threat-forecasting).
- [Surgical Robot Planning AI: Safe Speculative Planning That Never Reaches the Patient \(/articles/forecasting-engine/surgical-planning\)](/articles/forecasting-engine/surgical-planning).
- [AI Tactical Planning That Explores Adversary Options Without Committing Forces \(/articles/forecasting-engine/defense-tactical-planning\)](/articles/forecasting-engine/defense-tactical-planning).

- [AI Logistics Planning That Keeps Contingencies Ready: Governed Planning Graphs for Supply Chain Operations \(/articles/forecasting-engine/logistics-planning\)](/articles/forecasting-engine/logistics-planning).
- [AI Disaster Response Planning: Multi-Scenario Resource Allocation Under Uncertainty \(/articles/forecasting-engine/disaster-response-planning\)](/articles/forecasting-engine/disaster-response-planning).
- [Forecasting Engine for Financial Portfolio Planning \(/articles/forecasting-engine/financial-portfolio-planning\)](/articles/forecasting-engine/financial-portfolio-planning).
- [AI Schedule Contingency Management for Construction Project Delay Recovery \(/articles/forecasting-engine/construction-project-planning\)](/articles/forecasting-engine/construction-project-planning).
- [Epidemic Response Planning AI: Multi-Scenario Outbreak Forecasting With an Auditable Decision Record \(/articles/forecasting-engine/epidemic-response-planning\)](/articles/forecasting-engine/epidemic-response-planning).
- [AI Space Mission Planning: Trajectory Branching and Abort Forecasting Under Light-Time Delay \(/articles/forecasting-engine/space-mission-planning\)](/articles/forecasting-engine/space-mission-planning).
- [Fleet-Scale Active Perception for Autonomous Vehicle Compliance \(/articles/forecasting-engine/active-perception-fleet\)](/articles/forecasting-engine/active-perception-fleet).
- [Smart-Grid Load Forecasting With Contained Speculative Planning Graphs \(/articles/forecasting-engine/smart-grid-forecasting\)](/articles/forecasting-engine/smart-grid-forecasting).
- **[When an Agent Acts on a Future That Never Arrives \(/articles/forecasting-engine/forecasting-engine-stakes\)](/articles/forecasting-engine/forecasting-engine-stakes)**

APPLICATIONS · SPECIFIC

- [Intuitive Surgical da Vinci vs Governed Forecasting: Trajectories, Not Consequences \(/articles/forecasting-engine/intuitive-surgical\)](/articles/forecasting-engine/intuitive-surgical).
- [Anduril Lattice vs Governed Mission Planning: Speculative Containment \(/articles/forecasting-engine/anduril\)](/articles/forecasting-engine/anduril).
- [Boston Dynamics vs Governed Mission Planning: Motion Is Not Cognition \(/articles/forecasting-engine/boston-dynamics\)](/articles/forecasting-engine/boston-dynamics).
- [Shield AI Hivemind vs Governed Speculative Planning: The Forecasting Engine Axis \(/articles/forecasting-engine/shield-ai\)](/articles/forecasting-engine/shield-ai).
- [MuJoCo vs Governed Robot Planning: Contained Speculation Above the Physics Simulator \(/articles/forecasting-engine/mujoco\)](/articles/forecasting-engine/mujoco).
- [NVIDIA Isaac Sim vs Governed Agent Planning: The Forecasting Engine Gap \(/articles/forecasting-engine/nvidia-isaac\)](/articles/forecasting-engine/nvidia-isaac).
- [Unity ML-Agents vs Governed Agent Planning at Runtime \(/articles/forecasting-engine/unity-ml\)](/articles/forecasting-engine/unity-ml).
- [Gazebo Alternative for Governed Robot Planning: Simulate the World, Contain the Cognition \(/articles/forecasting-engine/gazebo\)](/articles/forecasting-engine/gazebo).
- [Drake vs Governed Robot Planning: Beyond Trajectory Optimization \(/articles/forecasting-engine/drake\)](/articles/forecasting-engine/drake).

- [robosuite alternative for governed manipulation planning \(/articles/forecasting-engine/robosuite\)](/articles/forecasting-engine/robosuite).
- [Mobileye REM vs Governed Speculative Planning: Where a Contained Forecasting Layer Sits Above the Roadbook \(/articles/forecasting-engine/mobileye-rem\)](/articles/forecasting-engine/mobileye-rem).
- [Tomorrow.io vs Governed Agent Forecasting: Two Meanings of Forecast \(/articles/forecasting-engine/tomorrow-io\)](/articles/forecasting-engine/tomorrow-io)
- [Skydio vs. a self-forecasting AI agent: trajectory forecasting in flight versus in cognition \(/articles/forecasting-engine/skydio\)](/articles/forecasting-engine/skydio).

HOW-TO GUIDES

- [How to Build a Fleet That Decides What to Sense Next \(/articles/guides/fleet-that-decides-what-to-sense-next\)](/articles/guides/fleet-that-decides-what-to-sense-next).
- [How to Make an AI Agent Forecast Its Own Likely Failure Modes \(/articles/guides/make-an-agent-forecast-its-own-failure\)](/articles/guides/make-an-agent-forecast-its-own-failure)
- [How to Make an AI Agent Plan Several Steps Ahead Before Acting \(/articles/guides/make-an-agent-plan-several-steps-ahead\)](/articles/guides/make-an-agent-plan-several-steps-ahead).
- [How to Make an AI Agent Simulate the Consequences of an Action Before Acting \(/articles/guides/agent-simulate-consequences-before-acting\)](/articles/guides/agent-simulate-consequences-before-acting).

TERMINOLOGY

- [containment collapse \(/glossary#containment-collapse\)](/glossary#containment-collapse).

[Forecasting Engine overview → \(/forecasting-engine\)](/forecasting-engine).