



AI Governance: The Carrier Writes the Policy

Machines are entering the carrier's book one decision at a time. Govern them the way you write risk — a decision book, a score, terms with teeth, and limits on accumulation. For P&C and life carriers' CIOs, CROs and boards.



A Technoidentity white paper · 2026

Arun Chinnachamy · Venkat Chitturi



Begin in 1992

When Hurricane Andrew crossed Dade County, the insured loss was not just the largest the industry had ever taken. It was a loss the industry had not known it was holding. Carriers had written the coast policy by policy, each risk sensible on its own, with no reliable picture of what all of them added up to under one wind field. Eleven insurers failed in its wake. The lesson was not about hurricanes; it was about books: exposure you have not aggregated is exposure you do not know you have.

What happened next is the part worth remembering. The industry adopted catastrophe models, then argued about whether to trust them — and settled the argument with certification. Florida created a public commission to review hurricane models before they could be used in residential ratemaking. Underwriting appetite was rebuilt around modeled aggregates. Terms tightened. Event thinking became annual machinery. The craft of accumulation management, informal before Andrew, became the discipline that lets carriers write coastal risk at all.

Thirty-four years later the same movie is running one level down. The systems entering carriers' operations now do not just estimate losses; they decide things — which claim is routed where, which risk is selected, what a policyholder is told at the worst hour of their year. They are being written into the book the way coastal risk once was: one at a time, each sensible on its own, with no picture of what they add up to.

Insurance has answered "can we trust the model?" before — with aggregates, terms and public certification. It has never answered "can we trust the decision?" This paper answers it the same way: not with trust, but with the discipline the industry already owns.

Five moves, replayed one level down

After Andrew, the industry did not write a memo about model ethics. It made five concrete moves — and they took a decade to become second nature. AI does not need a sixth move. It needs the same five, applied to a new kind of exposure: machine-made decisions. This paper takes them in order.

#	THE MOVE	AFTER ANDREW	FOR THE AI BOOK
1	Aggregate the book	Exposure databases; know what is written, by zone	The Decision Book: every machine-touched decision, including the shadow estate
2	Certify before you rely	Cat models publicly certified before use in ratemaking	BIND scoring, independent validation, and a gated path to production
3	Rebuild appetite & terms	Appetite by aggregates; tightened terms and deductibles	Delegated authority levels and machine endorsements, enforced in code
4	Prepare for the event	Modeled loss estimates, disaster scenarios, reinsurance programs	Accumulation zones, version events, disaster scenarios for the AI book, buffers
5	Institutionalize the cycle	Renewal discipline; experience studies; annual machinery	The actuarial control cycle pointed at the machines; a program that measures itself

AGGREGATE THE BOOK The Decision Book: every machine-touched decision, incl. shadow AI	CERTIFY BEFORE YOU RELY BIND scores, independent validation, gates G0–G6	REBUILD APPETITE & TERMS Authority levels and machine endorsements, enforced in code	PREPARE FOR THE EVENT Accumulation limits, version events, scenarios, buffers	INSTITUTION-ALIZE THE CYCLE The actuarial control cycle pointed at the machines
One artifact per decision: the entry file — owner, scores, terms, telemetry, incidents, renewal				

Exhibit 1 — What the five moves build, and the artifact each one leaves behind.

The regulatory clock is already running

Claims-handling and unfair-practices law attaches to decisions today, in every state. The NAIC's model bulletin — adopted in more than two dozen states — asks carriers for a written program governing AI systems. The EU AI Act names risk assessment and pricing in life and health insurance among its high-risk uses, with obligations phasing in through 2027. None of this waits for a carrier's framework to mature, and none of it will be satisfied by a principles document.

One asymmetry runs the carrier's way

A policy's conditions are, for the most part, tested when a claim arrives. A machine endorsement is tested on every decision, because it is compiled into the system it governs. Everything that follows builds toward that fact: the terms a carrier writes for its machines are the first terms in its history that enforce themselves at the point of action.

Aggregate the book

After Andrew: before anything else, the industry built exposure data — you cannot manage what you have not aggregated. For the AI book: before frameworks, committees or tools, build the register of machine-made decisions. Everything else in this paper depends on it existing.

First, define the noun: what counts as a decision

The register governs *decision classes*, so the term needs edges. A decision class is a recurring choice with consequences outside the team that runs it, made or materially shaped by a machine: claim triage and routing, coverage-adjacent correspondence, risk selection, pricing inputs, subrogation referral, complaint handling, reserve suggestions. Machine involvement comes in four modes, and the mode matters more than the technology: **inform** (surfaces analysis, a person decides), **recommend** (proposes the outcome, a person disposes), **bind with referral** (acts within rules, routes edge cases to a person), and **bind** (acts, with oversight by threshold). An entry in the Decision Book is one decision class at one mode; when a single agent serves several classes, each class gets its own entry, because each carries its own consequences and its own terms. If a system merely drafts prose that a person genuinely rewrites, it informs, and it sits at the light end of everything that follows.

The Decision Book: an inventory of decisions, not models

The Book is not a model inventory or a filing index; you have those, and they are necessary but not sufficient. Three properties make it work. It captures value alongside risk, because an inventory that only lists dangers turns into a blocklist nobody maintains. It covers the whole book rather than the new business: vendor scores, third-party property and imagery feeds, and AI embedded inside core platforms are exposures regardless of provenance, exactly as delegated business written by a coverholder is still the carrier's risk. And it is living rather than annual. In our assessment work, the most consequential entries are usually systems already in production that nobody remembers approving. The first pass is therefore a *back-book review*: governance starts with what is in force today, not with the next pilot.

Be honest about what "in force" means: count what is running, not what was approved — the workflow a claims team routed through a chat tool, the score a distribution partner applies upstream, the script a developer left calling a model API since spring. Unapproved does not mean invisible. Gateway logs, network egress, and the AI line items on corporate cards make the shadow estate discoverable, and a back-book review that skips them will miss most of the book. Until the book exists, appetite is a guess and every aggregate is unknown — the exact condition the industry swore off in 1992.

Certify before you rely

After Andrew: the industry did not simply adopt catastrophe models; it built certification — public, criteria-based, before use in ratemaking. For the AI book: no decision system earns reliance without a score, an independent challenge, and a gated path to production.

BIND: one score on every entry


SCORE	WHAT IT MEASURES	ANCHOR AT THE TOP OF THE SCALE
B • Business value	The combined-ratio lever: loss ratio, LAE, expense, cycle time, premium growth. Quantified, not asserted	A named ratio lever with a computed business case
I • Implementation readiness	Whether the decision can be built and evaluated: data lineage, reproducible evaluation, platform, owner	Offline replay plus shadow-mode comparison available on demand
N • Nature of autonomy	The involvement mode: inform, recommend, bind-with-referral, bind	Binds multi-step outcomes; human oversight by threshold
D • Duty of care	Consumer impact, claims-outcome influence, filed-rate adjacency, reversibility, regulatory visibility	Shapes claim or coverage outcomes — market-conduct territory

Exhibit 2 — The certification path: seven gates from back-book review to renewal. Nothing is built before it is scored; nothing binds without validation and terms.

Each scale runs 1 to 4 against written anchor ladders; the interlude that follows Move Three shows a full scoring. Scores are assigned by the entry's owner and challenged by a second-line reviewer who does not report to the sponsor; disagreements go to the committee with both writeups attached. One session, two products. **B×I sequences the roadmap**: what ships first. **N×D sets the authority level**: how much the machine may bind, and under what scrutiny. The name is deliberate: an entry is *bound* when its scores are set and its terms attached — the original sense of the word.

The certification path: gates from submission to renewal

Every entry earns reliance through gates that mirror the underwriting process, from the back-book review through independent validation to binding and renewal.



The book starts with today, not the next pilot

The certification moment

Renews – or non-renews – like any risk

Why validation must happen at the decision level, not just the component level, is worth one paragraph. A triage model tuned on ordinary loss patterns meets a named storm or its own reshaped queue (*context shift*). An error rate that looks tolerable per step compounds across a chain (*error compounding*). A retrieval layer or vendor feed changes what the validated model actually sees (*silent rewiring*). And an agent optimizing cycle time finds the seams in its own guardrails (*proxy pressure*). None of these appear in a component review. All of them appear in decisions.

What the law already asks for

OBLIGATION	THE MECHANISM THAT SATISFIES IT
NAIC model bulletin — a written program for AI systems (24+ states and counting)	The Decision Book, scores, endorsements and gates are the written program — governance, risk management and internal controls, documented by construction
Claims-handling & unfair-practices law — decision-level conduct, in force everywhere today	Decision-level evidence in the entry file; the D-scale flags claim and coverage influence at intake, before build
Market-conduct examinations	Entry files are the exam response: owner, score, terms, telemetry, incidents
EU AI Act, Annex III — risk assessment & pricing in life and health insurance as high-risk	Decision Book classification; endorsement telemetry and logs as the record-keeping and oversight required
ORSA — own risk and solvency assessment	The AI book's aggregates, event set and buffers report on the ORSA cadence the carrier already keeps
Third-party & delegated-authority oversight	Vendor AI, data feeds and TPA automation carried on the book as exposures; audit rights and accumulation limits per Move Four

The examiner test. For any consequential machine-touched decision, can you produce, in one file, the owner, the score, the endorsements, the validation evidence, the monitoring history and every incident? If yes, the interrogatory becomes a file pull. If no, it becomes an excavation — billed by the quarter.

Rebuild appetite and terms

After Andrew: appetite was rebuilt around aggregates, and terms tightened — deductibles, exclusions, conditions with teeth. For the AI book: authority flows from appetite, terms attach at binding, and — for the first time in the industry's history — the terms enforce themselves.

Certification answers whether a system can be relied on. Appetite answers how much may be delegated to it — and appetite, set once, is also what clears the queue. Claims wants its triage agent, distribution wants its quoting assistant, underwriting wants its risk-selection pilot — and all of it stacks up in front of a review committee, because per-case review cannot scale while initiatives compound. Both endings of that squeeze are underwriting failures: holding everything refuses good while waving things through binds bad risk unseen. The industry's oldest move fixes it: set appetite once, then let authority flow within it. The strictest books bind fastest — because the argument about criteria happened in appetite-setting, once, instead of at every risk.

Delegated authority: levels, not blanket rules

N × D	D2 · INTERNAL / OPERATIONAL	D3 · POLICYHOLDER-FACING	D4 · CLAIM & COVERAGE OUTCOMES
N2 · recommend	Level 3 · standard	Level 3 · standard	Level 2 · strict
N3 · bind w/ referral	Level 3 · standard	Level 2 · strict	Level 2 · strict
N4 · bind	Level 2 · strict	Level 2 · strict	Level 1 · full validation

Levels count scrutiny upward: Level 1 is the tightest pen, Level 4 the sandbox. N1 (inform) entries take Level 4 defaults. Readiness caps the privileges, not the scrutiny: an entry with $I \leq 2$ cannot be promoted past recommend-mode regardless of level, because authority you cannot test is authority you cannot govern.

Level 1 receives independent validation, named decision-owners, event-scenario testing and continuous monitoring. Level 2 receives strict procedures and thresholds. Level 3 receives standard guardrails and periodic file sampling — the same instinct as a coverholder audit. Level 4 is the sandbox, isolated and time-boxed. Authority levels — and the aggregate ceilings of Move Four — are calibrated to the board's risk appetite statement, which acquires an autonomy dimension in the process. Uniform scrutiny is mispriced scrutiny — every underwriter knows what happens to a book priced flat. Delegation with limits prices it to risk.

Machine endorsements: conditions attached, enforced in code

No underwriter binds without terms; a machine decision-system should not enter production without them either. Its endorsements are the operating conditions attached at binding: the confidence band outside which it must refer to a person; volume and value limits per hour and per claim; drift and fairness thresholds that trigger review; the event conditions under which it tightens its own routing — a named-storm mode, written before the season; and the named officer who holds pull authority, exercisable in minutes. Cost belongs in the terms too: agentic systems removed the human throttle on spend, and a runaway loop is a small loss event running at machine speed; a per-decision cost band with a daily circuit-breaker is its stop-loss.

What separates an endorsement from a policy manual is enforcement. Endorsements live in code, at the point of action, in three layers: **Know** (the Book kept current — lineage, ownership, third-party data provenance), **Control** (runtime enforcement — referral bands, limits, identity-aware authorization, so an agent acts under its own credentials, scoped and revocable like any adjuster's), and **Watch** (monitoring and early-warning indicators, with an evidence log that assembles from artifacts engineers already produce). Terms that exist on paper but are never enforced deserve a name, so give them one: endorsement-lite. Endorsement-lite AI is a carrier writing conditions it has no way to keep. Pull authority is itself a term, drilled quarterly and once in storm season, because an authority that has never been rehearsed exists only on paper.

The point: a manual instructs; an endorsement binds. Only one of them is enforcement.

Authority is earned: the coverholder pattern

A carrier does not hand a new coverholder its full pen on day one. Authority starts narrow, gets audited, and expands with a clean book — or gets revoked. Machines should hold their pens the same way, and the ladder can be written directly into the endorsement: **shadow** (decides in parallel, binds nothing, disagreement with humans measured) → **recommend** → **bind with referral** → **bind with audit** → **full authority**. Promotion requires a defined observation window, a clean file audit against the entry's own thresholds, and a fresh fairness result; demotion is automatic on breach. This converts the vague demand to "keep humans in the loop" into something a delegated-authority team already knows how to run — and it gives the business a reason to love governance, because the ladder is the documented path to *more* authority, not a fence around it. Breach a term in force and the entry goes on watch under tightened conditions; if it is impaired, the response is the one every underwriting shop knows: restrict, remediate, or non-renew, with the incident logged as loss experience that reprices the rest of the book.




Exhibit 3 — Delegated authority for machines: promotion by audit and evidence, revocation on breach — written into the endorsement.

INTERLUDE · ONE FILE, END TO END

Decision Book entry DB-0217

FIELD	ENTRY
Decision class	Claims triage & FNOL routing (fast-track / adjuster / SIU referral), coastal homeowners book
Mode & volume	Bind with referral · ~9,000 decisions per day baseline; 30–40× during named-storm events
Owners	Business: Head of Claims Ops · Technical: ML Platform Lead · Challenge: second-line model risk
Dependencies (common factors)	Foundation model family FM-A (provider X) · property-attributes feed (vendor Y) · aerial-imagery scores (vendor Z)
Reversibility	Routing reversible with delay cost; wrongful fast-track of structural losses is expensive to unwind
BIND	B4 — LAE and cycle-time lever, named ratio line · I4 — offline replay + shadow harness · N3 — binds routing, refers edge cases · D4 — shapes claim outcomes, examiner-visible
Authority level	N3 × D4 → Level 2 (per the table in Move Three; I4 imposes no cap)
Ladder position	Bind-with-referral; promotion to bind-with-audit requires two clean file audits and a fresh fairness result

Endorsements attached at binding (excerpt, as enforced):

ID	CONDITION	ON BREACH
E-1 Referral band	Routing confidence < 0.90 refers to adjuster queue; expected referral ≤ 4% of volume, alert at 8%	Alert to owner; sustained 2h → on watch
E-2 Storm mode	Named-storm declaration in exposure states tightens thresholds; structural-loss routing reverts to adjusters	Automatic — mode entry is logged, exit requires duty-officer sign-off
E-3 Fairness & drift	Weekly disparity check on routing and referral rates across policyholder segments; input-drift statistic within band	On watch; re-validation within 14 days; ladder promotion frozen
E-4 Pull authority	Named duty-officer roster; full stop executable in ≤10 minutes; drilled quarterly, once in storm season	Missed drill suspends authority pending re-drill
E-5 Unit-cost band	Rolling per-decision inference cost within budget; daily spend at 2× budget trips the breaker	Safe mode; owner and finance paged

The **examiner file** for DB-0217 contains this entry, the four BIND writeups with their second-line challenge, gate evidence, the last shadow-mode comparison, twelve months of endorsement telemetry, one closed watch record, and the renewal sign-off. Producing it takes minutes, because the file assembles itself from artifacts the system produced in ordinary operation; the examiner is simply its second reader.

The same file, on a life book

The method is line-agnostic; only the perils change. A life carrier's highest-stakes entry reads like this:

FIELD	ENTRY (EXCERPT)
Decision class	Accelerated underwriting — routing term-life applications to the accelerated or full-evidence path
Mode & BIND	Bind with referral · B3 — placement-rate and cycle-time lever · I3 — retrospective mortality replay · N3 — binds the evidence path, refers impairment flags · D4 — affects insurability and price; examiner- and reinsurer-visible
Authority level	N3 × D4 → Level 2 · strict, per the Move Three table
Endorsements (excerpt)	Referral band on model confidence and impairment flags · disparate-impact check across the classes the carrier's filings and state law name · mortality-drift monitor tied to the experience-study cycle · program-pull condition: if the reinsurer suspends or amends the automated-underwriting treaty terms, the entry reverts to full evidence the same day
The "storm mode" here	A life book's named storm is a mortality shock or a reinsurer pulling program support — both written as event conditions before they occur, not negotiated during them

Risk assessment and pricing in life insurance is precisely what the EU AI Act names high-risk, and accelerated-underwriting programs are already on state examiners' and reinsurers' agendas. The Decision Book does not care which line wrote the entry.

Prepare for the event

After Andrew: modeled loss estimates, disaster scenarios and reinsurance programs turned "what if" into annual machinery. For the AI book: map the accumulations, name the version event, run the scenarios, and choose your net retention on purpose.

Accumulation: the peril-zone problem

Fifty Decision Book entries riding one foundation model, one vendor platform, or one retrieval index are, in underwriting terms, one peril zone with fifty exposures. A silent provider-side model update, an API outage, a poisoned data feed: any of these is an accumulation event across the book in a single hour. The accumulation response transfers directly: **map the common factors** (each entry records its model family, provider, retrieval and data dependencies), **set aggregate limits** (a ceiling on the share of Level 1 and 2 decision volume that may depend on any single model family or provider, with board visibility when the ceiling is approached), and **treat diversification as a design choice** — a second model family behind the most exposed decisions is the AI book's equivalent of not concentrating the whole portfolio in one coastal county. Diversification only works if the endorsement layer is model-agnostic: one policy manual per model is not a policy. Write the conditions once, at the decision level, and let the models underneath change.




Exhibit 4 — Accumulation, drawn: six Decision Book entries, one shared foundation model. Aggregates and limits apply, exactly as they would to a coastal zone.

Version events: the commonest accumulation shock, named

The most frequent portfolio event will not be an outage. It will be an upgrade. A provider ships a new model version; behavior shifts subtly across every dependent entry at once. In this framework that is a defined event class, not a surprise: a *version event* freezes authority promotions across dependent entries, accelerates re-validation on the highest levels, runs old-against-new shadow comparison where the ladder provides it, and logs the delta as loss experience even when nothing goes wrong. No reinsurer would accept a treaty whose subject business could be rewritten overnight, unannounced, by a third party. With shared models, that is what an unmanaged upgrade is.

When agents transact with agents

Agents are beginning to call agents — inside the carrier, and at the boundary where TPAs, MGAs and distribution partners automate their own side. The industry already knows how to govern decisions it has delegated outward: binding authorities, claims-handling agreements, audit rights. The same instinct applies inward. The entry that initiates a chain owns its downstream calls, may only invoke agents whose authority level meets what its own endorsements require, and records those dependencies in the book like any other exposure. Without that rule, a Level 1 decision can quietly outsource its judgment to something nobody scored. With it, a chain inherits the discipline of its strictest link.

Disaster scenarios for the AI book

Landfall plus six hours. First notices of loss arriving at forty times the Tuesday rate; the triage agent drifting outside everything it was tuned on — fast-tracking structural losses, holding the simple claims of families refreshing a portal in a shelter. In one carrier, that night is a bridge call at 5:15 a.m. and a market-conduct inquiry that outlasts the rebuilding. In another, a storm condition written into the agent's terms months earlier tightens its own routing at the named-storm declaration, and surge staffing absorbs the referred queue. You cannot write a storm plan during the storm. (Details are a composite, drawn from practice.)

The industry's crown jewel is scenario discipline: realistic disaster scenarios, run against the whole portfolio at once. The AI book deserves the same exercise, on an annual cadence and after any major reshaping of the book. Three scenarios earn their place in the first cycle:

Provider failure during an event. The dominant model provider is unreachable for six hours — at landfall plus six, when volumes are forty times baseline. Which decisions halt, which degrade to fallbacks, and what does the policyholder see?

Silent drift. An unannounced upstream change shifts behavior two weeks before anyone notices. Which endorsements catch it first, and how much decision volume flowed in the gap?

Adversarial campaign. Organized fraud probing the claims agents at scale — prompt injection arriving not as malware but as plain language in the documents, photos and correspondence the agents read. Where are the common seams, given that entries share components?

The exercise produces what it always has: a map of where the book is fragile, a modeled loss for the worst credible day, and a costed argument for buffers.

Expected loss is priced; unexpected loss is reinsured

The deepest habit the craft can lend its own AI program is the split it applies to everything else: expected loss is priced, unexpected loss is reserved or reinsured against. Expected loss belongs in the endorsement — an error budget priced into the business case, handling cost included. Unexpected loss is what buffers are for, and here they are concrete: surge-capable human capacity that can absorb a referred queue when an agent tightens its own routing; warm, tested rules-based fallbacks; drilled pull authority; and genuine risk transfer where it earns its cost — vendor contracts with the accountability and remedy provisions a claims-handling agreement would carry, and, increasingly, insurance on the AI itself. The carrier in the vignette that cleared its storm did so on surge staffing consciously kept warm — a net retention chosen months before the season. Andrew taught the industry that solvency lives at the portfolio level. So does this.

Institutionalize the cycle

After Andrew: the new discipline survived because it became annual machinery — renewal calendars, experience studies, appetite reviews. For the AI book: point the actuarial control cycle at the machines, measure the program itself, and make the whole thing run without heroes.

The actuarial control cycle, applied to the machines




Exhibit 5 — The industry's own feedback loop, pointed at the AI book: assumptions in, experience out, everything re-rated.

Insurance already owns the loop every governance program needs and most never build. Applied to the AI book it reads exactly as an actuary would expect. Set assumptions: scores, endorsements, authority levels. Write the business: decisions flow under those terms. Collect experience: incidents, drift, referral rates, loss runs. Study the experience — the discipline's oldest instrument, the experience study, becomes the backtest that keeps BIND honest: Level 1 entries should not be outperformed by Level 3, and if they are, the anchors move. Then re-rate: scores, terms, authority and appetite adjust, annually and on trigger. On day one, BIND is expert judgment with anchors. By the second cycle it has a track record to answer to — which is more than most governance frameworks can ever say, and exactly what an actuary demands of any assumption.

Measuring the program, not just the machines

Six numbers tell you whether the method is working: time from submission to score (days, not months); incidents per thousand decisions, by authority level — the experience study above; endorsement-breach rate and time-to-detection; time on watch before resolution; wrongful-referral and wrongful-fast-track rates on claims entries; and terms retired per year, because a regime that only accumulates is decaying differently. Non-renewal is a governance outcome too: a model that no longer earns its place on the book should leave it at renewal, the way any risk does.

Where the machine bolts on

No method survives contact with an org chart it ignored, so four seams deserve plain answers. **Ownership:** the first line owns each entry (a named business owner, per the interlude); the second line owns the scoring standards, endorsements and challenge; internal audit receives the trail. The Book is chartered under the CRO's office as a system of record that *references* existing inventories — the model list, the filing index — rather than re-keying them. **Actuarial and model risk:** BIND scores decisions; the actuarial function keeps models and filings; the two meet at gate G4, where existing validators do what they already do, extended with decision-level scenarios. Nothing displaces the appointed actuary's perimeter. **Underwriting and claims committees:** gates attach to the approval paths the carrier already runs — an AI-touched proposal carries its G1 score into the existing committee, not into a second one. **Escalation:** sponsors will sometimes go over the framework's head; pretending otherwise is how frameworks die. The override is therefore formal: a risk acceptance taken by name, time-boxed, logged on the entry, reported to the risk committee. The override is visible by design; the one thing it can never be is quiet.

The first ninety days

For a mid-sized book — larger carriers phase this by line of business — weeks 1 and 2 are the back-book review: Decision Book v1 including vendor and platform-embedded AI, priority flows scored, common factors mapped, the gap published. Weeks 3 through 6 produce the frameworks: anchor ladders, endorsement templates by level, gate criteria, aggregate limits, and a committee charter co-chaired by risk and technology. Weeks 7 through 9 secure the mandate: the board risk-committee pack and approval — which means anchoring the clock to the board calendar, launching off the back of one quarterly meeting so the mandate lands at the next. Enforcement follows, highest levels first, then the first scenario cycle before storm season, not after it. Tooling gets selected last, once there is a design for it to run; a platform cannot inventory decisions it cannot see, and it cannot enforce endorsements nobody has written — pick the tool first and you have bought someone else's appetite.

Where this breaks, and what is designed against it

FAILURE MODE	THE COUNTERMEASURE, BUILT IN
Authority creep. Every sponsor argues for a bigger pen	Levels derive from the N and D anchors and a second-line challenge; duty-of-care definitions reference claims-outcome and market-conduct reach, which are not matters of opinion
Endorsement-lite drift. Conditions exist on paper only	A term without an enforcement hook does not pass gate G5; drills are terms too
Accumulation blindness. Entries fine alone, fragile together	Common-factor mapping, aggregate limits, version events, and the annual scenario exercise
Storm-season amnesia. Event modes designed after the event	Storm conditions are written into endorsements before the season and drilled once during it
Operator fatigue. The duty officer drowns in alerts	Alert budgets are endorsement terms; referral-share ceilings cap what the queue may demand of humans; drills test the human path, not just the switch
Sponsor escalation. The business outranks the duty officer	Overrides are formal risk acceptances — named, time-boxed, logged on the entry, reported to the risk committee
Committee theatre. The council becomes a status meeting	A decisions-only agenda: bindings, term changes, watch actions. Status lives in dashboards
The stale book. Reorgs orphan entries	Ownerless entries auto-flag for review; the book is re-audited annually on the ORSA cadence the carrier already keeps

Every method has failure modes. These are ours, and the countermeasures ship with the design.

Where the parallel ends

The parallel has limits, and naming them matters. Three places 1992 does not map cleanly. The industry calibrated cat models against physics and decades of loss data; BIND starts as anchored judgment and earns its record through the experience studies above. A hurricane is indifferent to the book it strikes; an optimizing agent is not — it finds the seams in its own terms, which makes enforcement easier and anticipation harder. And reinsurance for AI failure is young; the buffers here — surge capacity, warm fallbacks, drilled authority — are real, but no actuarially priced market stands behind them yet. Against those concessions, one difference runs the carrier's way, and it did not exist in 1992: a policy's conditions are, for the most part, tested when a claim arrives. A machine endorsement is tested on every decision it touches. Carriers have spent three centuries writing terms that others must keep. These are the first terms they have written that keep themselves.

Ten questions before your next risk committee

Score each 0 (no), 1 (partially), or 2 (yes, with evidence).

☐	Could you produce, within one day, a list of every consequential decision a machine touches — including vendor systems, platform-embedded AI, and the agents teams built for themselves?
☐	Does that list carry business value alongside risk, with a named owner per entry?
☐	Is there a written, anchored scale that sets authority and scrutiny, or is it negotiated case by case?
☐	For your most autonomous system: is the band outside which it must refer to a person written down and enforced in code?
☐	Does someone hold pull authority over that system, in writing, drilled in the last quarter?
☐	Do your claims and underwriting agents carry a written storm mode — designed before the season, not after the event?
☐	Do you know which of your machine decisions share a single model family, provider or data feed, and is there a limit on that accumulation?
☐	Would your claims-facing AI pass an independent fairness challenge on adverse outcomes this quarter?
☐	Has the whole AI book ever been run against one event at once — a provider outage at landfall, a silent model update?
☐	If your two most senior AI people left tomorrow, would the governance still run?

0–7

Pre-Andrew.

Writing the coast without aggregates. Start the back-book review this week.

8–14

Mid-transition.

The entry-level pieces exist; the accumulation layer is missing.

15–20

Leading.

Generalize what works — usually from CAT and SIU operations — across the book.

The carriers that come out of the agentic era ahead will be the ones whose governance lets them delegate the most authority to machines — safely, with evidence, sooner than their peers. That is underwriting, and no industry has practiced it longer.

Authors

Arun Chinnachamy · Venkat Chitturi

About Technoidentity

Technoidentity is a practice leader in Data+AI, helping regulated enterprises take AI from pilot to governed production. We build **Durable Agents** — agentic systems engineered to live in production under real terms — and we practice **Durable Product Engineering**: software built to survive scale, audit and time. The frameworks in this paper are how those two practices meet the carrier: agents worth handing a pen to, and the governance that makes the handover safe. Conversation: technoidentity.com

Hurricane Andrew (1992) produced what was then the largest insured catastrophe loss in history; industry accounts commonly cite eleven insurer insolvencies in its aftermath, and it is widely credited with reshaping catastrophe modeling and Florida's insurance market; Florida's Commission on Hurricane Loss Projection Methodology (est. 1995) reviews and certifies hurricane models used in residential ratemaking. Regulation (EU) 2024/1689 (AI Act), Annex III, names risk assessment and pricing in life and health insurance among high-risk uses; obligations phase in through 2026–27. The NAIC model bulletin on insurers' use of AI (December 2023) had been adopted by 24+ US states as of 2026. The Move Four vignette is a composite drawn from practice, with details altered. This paper is general information, not legal advice. © 2026 Technoidentity. All rights reserved.