

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

THE REVEAL

A Premortem of the Collapse of
Complex Incident Management

JAYSON COIL

The Reveal

COGPLEXITI · SYSTEMS & DECISION ANALYSIS
PREMORTEM ANALYSIS

The Reveal

A premortem of the collapse of Complex Incident Management

JAYSON COIL

Cogplexiti LLC · Flagstaff, Arizona

July 2026

COPYRIGHT AND LICENSE

© 2026 Jayson Coil. All rights reserved.

ISBN 979-8-9959056-8-4

Published by Cogplexiti LLC.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means without the prior written permission of the publisher, except for brief quotations used in reviews and other uses permitted by copyright law.

Use for Artificial Intelligence and Machine Learning. This work may not be used to train or develop artificial intelligence or machine learning systems, or be included in any dataset used for such purposes, without the prior written permission of the publisher.

Permissions. Direct requests to Cogplexiti LLC at permissions@cogplexiti.com.

EDITION INFORMATION

July 2026.

Published by Cogplexiti LLC, Flagstaff, Arizona · jayson@cogplexiti.com · www.cogplexiti.com

Part of the Cogplexiti Analysis line (Systems & Decision Analysis), alongside The System Tier and The Incident Tier. This premortem has a companion discussion paper, Sustaining Decision Performance, which isolates the individual-level certification gap.

DISCLAIMER

The views, frameworks, and recommendations expressed in this publication are those of the author and do not necessarily reflect the official policy or position of the Sedona Fire District, Southwest Incident Management Team 1, the National Wildfire Coordinating Group, the U.S. Forest Service, or any other agency or organization. This document is not doctrine. It is a professional analysis informed by operational experience, research, and study.

The 2031 collapse date is a model output, not a prediction; see How to Read This Paper.

How to Read This Paper

This paper belongs to the Cogplexiti Analysis line — systems- and incident-level studies that sit alongside The System Tier and The Incident Tier. It is a premortem, in the sense Gary Klein formalized. It stipulates that the interagency Complex Incident Management Team system has already failed — shuttered in 2031 — and works backward to explain why. The date is not a prediction. It is the central estimate of a workforce stock-and-flow model, and it moves when the model's parameters move. The stipulation exists to remove the social cost of describing failure, so the failure can be described accurately. Stipulated history is marked as stipulated throughout, and the data program in Section 5.2 identifies which posited facts the enterprise's own records can confirm or refute.

The paper is built in the Cogplexiti house analytical format — Executive Summary, Support & Analysis, Red Team Assessment, Dissenting Alternative, and Recommended Course of Action. Its argument rests on established theory developed in other high-consequence domains, carried in the annotated Theory Anchors that close the volume. The register is deliberately forceful — a premortem is written to make failure describable — but the standard of evidence does not change with the register: calibrate language to the evidence, distinguish evidence from opinion, and apply local rationality symmetrically. Explain, don't judge.

This premortem has a companion discussion paper. Sustaining Decision Performance isolates the individual-level version of this paper's first finding — the qualified individual is never re-evaluated — in a measured register prepared for interagency review, anchored to primary regulatory sources. The two papers stand independently. Read together, the companion supplies the concrete, verifiable instance; this paper supplies what that instance becomes at enterprise scale, and the mechanism-level answer to the question the companion deliberately leaves open.

Contents

How to Read This Paper	v
1 · Executive Summary	1
2 · Support & Analysis	3
2.1 · Method: The Premortem and the Incubation Frame	3
2.2 · The Incubation Record: Symptoms Crosswalked to Established Theory	4
2.3 · The Hoarding Paradox: Multi-Echelon Root Causes	7
2.4 · The Pipeline Arithmetic: Deriving the Collapse Date	8
2.5 · Core Mechanisms of Failure	11
2.6 · The HRO Gap: The Claim Versus the Criteria	12
3 · Red Team Assessment	14
3.1 · Biases and Weaknesses in This Analysis	14
3.2 · Low-Probability, High-Impact Risks	15
4 · Dissenting Alternative	16
5 · Recommended Course of Action	17
5.1 · Name the Method, Publish the Model	17
5.2 · Close the Data Gaps First	17
5.3 · Restore the Signal	18
5.4 · Govern Polycentrically, Not by Leviathan	18
5.5 · Architect the Successor as a Legitimate HRO	18
5.6 · Implementation Considerations	19
Theory Anchors	20

1 · Executive Summary

The degradation and structural failure of the national Complex Incident Management Team enterprise is not a novel, domain-specific crisis. It is a predictable manifestation of organizational and cognitive failure modes that have been documented, formalized, and in several cases mathematically modeled across complex, high-consequence industries. Barry Turner's disaster incubation model supplies the frame this paper is built on:

organizational disasters incubate over years, with warning signals present, visible, and systematically discounted, until a precipitating event makes the failure undeniable. By that standard, the collapse stipulated here for 2031 is not the failure — it is the reveal. The failure occurred years earlier, arguably by 2025 or 2026, while the enterprise's own health metrics continued to report success.

The method is a premortem, in the sense Gary Klein formalized: rather than asking what could go wrong, the analysis stipulates that the system has already failed and asks why. Research on prospective hindsight indicates this framing materially improves the identification of failure paths, because it removes the social cost of predicting failure. No reader is asked to agree that the system will collapse; they are asked only to explain a stipulated collapse. That is the same permission structure a well-run after-action review creates, applied in advance.

The analysis reaches five principal findings. First, the qualification system has suffered a signaling failure: credentials that once carried information about capability have decoupled from it, so the system can no longer distinguish the qualified from the credentialed — a failure that was visible before the CIM consolidation in the divergence between team type and team performance, and one the consolidation entrenched rather than repaired. (The companion paper, *Sustaining Decision Performance*, documents this failure at the level of the individual position-holder.) Second, the workforce pipeline is arithmetically insolvent — the time required to produce a fully qualified Operations Section Chief or Incident Commander now exceeds the remaining service life of the cadre expected to produce them, and inflow is below attrition. The 2031 date is derived from this arithmetic, not chosen for effect. Third, resource hoarding is not one behavior but three, with distinct game-theoretic structures at the state, geographic-area, and incident echelons; interventions that misdiagnose the game are guaranteed to fail. Fourth, the enterprise's health metrics have been captured by the phenomenon Goodhart's Law describes — team counts, roster minimums, and preparedness levels became targets, ceased to measure capability, and now actively mask its depletion. Fifth, the enterprise routinely describes itself as a high reliability organization while meeting none of the five criteria the high-reliability literature actually specifies. The gap between the claimed identity and the audited state is itself a warning signal of the kind Turner's model predicts will be discounted.

The recommended course of action follows from the diagnosis rather than from culture: publish the workforce model and invite refutation; close the data gaps that simultaneously parameterize the model, adjudicate the dissent, and solve the geographic-area trust problem; restore the informational value of qualification through forced-failure simulation vetting; reject the intuitive central-authority demobilization fix in favor of polycentric governance, which the empirical record favors; and architect the successor system explicitly against the five high-reliability criteria — that is, build the organization the current one already claims to be. Above all: stop treating these failure modes as special. They are not. Treating them as special is what has made every intervention to date a half-measure, and half-measures consume the one resource the arithmetic says is running out — time.

THE STIPULATION — MARCH 2031

It is March 2031. The National Multi-Agency Coordinating Group has suspended the national Complex Incident Management Team rotation. The final season, 2030, opened with fewer than half of rostered CIMTs able to fill Command and General Staff positions without waivers; three teams declined national assignment outright for lack of a qualified Operations Section Chief. A destructive wildland-urban interface fire in the 2029 season — contained late, at extraordinary cost, after days of tactical stagnation later characterized as command paralysis — became the precipitating event. In the aftermath, four states formalized compact-first mobilization policies, formally embargoing state and municipal resources from national dispatch during declared state emergencies. The interagency roster model was not repaired; it was replaced by a pilot program of regional standing task forces.

Posited history. Every item in this block is stipulated for premortem purposes. Section 2.4 identifies which elements are outputs of the workforce model and which are placeholders pending the data program in Section 5.2. The stipulation is falsifiable by design.

2 · Support & Analysis

2.1 · Method: The Premortem and the Incubation Frame

Two established tools structure this analysis, and both come from outside the fire service — which is deliberate.

The premortem. Klein’s premortem technique inverts the usual planning question. Instead of asking a team to critique a plan — which invites defensiveness, because critiquing the plan means critiquing its owners — it stipulates that the plan has already failed and asks each participant to explain why. Research on prospective hindsight, the underlying mechanism, found that assuming an outcome has already occurred substantially improves people’s ability to identify reasons for it. The stipulated 2031 collapse performs the same function for an entire enterprise: it makes failure discussable without requiring anyone to defend a prediction, and without requiring anyone in the current system to be wrong.

The incubation model. Turner’s analysis of man-made disasters, later extended with Pidgeon, found that large organizational failures share a common structure: a long incubation period in which discrepant events accumulate but are discounted, misinterpreted, or absorbed by ‘decoy’ problems that attract attention away from the developing failure; a precipitating event; and only then, cultural readjustment. The critical implication for this paper is temporal. The disaster date is when the failure becomes undeniable, not when it occurs. If the CIMT enterprise collapses in 2031, Turner’s model predicts the enterprise was already insolvent years earlier — and predicts, further, that the intervening years will have been spent processing waivers, adjusting roster definitions, and running studies: administrative absorption of what were in fact terminal signals.

The date. The 2031 collapse date is not rhetorical. It is the central estimate of the stock-and-flow model in Section 2.4, which computes the year the stock of fully qualified command personnel crosses below a minimum viable national rotation under current inflow and attrition. The estimate is deliberately published with its parameters so it can be attacked, tightened, or refuted with the enterprise’s own records. If better actuarial data moves the date to 2034 or 2029, the model has done its job either way. What the model does not permit is the current default assumption — that the date is comfortably far away — because no plausible parameterization produces that result.

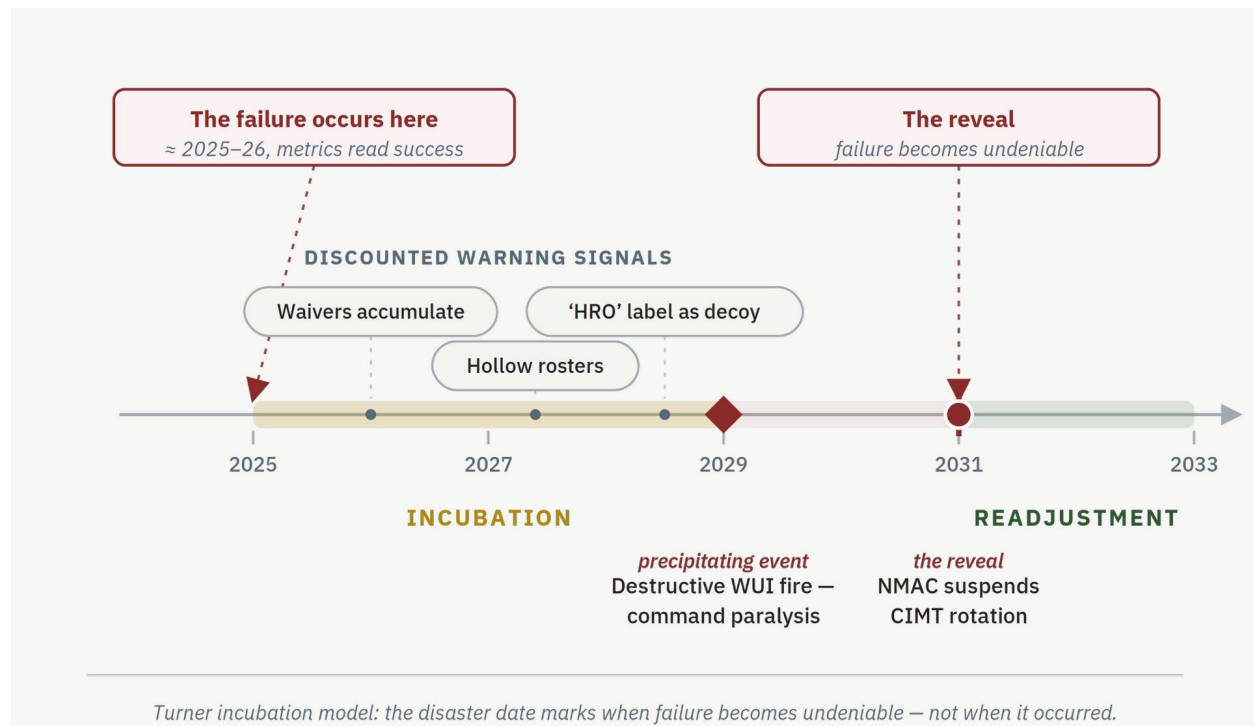


Figure 1 • The failure incubates years before it becomes visible

2.2 • The Incubation Record: Symptoms Crosswalked to Established Theory

Table 1 maps the observed symptoms of enterprise degradation to the established theoretical constructs that describe them. Each of these constructs was developed, tested, and in most cases formalized in another industry. None was developed for wildland fire, and none requires modification to apply to it. That is the point: these are not analogies, they are instances.

TABLE 1

Crosswalk of core systemic symptoms to established theory

Observed symptom (the 'band-aid')	Established theory	Cross-disciplinary definition
Merging Type 1 and Type 2 into 'Complex' (CIM) as a deliberate policy act, restoring roster legibility without increasing simulation rigor or capability validation.	Migration toward the boundary / practical drift Rasmussen (1997); Snook (2000)	Under sustained resource and efficiency pressure, work systems migrate incrementally toward the boundary of acceptable performance. Each adaptation is locally rational and individually defensible; collectively they consume the safety margin. The merger was one large, visible step in a migration made of many small ones.
Accumulation of position waivers, trainee-heavy	Normalization of deviance	The gradual process by which anomalies are reclassified as acceptable because they have not

rosters, and relaxed currency standards — each individually justified, none individually alarming.	Vaughan (1996)	yet produced catastrophe. The standard does not fall in one decision; it is renegotiated one waiver at a time, and each successful season is read as evidence the new standard is safe.
Operators building heuristic mental models (‘sets and reps’) on low-complexity, quasi-linear fires, then exhibiting command paralysis on non-linear megafires.	Negative transfer and the boundary conditions of recognition-primed decision-making Dörner (1996); Klein (1998); Woods	Schemas built in linear environments fail — sometimes catastrophically — when transferred to dynamic, non-linear ones. Dörner demonstrated this experimentally: capable operators managing simulated complex systems defaulted to over-control, thematic vagabonding, or freezing. Klein’s own work specifies that recognition-primed decision-making is valid only where the decision-maker’s experience base matches the environment; the megafire is precisely where it does not.
Promoting personnel to CIM command roles on time-in-grade and administrative taskbook sign-offs rather than validated capability.	Signaling theory Spence (1973)	Credentials are signals of capability, not measurements of it. A signal carries information only while its acquisition cost is coupled to the underlying attribute. When sign-offs can be accumulated administratively, the cost decouples, the signal degrades, and selection on the signal fails. (The Peter Principle is the folk version of this claim; signaling theory is the formalized, Nobel-recognized version, and it specifies exactly what the fix must do: re-couple cost to capability.)
Reliance on an aging cadre of Administratively Determined retirees to fill critical Command and General Staff roles.	Stock-and-flow insolvency Forrester; Meadows	When the production latency of a qualified asset exceeds the remaining service life of the existing stock, and inflow runs below attrition, depletion is not a risk — it is arithmetic with a date. Section 2.4 builds this model explicitly.
The assumption that unmanaged complexity can be absorbed by adding administrative tools, checklists, and standardized process.	Declining marginal returns on complexity; brittleness in stretched systems Tainter (1988); Woods	Societies and organizations solve problems by adding complexity, and each added layer costs more and returns less until simplification — voluntary or forced — becomes the economizing option. Systems run permanently at capacity do not degrade gracefully; they are brittle, appearing functional until they fail suddenly.
Gauging enterprise health by team counts, roster minimums, and preparedness levels.	Goodhart’s Law / Campbell’s Law	When a measure becomes a target, it ceases to be a good measure. Institutions under pressure defend the metric, not the capability the metric was meant to track — including by redefining what counts as a ‘full’ team. Bureaucratic systems rarely report their own demise; they redefine success.

The merger’s evidentiary basis — and the inference error. Local rationality applies to the consolidation’s advocates as much as to anyone else in this record, because they were not merely gaming rosters — they had genuine evidence in hand. Prior to the merger, it was easy to find incidents assigned to Type 2 teams that a complexity analysis

would have ranked as Type 1, and in high-occurrence geographic areas — the Southwest and Pacific Northwest among them — Type 2 teams routinely accumulated more complex-incident exposure than their designation implied. At the same time, some Type 1 teams performed at mediocre levels. The type labels demonstrably failed to track capability: the variance in performance within each type exceeded the difference between the types that the labels asserted. That is a real finding, and it points at a real root cause — the enterprise never built and equally applied an objective competence standard, so team type was always a proxy (for organizational pedigree, roster composition on paper, and qualification-card counts) rather than a measurement. Note also what the exposure pattern implies: capability tracked complex-incident reps, which tracked geography and fire occurrence — not designation. The signal the enterprise relied on was broken before the merger touched it; the team-type divergence was the qualification system’s signaling failure showing up one echelon higher.

The inference error is what was done with the finding. From ‘the signal is broken,’ two conclusions were available: repair the signal — build the competence standard and measure every team against it — or eliminate the signal by collapsing the categories. The enterprise chose elimination, and the case for it was assembled by selection on the favorable tail: the high-performing Type 2 teams were championed as proof the distinction was meaningless, while the remainder of the Type 2 distribution — the teams performing at or below their designation — was left out of the argument. Selecting the tail and generalizing to the population is a named methodological failure, and it produced a predictable result: the consolidation removed the between-type distinction without touching the variance the evidence had actually revealed. A single CIM label now spans the entire capability distribution that two labels spanned before, so the post-merger signal carries strictly less information than the broken one it replaced, and the unmeasured variance rides along undisturbed. The merger, in short, responded to evidence of a measurement failure by abolishing the measurement — which is why Section 5.3’s remedy is a capability standard applied to every team, not a re-drawing of category lines.

A note on Normal Accident Theory. Perrow belongs in this inventory, but only if the theory is pointed at the right object — and the common application points it at the wrong one. Perrow’s claim is that in interactively complex, tightly coupled systems, catastrophic accidents are inevitable regardless of organizational quality — a claim that sits in direct tension with any set of recommended organizational fixes. An analysis that invokes Normal Accident Theory as a general frame and then prescribes organizational reform is arguing against itself. The tension resolves once the theory is applied to the right object. The megafire incident environment is Perrowian: interactively complex, tightly coupled, capable of cascading faster than any command structure can track. The national mobilization enterprise is not — it is loosely coupled and degrades slowly, which is why its failure incubates rather than explodes, and why Woods’ concept of brittleness under

stretch describes it better than Perrow does. Perrow, properly placed, describes the problem environment the successor system must operate in. The competing school — high reliability organization theory — describes what that successor system must be. Section 2.6 develops this, and it turns out the enterprise has already conceded the point in its own rhetoric.

2.3 · The Hoarding Paradox: Multi-Echelon Root Causes

Resource hoarding produces an identical operational signature at every echelon — orders returned Unable to Fill while capable resources sit held elsewhere. The identical signature is exactly what makes it dangerous to diagnose, because the drivers are structurally different at each level, and each driver belongs to a different, well-characterized class of problem with a different solution set. Two of the three echelons are routinely misdiagnosed — state withholding gets labeled a tragedy of the commons, and geographic-area behavior gets labeled a Prisoner's Dilemma. Neither label survives contact with the actual incentive structure, and the distinction is not pedantic; it determines what will actually work.

TABLE 2

The hoarding paradox — multi-echelon root causes

Echelon	Operational execution	Governing structure	Mechanism and solution class
State / municipal (e.g., Cal OES, AZ DFFM)	Embargoing state and local government resources from national dispatch; prioritizing intra-state and compact responses.	Free-rider problem / collective action failure — Olson (1965).	Hardin's commons is depleted by overconsumption of a shared resource. What states do is the mirror image: under-contribution to a shared capability. National surge capacity is a public good — non-excludable — so a state that withholds contribution still receives protection when its turn comes. Defection is individually costless and politically rewarded. Solution class: contribution mechanisms — conditional funding, reciprocity accounting, exclusion from benefits — not appeals to stewardship of a 'commons' that is not being consumed but starved.
Geographic area (GACC) (e.g., SWCC)	Declining to release assets to NICC when local fire danger is elevated but actual fire activity is low.	Stag Hunt (assurance game)	In a Prisoner's Dilemma, defection dominates no matter what others do, and only enforcement fixes it. A GACC's situation is different: it would prefer mutual release if it trusted reciprocity. Cooperation is the preferred equilibrium — it is merely fragile under uncertainty about what other GACCs are holding. That is an assurance game, and assurance games are solved with credible information, not coercion: transparent, real-time, enterprise-wide visibility of every GACC's

			held-resource and UTF ratios converts the game. The data program in Section 5.2 is therefore not just measurement — it is the intervention.
Incident (IC / C&G staff) single-incident scope	Padding resource orders, holding unassigned resources in staging, aggressively delaying demobilization.	Principal-agent problem, amplified by loss aversion — Jensen & Meckling; Kahneman & Tversky (1979).	The IC (agent) optimizes for their incident; the enterprise (principal) bears the cost, and the accountability structure is asymmetric — intense scrutiny if an incident escapes for want of resources, zero institutional penalty for containing a fire with thirty percent more resources than tactically necessary. Loss aversion then does the rest: the vivid, attributable risk of releasing a crew outweighs the invisible, unattributable benefit of that crew saving a structure elsewhere. Solution class: realign the incentive structure and make utilization observable — not exhortation, which asks agents to act against their own rational interest.

The taxonomy carries one further implication, and it determines intervention design. Because the three echelons are playing three different games, no single policy instrument can address ‘hoarding.’ A transparency regime that solves the GACC assurance problem does nothing about state free-riding; a contribution mechanism that fixes state incentives does nothing about incident-level loss aversion. An enterprise that continues to treat hoarding as one behavior — usually as a character defect — will continue to produce interventions that fail at two echelons for every one they touch, and the failures will be read as evidence that ‘nothing works,’ deepening the fatalism that already pervades the system.

2.4 · The Pipeline Arithmetic: Deriving the Collapse Date

The phrase ‘demographic cliff’ is a metaphor, and metaphors are contestable. A stock-and-flow model is arithmetic, and arithmetic is not. The model has four elements, all of them estimable from records the enterprise already holds.

The stock. The number of fully qualified, currently available CIM-level Operations Section Chiefs and Incident Commanders in the national rotation — qualified without waiver, physically available for assignment, and actually accepting rotations. This is the quantity the roster metrics are supposed to track and, per Table 1’s final row, no longer do: a roster count that includes waived positions and non-mobilizing names overstates the stock by a margin the enterprise has never measured.

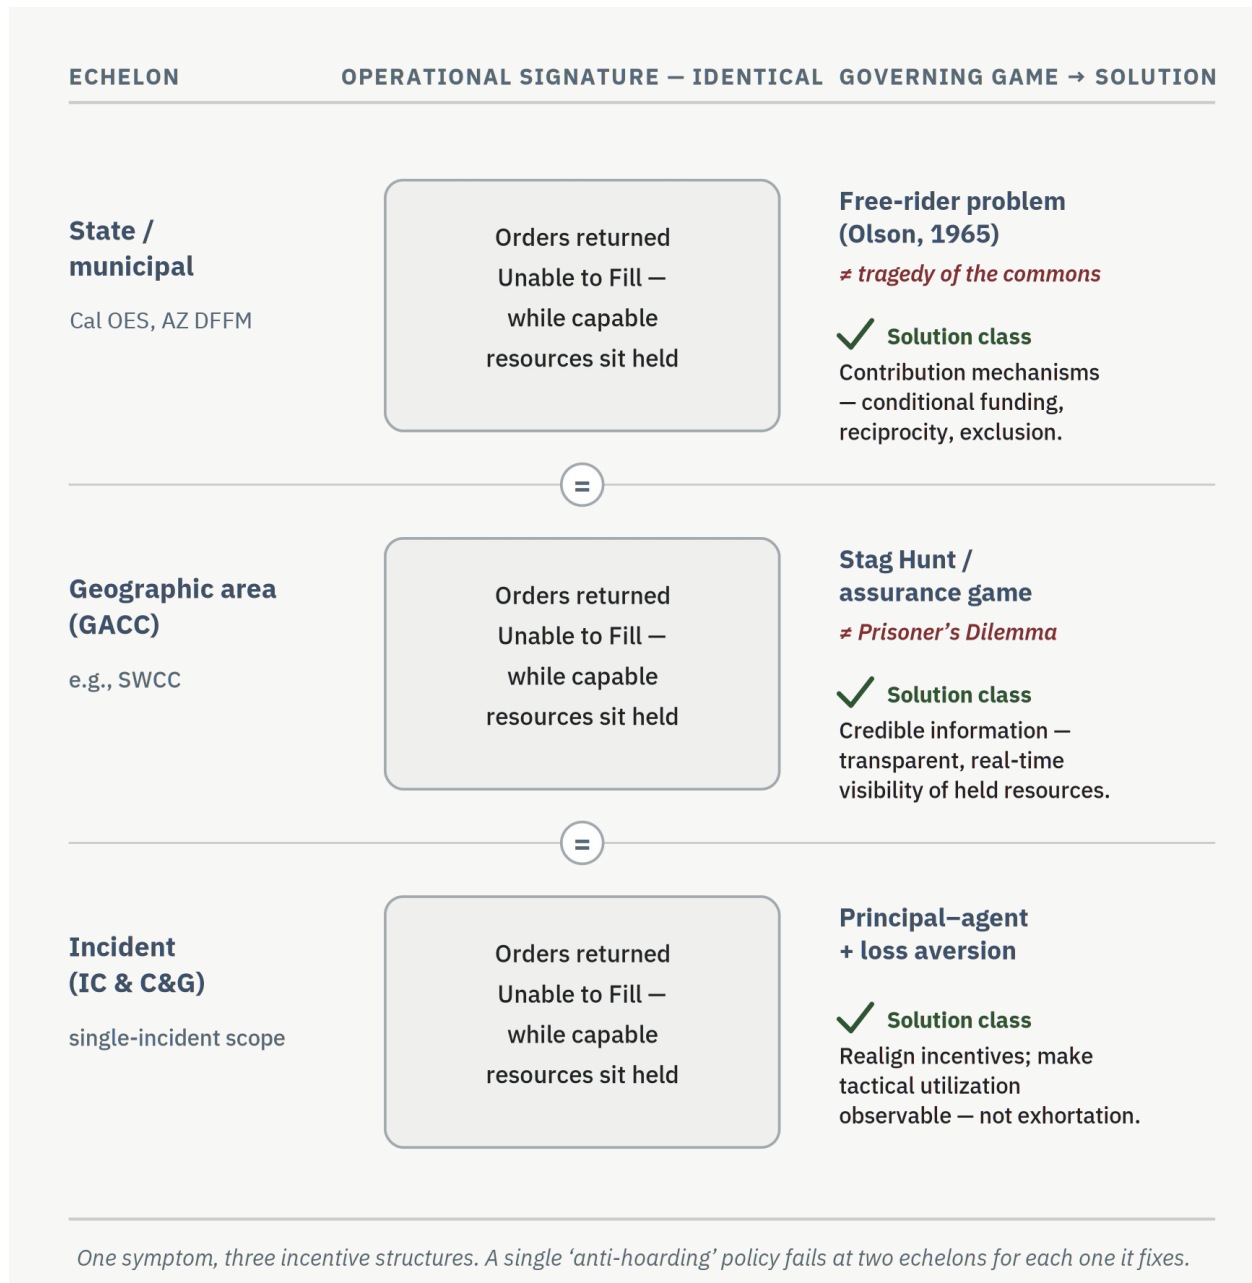



Figure 2 · One signature, three games, three solution classes

The outflow. Attrition, dominated by retirement. The cadre currently carrying the load is heavily weighted toward Administratively Determined retirees whose remaining participation horizon is plausibly five to eight years — a parameter the actuarial audit in Section 5.2 will replace with a measured distribution. Every year of delay converts an estimate into a smaller measured number.

The inflow. Qualification production. Producing a no-kidding CIM-qualified OSC or IC requires fifteen or more years of progressively complex assignments, and the experience economics have worsened in a compounding way: the complex assignments that produce

qualification increasingly depend on the presence of the same senior cadre that is leaving. The system is consuming its instructors faster than it graduates their replacements.


The threshold. A minimum viable national rotation — the smallest stock that can sustain coverage through a severe season without cannibalizing recovery time. The enterprise has never formally defined this number, which is itself a finding: an organization that cannot state its minimum viable capability cannot know its distance from it.



Read: while inflow lags outflow, the fill falls toward the red line — see Figure 4.

Figure 3 · The qualified-cadre pipeline as a stock and flow

Under any plausible parameterization, inflow is below outflow now — the disagreement is only about the crossing date. The central estimate, using mid-range assumptions for each parameter, puts the stock below a defensible minimum viable rotation around 2031, with a sensitivity band of roughly 2029 to 2034. Two properties of this result matter more than the specific year. First, the model computes insolvency, not collapse: collapse is how insolvency is discovered, and Turner’s model plus the ghost-phase mechanism in Section 2.5 explain why discovery lags reality — waivers and hollow rosters keep the metric above zero after the capability has crossed it. Second, the model is falsifiable with the enterprise’s own data, and Section 5.2 specifies exactly which records do it. A reader who believes the date is wrong is invited to say which parameter is wrong, and by how much. That is a different kind of argument than the enterprise is used to having, and it is the kind this paper exists to force.



Axis values are model output on estimated parameters until the Section 5.2 actuarial audit; the curve shows structure, not a forecast to a specific year.

Figure 4 · Projected qualified stock against minimum viable rotation

2.5 · Core Mechanisms of Failure

Three mechanisms convert the structural conditions above into operational failure. Each is a named phenomenon; none is unique to fire.

The capability illusion, restated as signaling failure. The current system confers authority without conducting objective, stress-inoculated capability validation — it certifies the signal, not the attribute. The operator is frequently the first to sense the gap: a commander who holds the qualification but not the validated capability behind it recognizes the deficit at some level, and that recognition manifests as extreme operational risk aversion. This is the mechanism behind the ‘safe catastrophe’ scenario in Section 3.2 — tactical stagnation presented, sincerely, as risk management. The remedy is specified by the theory that names the problem: restore the coupling between the cost of the credential and the capability it claims (Section 5.3).

Decoupled accountability in logistics. There is intense scrutiny on an Incident Commander if an incident escapes containment for want of resources, and zero institutional penalty for containing a fire using thirty percent more resources than tactically necessary. The asymmetry makes gross over-ordering the rational strategy for every individual agent, in the same move that starves the enterprise. This is the incident-

level principal-agent structure of Table 2 expressed as an accountability design flaw — and note that it persists partly because utilization is unmeasured, which is a Goodhart problem: the enterprise tracks what was ordered and filled, not what was used.

The systemic ghost phase. The mobilization system currently masks insolvency by fielding functionally hollow forces — high trainee ratios, extreme waiver usage, rosters that are full on paper and thin on the line. In Woods’ terms this is brittleness under stretch: a system running permanently at capacity, with its margin consumed, that continues to look functional right up until it does not. The ghost phase is the specific mechanism by which Turner’s incubation is sustained — it is how an insolvent enterprise keeps generating metrics that read as solvent, and it is why the collapse date in the stipulation arrives as a surprise to the metric-readers and as a long-expected confirmation to the field.

2.6 · The HRO Gap: The Claim Versus the Criteria

The wildland fire enterprise routinely describes itself as a high reliability organization. The term appears in doctrine, in training curricula, and in leadership rhetoric as a settled identity. This creates an unusual analytical opportunity: the HRO literature — principally Weick and Sutcliffe, building on La Porte and the Berkeley school — does not offer a label, it offers five operational criteria. An organization either exhibits them or it does not. The enterprise can therefore be audited against a standard it has already accepted, and the audit requires no imported values whatsoever.

TABLE 3

HRO criteria audit — the claimed identity versus the operating state

HRO principle (Weick & Sutcliffe)	What it requires	The enterprise’s audited state
1. Preoccupation with failure	Near-misses, anomalies, and small discrepancies are treated as data about the system’s health and are actively hunted. Success is treated as a reason for suspicion, not reassurance.	Reporting channels exist (SAFENET, RLS), but the metric layer above them is Goodhart-captured: waivers are processed as administration rather than read as anomalies; the redefinition of a ‘full’ CIMT roster is the precise inverse of this principle — a preoccupation with the appearance of success. An organization preoccupied with failure would treat every waiver as a data point in the stock-and-flow model. This one treats each as paperwork to be ignored.
2. Reluctance to simplify	Resistance to single narratives; deliberate cultivation of diverse interpretations; skepticism toward received explanations.	The enterprise’s default diagnosis for every symptom is one narrative — more training, more checklists, more standardized process — applied universally. Per Tainter, it is also an increasingly expensive narrative. The multi-echelon hoarding taxonomy in Table 2 is exactly the kind of

		differentiated diagnosis a simplification-reluctant organization would have produced internally; it did not.
3. Sensitivity to operations	A continuously updated, real-time picture of what the front line is actually doing — the ‘bubble’ — held at the organizational level, not just the incident level.	The enterprise cannot currently state what its committed resources are doing during a Preparedness Level 5 draw — how many operational shifts are spent in staging or low-priority mop-up, or operations that have no chance at success? What its true tactical utilization rate is. Its situational awareness of its own posture is aggregate and lagged. The data gaps in Section 5.2 are, item for item, a sensitivity-to-operations deficit list.
4. Commitment to resilience	Investment in slack, redundancy, and recovery capacity; the ability to absorb a hit and reorganize. Resilience is budgeted, not hoped for.	The enterprise runs structurally hollow: high trainee ratios, waiver-sustained rosters, and a surge model dependent on a retiree cadre it does not control. This is not merely an absence of slack — the AD dependency is anti-slack, a load-bearing element with a known expiration and no replacement in production. Brittleness is not a deviation from the design; it currently is the design.
5. Deference to expertise	In anomaly and crisis, decisions migrate to the person with the most relevant expertise regardless of rank; in calm, hierarchy resumes. Rank and expertise must therefore remain coupled enough that the migration is possible.	The signaling failure in Table 1 severs the precondition: when qualification is acquired by time-in-grade and administrative sign-off, rank no longer reliably locates expertise, so there is nowhere for decisions to defer to. Time-in-grade promotion is not merely non-HRO; it is structural anti-deference.

The audit supports two conclusions, one diagnostic and one architectural. Diagnostically, the gap between the claimed identity and the audited state is itself a Turner-class warning signal — arguably the most dangerous one in the inventory, because an organization that believes it already is what it needs to become has no reason to build it. The HRO label functions as a decoy phenomenon in Turner’s exact sense: it absorbs the anxiety the underlying signals should generate. Architecturally, the audit hands the recommended course of action its frame at no political cost. The successor system does not need to be sold as a foreign import or a repudiation of the culture. It is the version of the enterprise that meets the criteria of the label the enterprise already claims. Every one of the five rows in Table 3 converts directly into a design requirement, and Section 5.5 performs that conversion.

3 · Red Team Assessment

3.1 · Biases and Weaknesses in This Analysis

Retrospective coherence — the premortem’s own failure mode. The premortem’s strength is fluent causal narrative, and fluency is not evidence. Prospective hindsight debiases the audience but can bias the author: a stipulated collapse invites a story in which every symptom points the same direction, when real incubation records are noisier and include signals that pointed nowhere. Mitigation is procedural — every stipulated fact remains tagged as posited until the Section 5.2 data confirms or kills it, and the paper’s claims are deliberately structured so that each theory application can fail independently without collapsing the thesis.

Measurement bias cuts both ways. Goodhart’s Law is invoked here against the enterprise’s health metrics — but the same law forbids using those metrics to prove decline. If team counts cannot demonstrate health, they cannot demonstrate collapse either. The analysis is therefore only as strong as its independent telemetry, which does not yet exist. Until the utilization, held-resource, and actuarial data are mined, the stock-and-flow model runs on estimated parameters and the correct epistemic posture toward the 2031 date is: derived, disclosed, and unvalidated.

The date will be read as a prediction. However carefully framed as a model output, ‘2031’ will circulate as a forecast, and if 2031 passes without a visible collapse the analysis will be dismissed — even though Turner’s model says a quiet 2031 is the expected observation during incubation, and the ghost-phase mechanism says the metrics will read as solvent after the capability is gone. Mitigation: publish the model and its sensitivity band, never the date alone, and pre-commit to updating the estimate on a schedule as parameters are measured. The failure condition for this analysis is not ‘no collapse by 2031’; it is ‘measured inflow exceeds measured outflow.’

Conflating macro with micro. The most consequential internal risk: the entire intervention logic depends on the echelon taxonomy in Table 2 being right. If state-level withholding turns out to be driven by incident-style loss aversion, or GACC behavior turns out to be a true Prisoner’s Dilemma rather than an assurance game, the matched solutions fail — and fail in ways that look like confirmation that nothing works. The taxonomy should be treated as a hypothesis set, testable against the same data program.

Survivorship in the baseline. The ‘capable golden age’ against which decline is implicitly measured is, in part, the remembered prime of the cohort now retiring — a baseline subject to availability bias and survivorship. This is not a footnote; it is the load-bearing premise of the Dissenting Alternative, and it is granted there in full.

3.2 · Low-Probability, High-Impact Risks

The safe catastrophe. An artificially qualified CIM Operations Section Chief, sensing the gap between credential and capability, defaults to hyper-conservative tactical stagnation on a fire that punishes delay. The community is lost not to recklessness but to command paralysis wearing the uniform of risk management. This scenario is the signaling failure and the negative-transfer mechanism arriving on the same incident, and it is the single most likely precipitating event for the stipulated 2031 suspension — which is why the stipulation includes it.

Hoarding contagion and the assurance collapse. Framed in Stag Hunt terms, the risk is sharper than ‘trust erodes.’ Assurance games fail on information: a single verified instance of a GACC or IC holding critical assets while another region absorbs a mass-casualty WUI event on a documented UTF order does not merely damage trust — it flips the rational strategy for every player simultaneously, because each now has evidence that cooperation is unreciprocated. The cooperative equilibrium does not erode; it snaps, and the Balkanization of the national system follows as the new equilibrium, not as a breakdown.

The premature successor. A rushed replacement architecture built on the misdiagnosed game structures — most plausibly, a central authority empowered to strip resources unilaterally — validates every echelon’s defection story at once and accelerates the withdrawal it was meant to prevent. Section 5.4 addresses why the empirical governance record argues against the Leviathan design even when the diagnosis is right.

4 · Dissenting Alternative

The speciation thesis. The system is not dying. It is adapting — rationally, successfully, and without a decision memo — to the mission it actually has. Tactical containment of climate-driven WUI megafire with legacy linear doctrine is physically unwinnable; the fires exceed the control envelope of the tools. An enterprise cannot sustain investment in a capability that no longer produces outcomes, and this one has quietly stopped: it has reallocated from the production of tactical commanders to the functions that still generate value on a modern incident — logistics throughput, financial administration, liability processing, and interagency coordination. What the main analysis reads as a training-pipeline deficit is, under this interpretation, disinvestment in an obsolete capability. What it reads as hoarding is inventory management under honest scarcity. What it reads as a demographic cliff is a bridge cadre carrying a legacy function to the end of its useful life. The Incident Management Team did not fail to remain what it was; it speciated into what the incident environment now selects for — an Incident Administrative Team — and the collapse narrative is an availability-bias artifact of a retiring generation measuring the system against its own remembered prime, a golden-age baseline that may never have existed as remembered.

Note carefully what the dissent does not dispute. It concedes the signaling failure, the metric corruption, and the pipeline arithmetic in full — it disputes only whether they matter, because it disputes that the capability being lost is the capability the mission requires. That makes the disagreement empirical rather than cultural, and it makes the falsifier shared: the utilization and outcome data in Section 5.2 adjudicate between thesis and dissent. If tactical command capability no longer drives incident outcomes at the margin — if the variance in outcomes is explained by fuels, weather, and pre-incident mitigation rather than by command quality — the dissent stands, and the honest response is to design the successor architecture for the administrative mission explicitly, fund it as such, and stop performing a tactical identity the enterprise no longer resources. The most damaging outcome is not either answer; it is the current one, in which the enterprise claims the tactical mission, staffs the administrative one, and validates neither.

5 · Recommended Course of Action

The recommendation is sequenced deliberately: evidence gates structure. The data program comes first because it parameterizes the model, adjudicates the dissent, and — at the GACC echelon — is itself the intervention. Structural commitments follow validation, not the reverse.

5.1 · Name the Method, Publish the Model

Release this analysis explicitly as a premortem, with the method stated in the first paragraph — as this version does. The declaration is not decoration; it is the inoculation against the two predictable dismissals ('you fabricated a history' and 'you predicted a date'), and it models the practice the enterprise should adopt. Publish the stock-and-flow model with its parameters and sensitivity band, invite refutation on the parameters, and pre-commit to a published update cadence as measured values replace estimates. The paper should be attackable in exactly one currency: evidence.

5.2 · Close the Data Gaps First

Three data streams, mined from IROC/ROSS and workforce records, each serving multiple functions simultaneously.

Incident-level utilization rates. Quantify operational shifts spent in staging or low-priority mop-up/suppression repair prior to demobilization during PL 4/5. Function one: parameterizes the over-ordering claim and the decoupled-accountability mechanism. Function two: primary falsifier between thesis and dissent — this is where command quality either shows up in outcomes or does not.

GACC-level UTF versus held-resource ratios. Determine the historical delta between unassigned regional reserves and UTF orders generated back to NICC, and then publish it continuously, enterprise-wide, in near-real time. Function one: model input. Function two: the transparency mechanism that converts the Stag Hunt — credible information about what every GACC is holding is the assurance the cooperative equilibrium requires. One data product, and it is both the measurement and the fix. Accept that some orders are never placed because they will not get filled.

The actuarial audit. Median age, qualification currency, actual mobilization acceptance rate, and projected participation horizon of fully qualified CIMT positions in the national rotation — the measured distribution that replaces the estimated outflow parameter and moves the 2031 date to wherever the evidence puts it.

5.3 · Restore the Signal

Per Spence, a degraded signal is repaired one way: re-couple the cost of acquiring it to the capability it claims. For command qualification that means objective, forced-failure, stress-inoculated simulation vetting — assessment events that a meaningful fraction of candidates do not pass, designed around non-linear scenarios precisely because the recognition-primed heuristics built on linear fires must be shown to fail before the operator faces the failure on a real incident. Dörner’s experimental paradigm is directly adaptable here. As an immediate step, initiate the Red Cell hollow-force exercise: a multi-ignition WUI scenario run under the assumption of zero out-of-GACC IMT availability and fifty percent municipal declination — a stress test of the enterprise’s actual posture rather than its rostered one.

5.4 · Govern Polycentrically, Not by Leviathan

The intuitive fix at the enterprise level — and the one most often proposed — is a MAC-level entity with unilateral authority to strip idle resources from an IMT without the IC’s consent during extreme national drawdowns. That design should be rejected in its pure form, for an evidence-based reason: it is Hardin’s ‘Leviathan’ solution, and Ostrom’s Nobel-recognized empirical work on common-pool and public-good regimes found that centralized coercive authority is routinely outperformed by polycentric governance — nested decision-making with transparent monitoring, graduated sanctions, and participation by the governed in rule-setting. Worse, the Red Team’s premature-successor risk applies directly: unilateral stripping authority confirms every echelon’s low-trust defection story and accelerates withdrawal from the national system. The sound design retains a dynamic demobilization mechanism but embeds it as the top rung of a graduated-sanction ladder inside a polycentric structure: transparent held-resource reporting does the bulk of the work (Section 5.2), peer-level accountability among GACCs does the next increment, and compulsory demobilization exists as the visible, rarely used backstop — legitimate because the governed echelons participated in defining its triggers. The instrument survives; the governance model around it changes to the one with the better empirical record.

5.5 · Architect the Successor as a Legitimate HRO

Cease attempts to repair the interagency roster model; the pipeline arithmetic forecloses it. Architect regional standing task forces designed explicitly against the five criteria of Table 3, each converted into a requirement.

Preoccupation with failure → near-miss and waiver telemetry as a core system function with dedicated analytic staffing — every waiver an automatic input to the workforce model, every UTF a logged system-health event.

Reluctance to simplify → an institutionalized red team and dissent channel at the coordination level, with the multi-echelon diagnostic discipline of Table 2 as standard practice: name the game before prescribing the fix.

Sensitivity to operations → the utilization and held-resource telemetry of Section 5.2 as a permanent, real-time operating picture — the enterprise-level ‘bubble’ — not a one-time study.

Commitment to resilience → structural slack budgeted as a line item: bench depth requirements, surge reserve targets, and a prohibition on waiver-sustained rosters, so that hollowness registers as a funding breach rather than an administrative accommodation.

Deference to expertise → qualification exclusively by validated capability (Section 5.3), so that rank once again locates expertise and decision migration under anomaly has somewhere to go.

The framing for the enterprise should be stated exactly once and then left alone: this is not the imposition of a foreign model. It is the enterprise becoming, in fact, what it already claims to be. Every criterion in Table 3 is drawn from a literature the fire service cites about itself. The successor system simply takes the citation seriously.

5.6 · Implementation Considerations

Sequencing. Data program (5.2) and signal restoration pilot (5.3) proceed immediately and in parallel; model validation and thesis-versus-dissent adjudication follow the data; structural architecture (5.5) is designed now but committed only as evidence gates clear. Parallel structural commitments ahead of evidence would repeat the enterprise’s signature error at the moment of trying to correct it.

Handling the date. In every briefing and document, the 2031 estimate travels with its derivation and sensitivity band, never alone. The one-sentence version: ‘the model computes when the qualified stock crosses minimum viable rotation; the current central estimate is 2031; here is what would move it.’

The paper as instrument. This document doubles as a premortem exercise instrument. Convened groups — a MAC group, an IMT C&G staff, an agency leadership team — are given the stipulation block and asked to extend the incubation record from their own echelon’s experience before reading Section 2. The failure paths they add are data, and the exercise builds the constituency the successor architecture will need.

Theory Anchors

These are the established constructs on which this paper is built. Each was developed, tested, and in most cases formalized in another high-consequence domain; none was developed for wildland fire, and none requires modification to apply to it. The companion paper, *Sustaining Decision Performance*, carries its own back matter of primary regulatory and standards sources supporting its comparative claims.

- Dörner, D. (1996). *The Logic of Failure: Recognizing and Avoiding Error in Complex Situations*. Experimental demonstration of operator failure modes in non-linear dynamic systems; anchors the negative-transfer and command-paralysis analysis.
- Forrester, J. (1961); Meadows, D. (2008). *Industrial Dynamics; Thinking in Systems*. Stock-and-flow modeling; anchors the pipeline insolvency model of Section 2.4.
- Goodhart, C. (1975); Campbell, D. (1979); Strathern, M. (1997). “When a measure becomes a target, it ceases to be a good measure” (Strathern’s phrasing); anchors the metric-capture analysis.
- Kahneman, D., & Tversky, A. (1979). *Prospect Theory: An Analysis of Decision under Risk*. Loss aversion; anchors the incident-level demobilization analysis.
- Klein, G. (1998; 2007). *Sources of Power; “Performing a Project Premortem,” Harvard Business Review*. Recognition-primed decision-making and its boundary conditions; the premortem method used by this paper.
- Olson, M. (1965). *The Logic of Collective Action*. Free-rider dynamics in public-good provision; anchors the state-echelon analysis, rather than the commonly applied tragedy-of-the-commons framing.
- Ostrom, E. (1990). *Governing the Commons*. Empirical case against Leviathan solutions; polycentric governance, monitoring, and graduated sanctions; anchors Section 5.4.
- Perrow, C. (1984). *Normal Accidents: Living with High-Risk Technologies*. Interactive complexity and tight coupling; applied here to the megafire environment, not the mobilization enterprise.
- Rasmussen, J. (1997). “Risk Management in a Dynamic Society,” *Safety Science*. Migration toward the boundary of acceptable performance under efficiency pressure; anchors the CIM-merger analysis.
- Skyrms, B. (2004). *The Stag Hunt and the Evolution of Social Structure*. Assurance-game dynamics; anchors the GACC-echelon analysis, rather than the commonly applied Prisoner’s Dilemma framing.
- Snook, S. (2000). *Friendly Fire: The Accidental Shootdown of U.S. Black Hawks over Northern Iraq*. Practical drift — locally rational procedural adaptation accumulating into system failure.
- Spence, M. (1973). “Job Market Signaling,” *Quarterly Journal of Economics*. Signaling theory; anchors the qualification-system analysis, rather than the folk Peter Principle framing.
- Tainter, J. (1988). *The Collapse of Complex Societies*. Declining marginal returns on complexity; anchors the checklist-accretion analysis and the collapse frame.
- Turner, B. (1978); Turner, B., & Pidgeon, N. (1997). *Man-Made Disasters*. Disaster incubation, decoy phenomena, failure of foresight; the retrospective spine of this paper.

- Vaughan, D. (1996). The Challenger Launch Decision. Normalization of deviance; applied here to waiver accumulation and standards erosion, not to the discrete merger decision.
- Weick, K., & Sutcliffe, K. (2007). Managing the Unexpected: Resilient Performance in an Age of Uncertainty. The five HRO principles; the audit standard of Section 2.6 and the design standard of Section 5.5.
- Woods, D. (with Hollnagel, E., & Leveson, N., 2006). Resilience Engineering: Concepts and Precepts. Brittleness, decompensation, and the law of stretched systems; anchors the ghost-phase analysis.

“The 2031 collapse is not the failure. It is the reveal. The failure occurred years earlier, while the metrics still read success.”

It is March 2031, and the national Complex Incident Management Team rotation has been suspended. That history is stipulated — this is a premortem, in the sense Gary Klein formalized: assume the system has already failed, and explain why. The stipulation removes the social cost of describing failure, so the failure can be described accurately. No reader is asked to agree the system will collapse. They are asked only to explain a collapse — the same permission an after-action review creates, applied in advance.

The collapse requires no villain and no novelty. It is the compound of failure modes named, formalized, and largely solved in other high-consequence industries: a qualification signal decoupled from capability, a workforce pipeline that is arithmetically insolvent, three distinct hoarding games misdiagnosed as one behavior, health metrics captured by their own targets, and a high-reliability identity claimed in doctrine while meeting none of the five criteria the literature specifies. Treating these failure modes as unique to fire is the specific reason every intervention to date has been a half-measure.

Built in the house analytical format — Executive Summary, Support & Analysis, Red Team Assessment, Dissenting Alternative, Recommended Course of Action. The 2031 date is not a prediction; it is the central estimate of a workforce stock-and-flow model, published with its parameters, its sensitivity band, and an invitation to refute it with the enterprise’s own records. The dissent is carried at full strength, the data program that adjudicates it is specified, and the document doubles as a premortem exercise instrument for any MAC group or command staff willing to extend the stipulation from their own echelon’s experience.

The Reveal belongs to the Cogplexiti Systems & Decision Analysis line, alongside The System Tier and The Incident Tier. Its companion discussion paper, Sustaining Decision Performance, isolates the individual-level instance of the signaling failure — the qualified decision-maker who is never re-evaluated — in a measured register prepared for interagency review. Read together, one paper supplies the verifiable instance, the other what the instance becomes at enterprise scale.

ABOUT THE AUTHOR

Jayson Coil is the founder of Cogplexiti LLC, an Assistant Chief with the Sedona Fire District, and a Type 1 Operations Section Chief on Southwest Incident Management Team 1. A U.S. Army veteran and red team instructor, he works at the intersection of operational command and the study of how organizations decide under uncertainty. He has helped develop the systems agencies use to manage major incidents and has presented for the National Fire Academy, the U.S. Forest Service, the Department of the Interior, FEMA, the Haas School of Business at UC Berkeley, and Los Alamos National Laboratory.

