

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

ORDER IT AND IT COMES

The Mutual Aid Assumption
and the Indicators That Would
Tell Us It's Failing

JAYSON COIL

COGPLEXITI · SYSTEMS & DECISION ANALYSIS

WHITE PAPER

Order It and It Comes

The Mutual Aid Assumption and the Indicators That Would Tell Us It's Failing

JAYSON COIL

Copyright © 2026 by Jayson Coil. All rights reserved.

Published by Cogplexiti LLC · Flagstaff, Arizona · cogplexiti.com

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

Cogplexiti · Systems & Decision Analysis

ISBN:979-8-9970824-0-6

First edition · July 2026

Suggested citation: Coil, Jayson. 2026. *Order It and It Comes: The Mutual Aid Assumption and the Indicators That Would Tell Us It's Failing*. Flagstaff, AZ: Cogplexiti.

How to Read This Paper

This white paper belongs to the Cogplexiti Systems & Decision Analysis line and follows the house analytical format. It operationalizes the line's alternative-futures and premortem work: those analyses described where the system may be going and how its plans may fail; this one builds the watching. The register in Exhibit A is the deliverable; the body is the argument for it. Production insertion points for current-season figures are marked in gold throughout.

Contents

How to Read This Paper	3
Contents	3
1. Executive Summary	4
2. Support and Analysis.....	5
2.1 One Order, One Morning	5
2.2 The Assumption Is a Stack	6
2.3 The Failure Signature.....	8
2.4 The Register	9
3. Red Team Assessment	10
4. Dissenting Alternative.....	11
5. Recommended Course of Action	12
Exhibit A — The Mutual Aid Assumption Register (v1)	13
References.....	14

1. Executive Summary

The national wildland fire system rests on one assumption so load-bearing that it is almost never stated: resources scale with demand. Place the order and something arrives. Every incident strategy that holds a fire small "until reinforcements get here," every community evacuation plan that budgets time against an expected suppression response, every cooperator agreement, every insurance model, and every homeowner's implicit bargain with the ridge behind the house prices that assumption in. The assumption was true when the architecture was built. Whether it is still true — on the days that matter, which are by definition the days everyone orders at once — is a question the system does not currently instrument.

The assumption does not fail with an announcement. It fails at the margin, silently, worst-day-first: as unable-to-fill responses on resource orders, as teams rostered short and rotations consumed early, as aircraft on contract but not crewed, as strategies quietly re-scoped to fit what actually arrived. The field absorbs each increment as workload and improvisation, so the degradation is experienced everywhere and reported nowhere as what it is — system state. Every autumn the same argument restarts from zero: one side says the system nearly broke, the other says the system adapted as it always has, and both are describing the same season, because no agreed set of observables exists against which the claim could be settled.

This paper's deliverable is that set of observables. Using two instruments — the Key Assumptions Check and Signposts of Change (UFMCS 2019) — it decomposes the master assumption into its load-bearing components, interrogates each, and pairs each with indicators that are unique, valid, observable, and collectable, and with a pre-decided response for the day the indicator fires. The register converts a thirty-year argument into a monitoring problem. It does not claim the assumption has failed. It claims the question "is the surge assumption still holding?" should be checkable rather than arguable — and that a system which cannot check it is making promises, to communities and to itself, that it cannot know it can keep.

Two companion analyses stand behind this one. The alternative-futures work established where the system may be going; the premortem established how its plans may fail. Both produce their value only if someone is watching for the signatures they describe. This register is the watching.

2. Support and Analysis

2.1 One Order, One Morning

The scene is Pine Bench — a constructed mountain town, but the arithmetic is anywhere. A start on the bench above town, second burning period coming, fuels running unbroken to the ridge, and one road in. The duty officer does what the system trained him to do and what the plan on the shelf assumes he will do: he builds the order. Two Type 1 crews, a strike team of engines, a Type 3 helicopter, structure protection overhead. It is a competent order against a familiar problem, and the plan for the bench — the trigger points, the evacuation timing, the defensible stretch along the one road — was written on the assumption that an order like this one gets filled.

The responses come back through the morning. One crew, unable to fill, no date. The engines: two of six, the rest UTF because three forests south the same wind event is producing the same orders. The helicopter is on contract and not crewed today. Dispatch is apologetic and precise; everyone in the chain is doing their job, and the region is at a preparedness level where the ordering system is functioning exactly as designed — which is to say, rationing. By 1400 the duty officer is doing the other thing the system trained him to do, the one that appears in no plan: re-scoping the strategy to fit what arrived. The line on the bench becomes point protection if resources allow. The defensible stretch along the road becomes a trigger point for earlier evacuation. Nobody is endangered by the revision; that is what the revision is for. The cost lands as arithmetic — a bench that the plan defended and the strategy now concedes, contingent on weather. If the wind lies down, no one will ever know the assumption failed that morning. That is the point. The assumption fails silently on the mornings the wind lies down, and the record keeps no ledger of them — so the season's luck and the system's capacity remain indistinguishable in the data, right up until a morning when they are not.

The rest of this paper is about instrumenting that morning: naming the assumption the duty officer's plan was priced on, decomposing it, and building the indicator set that would let the system distinguish, before the fact and in public, between a bad day and a failing assumption.

2.2 The Assumption Is a Stack

FIGURE 1

The assumption is a stack

"Resources scale with demand" is seven load-bearing assumptions on one unexamined foundation. Each is paired in Exhibit A with observable indicators and a pre-decided response.



Silence is not evidence of health.

"Resources scale with demand" is not one assumption. It is a stack of seven, each load-bearing, each with its own failure mode, and the Key Assumptions Check discipline — why must this be true, what is it based on, under what conditions does it fail (UFMCS 2019) — applies to each in turn. Stated as they would be stated by the system if the system said such things out loud:

A1 — Local capacity exists. Initial attack succeeds locally most of the time, so the national system only sees the residue. *Based on:* home-unit staffing models built in a different labor market. *Fails when:* vacancy rates hollow out modules that exist on paper;

the militia tradition thins as land-management workforces shrink; local capacity becomes a roster fiction that the ordering system treats as fact.

A2 — The geographic area can backfill. Neighbors inside the geographic area cover each other because their bad days are assumed not to coincide. *Based on:* historical weather independence between neighboring units. *Fails when:* the synoptic events driving modern fire seasons are regional by nature — the same wind event, the same lightning bust, the same critically dry fuel bed across an entire geographic area at once. Correlated demand is the quiet killer of mutual aid arithmetic; insurance learned this lesson about catastrophe risk decades ago, and the fire system's version is a wind event that turns every lender in the compact into a borrower on the same afternoon.

A3 — The nation can backfill the geographic area. When one GACC exhausts its available resources, the others feed it. This assumption depends on seasonal asynchrony among regions—for example, the Southwest traditionally peaks and begins releasing resources as activity builds in the Northwest. It fails when the national season lengthens at both ends, regional peaks overlap, and multiple GACCs simultaneously compete for the same scarce resources. Multi-GACC simultaneity at high preparedness levels is therefore the most direct measure of this assumption's health, and the trend in that measure may be the single most important number in this paper. Over the last ten completed years, the nation reached PL5 in five years; the median first PL5 declaration among those years was July 28. The highest annual duration was 68 days at PL5 in 2021. The median annual duration was 17 days when all ten years are included, or 45 days when limited to years that actually reached PL5.

A4 — Command capacity exists to run what arrives. Resources without teams are a crowd. *Based on:* an IMT bench sustained by a militia model in which team service is a collateral duty employers release people to perform. *Fails when:* rosters fill on paper and travel short; the rotation depth that looked like resilience is consumed by August; the pipeline that replaces retiring command-level experience runs an order of magnitude slower than the attrition drawing it down — the pipeline–attrition loop the systems analysis in this line has already documented (Coil 2026c).

A5 — Aviation arrives contracted, crewed, and current. *Based on:* an industrial base of vendors, airframes, and pilots that renews itself. *Fails when:* contract counts hold while crewed-and-available days quietly fall; the pilot and mechanic labor market tightens; "on contract" and "mission-capable today" diverge — a divergence invisible in any number the public ever sees.

A6 — Surge beyond the system exists. Guard and active military augmentation, cooperator capacity, compact resources. *Based on:* the assumption that the emergency is ours alone — that the augmenting institutions are not simultaneously tasked. *Fails when:* the same summer that maxes the fire system floods a coast or fills the Guard's calendar;

surge capacity shared across emergencies is surge capacity that arithmetic can assign to only one of them.

A7 — International augmentation is available. *Based on:* hemispheric asynchrony — their winter is our summer. This is the worked example of an assumption that was true when the architecture adopted it and is visibly expiring: the partner seasons are lengthening toward ours just as ours lengthens toward theirs, and the northern partner increasingly competes for the same airframes and crews in the same months instead of supplying them. An assumption of non-overlap, adopted when it was true, does not announce its own expiration; it simply fills fewer orders each year while remaining on the books.

A8 — Demand is a given. The stack above is supply-side, and the register would be dishonest without naming the assumption underneath all of it: that the demand placed on the system is exogenous — weather and ignitions — rather than partly chosen through strategy. It is not. Strategy selection sets demand, suppression-default settings concentrate national resources on fires where they buy the least, and the suppression-fuels loop converts this season's successes into future seasons' demand (Calkin, Thompson, and Finney 2015; Coil 2026c). A register that watched only supply would measure the rope while ignoring the hands pulling on both ends. The dissenting alternative in Section 4 presses this point to its strongest form, and the register concedes it a seat: the exhibit carries demand-side indicators alongside the supply stack.

2.3 The Failure Signature

Assumptions of this class do not fail the way equipment fails. Three properties of the failure mode determine what the indicators must look like.

It fails silently. No component of the stack reports its own erosion. A7 fills fewer orders; nobody publishes "the hemispheric asynchrony assumption is now false." A4's rotation depth is consumed; the sit report shows teams committed, not the bench behind them. The Pine Bench morning generates no record distinguishable from an ordinary busy day. Silence is not evidence of health; it is a property of the architecture, which was built to fill orders, not to report on its own capacity to keep filling them.

It fails at the margin, worst-day-first. Average-day performance is nearly useless as a health measure, because the assumption only matters on correlated-demand days — and those are precisely the days when every indicator is hardest to collect and every operator is too committed to report. A system can post acceptable annual fill rates while failing catastrophically on the eight days that produced the year's losses. The indicators must therefore be conditioned on season position and preparedness level, not annualized.

Its degradation is absorbed as workload. This is the property that defeats ordinary reporting. When the order comes back short, the field does what the field does: extends shifts, re-scopes strategy, improvises, copes. The gap between what was assumed and what arrived is converted, on contact, into fatigue and improvisation — experienced by everyone, recorded by no one, and feeding the same attrition loop that thins the bench for next season. The system's adaptive competence is real, and it is exactly what launders the signal. Any monitoring design that waits for the field to report assumption failure will wait forever, because the field's professional pride and operational tempo both point the other way. The indicators must be drawn from transactional exhaust — orders, fills, UTFs, rosters, contract status — not from testimony.

2.4 The Register

The instrument is the pairing discipline this line has used before: assumption → indicator → pre-decided response, with each indicator held to the signpost standard — unique, valid, observable, collectable (UFMCS 2019). Exhibit A carries the full register; three design rules govern it and need stating in the body.

Indicators are conditioned, not annualized. Every rate is measured by season position and preparedness level — UTF rate *at PL4-plus*, crew availability *in the peak window* — because Section 2.3's failure signature lives in the conditional distribution and vanishes in the annual average.

Indicators come in pairs where gaming is cheap. A fill rate can be improved by filling orders or by discouraging them; the register therefore pairs UTF rates with orders-placed trends, and contract counts with crewed-available days, so that managing one number without the other becomes visible. Section 3.2 develops the point.

Every trip-wire carries a pre-decided response. An indicator without a response tier is a dashboard, and dashboards are where warnings go to be acknowledged. The register's response column is deliberately modest — it prescribes decisions about *decisions*: at Tier 2, strategy defaults and pre-positioning change and public expectations are reset before the season; at Tier 3, the assumption is declared void for planning purposes, and every plan priced on it — community evacuation timing, cooperator agreements, the bench above Pine Bench — is formally notified that the surge it budgets for should no longer be assumed. Tier 3 is not a fire-season press release. It is a preseason act of honesty with the parties who built plans on the system's promise, and the communication architecture it requires — who convenes whom, inside what emotional window — is the machinery the companion insurance-crisis analysis has already designed (Coil 2026d).

3. Red Team Assessment

3.1 The data problem: a register built on unavailable data is a wish list. Several proposed indicators are inconsistently tracked, definitionally unstable across years, or not public. UTF data exists in the ordering system's transactional record but its extraction and definitions have shifted across platform generations; aviation crewed-available days are vendor- and contract-administration data that no public product aggregates; rotation-depth on the team bench is knowable inside coordination centers and published nowhere. Worse, the cleanest indicator is contaminated at the source by *shadow demand*: experienced ordering managers stop placing orders they know will UTF, so the measured gap shrinks exactly as the real gap grows. The paper's answer is structural, not rhetorical: the register separates Tier-A indicators (collectable today from published products — preparedness-level records, sit-report team and crew commitments, simultaneity counts) from Tier-B indicators (requiring a reporting decision — conditioned UTF extracts, crewed-available aviation days, bench-depth reporting), and states plainly that the Tier-B column is itself a finding. A system that wanted to know whether its load-bearing assumption still held would collect these six numbers; the fact that collecting them requires a decision nobody has made is evidence about the system's relationship to the question. Shadow demand gets its own paired indicator — orders-placed trend alongside fill rates — imperfect, but the pairing at least makes demand suppression visible as divergence.

3.2 Goodhart's law: a published register invites managing the indicator. Once UTF-at-PL4 becomes a number leadership is judged on, the cheapest response is not capacity — it is definitional drift, order discouragement, reclassification. This is not a hypothetical failure mode; it is the standard fate of published performance measures. Three countermeasures are built in, and their limits acknowledged. The pairing rule (2.4) makes single-number management visible as divergence between paired series. The register's ownership design (5.2) separates the measuring function from the managed function — the entity reporting the indicators cannot be the entity whose performance they describe. And the response tiers attach to *decisions about planning assumptions*, not to performance judgments, which lowers the incentive to game by lowering what gaming buys. The honest residual: no indicator regime survives a determined institutional decision to look good rather than see clearly. The register raises the cost of that decision from free to visible. It cannot raise it higher.

3.3 The crying-wolf problem: the system has been "at the breaking point" for thirty years. Every retrospective since the 1990s has declared the system stretched past its limits, and every following season the system adapted, absorbed, and held. A skeptical reader — especially one who has personally absorbed thirty years of the adapting — can file this paper with the other alarms. The reply has to concede the record to keep its

credibility: the system's adaptive capacity is real, and the alarms were wrong in form even where right in direction, precisely because they were unfalsifiable — "nearly broke" and "adapted as always" describe the same season when no observables are specified in advance. That is the paper's actual claim against the crying-wolf history: this register is the *falsifiability instrument* the previous alarms lacked. If the indicators hold flat across a decade, the register will say the assumption is holding, in public, and this paper's concern will be shown overstated — an outcome the design accepts on purpose. What the register refuses to accept is the third possibility the current arrangement permits: that the assumption fails by degrees while the absorption described in 2.3 launders every increment into workload, and the argument restarts from zero each autumn because nobody wrote the criteria down when it could still be done calmly. Adaptation is not the rebuttal to the failure signature. Adaptation *is* the failure signature, measured from inside.

3.4 The register may understate by design. A structural challenge the drafting must carry rather than resolve: indicators built from transactional exhaust see the formal system only. Cooperator fatigue, the quiet non-renewal of a rancher's dozer agreement, the volunteer department that stops answering the second page — the informal capacity that correlated-demand days actually run on degrades invisibly to every column in Exhibit A. The register measures the load-bearing wall it can see. The reader should hold, and the paper should state, that the building also stands on framing the instrument cannot inspect.

4. Dissenting Alternative

The assumption is not failing — it is being spent badly. The credible competing interpretation runs: surge capacity is adequate for what the system should be doing and inadequate for what it is choosing to do. National resources are routinely committed to fires threatening little, held by suppression-default strategy settings long after the marginal crew stopped buying risk reduction; the highest-capability crews in the country spend weeks of peak season on assignments a candid assessment would not have staffed at that level. On this reading, the Pine Bench morning is real but misdiagnosed — the crews that UTF'd were not absent from the system; they were committed elsewhere, by choice, against lower stakes. The binding variable is demand strategy, not supply; the register instruments the wrong side of the ledger; and the fix is decision discipline about which fires get what — a strategy argument, in a decade the dissent says needs strategy arguments, not another capacity argument that ends in a procurement request. The suppression-fuels loop sharpens the point: capacity added without demand discipline is spent by the same defaults that created the shortage, and partially converts into future demand (Calkin, Thompson, and Finney 2015).

The dissent is right about the ledger and wrong about the remedy's sufficiency, and the register has already paid it the concession it earns: A8 and the demand-side indicators exist because of this argument. Two rebuttals bound it. First, the dissent's own remedy *requires* the register — strategy discipline about which fires get staffed is only defensible, politically and legally, if the scarcity it responds to is documented in public observables rather than asserted by the officials rationing the response; a demand-side fix without a supply-side record is indistinguishable, to a burned community, from abandonment by preference. Second, the correlated-demand days that break the assumption break it even under disciplined demand — the wind event that produces forty legitimate simultaneous demands does not care how many illegitimate ones were declined in June. Demand discipline buys margin; it does not repeal simultaneity. The corrected conclusion, adopted in the COA: the register carries both sides of the ledger, and its publication serves the dissent's program as much as the paper's own.

5. Recommended Course of Action

5.1 Publish the register, tiered by collectability. Exhibit A ships in two visible tiers: Tier-A indicators assembled from published products beginning with the current season, so the register produces its first reading immediately; Tier-B indicators published as a named reporting gap — the six numbers the system would collect if it wanted to know, stated as a request with an owner, not as a complaint. Publishing the gap is half the register's work.

5.2 Name a monitoring owner outside the managed function. The register is read on a cadence — preseason baseline, biweekly at national PL3 and above, postseason against the decade trend — by an entity that reports on the system rather than for it. Candidates exist at several altitudes (coordination-system staff work, external research capacity, a standing review function); the design requirement is independence of the measuring function, per Red Team 3.2, and continuity across seasons, because the register's entire value is the trend line no single season can produce.

5.3 Pre-decide the response tiers, in writing, out of season. Tier thresholds and their consequences are set in the calm months by the parties who will live with them: Tier 1, monitor and report; Tier 2, strategy defaults and pre-positioning adjust, and public expectations are reset before the season opens; Tier 3, the assumption is declared void for planning purposes and every downstream plan priced on it is formally told so. The tier definitions are the register's hard political work and the reason it cannot be a dashboard: a threshold negotiated during a smoke event will be negotiated by the smoke.

5.4 Build the Tier 3 communication before it is needed. Telling communities, cooperators, and elected officials that the surge their plans assume should no longer be assumed is the single most consequential act the register can trigger, and it cannot be

improvised inside a fire season's emotional window. The convening architecture, the sequencing, and the window discipline for exactly this class of message are already designed in the companion insurance-crisis analysis (Coil 2026d); the register's Tier 3 plugs into that machinery rather than inventing its own. The alternative to a designed Tier 3 conversation is the undesigned one — conducted street by street, in the language of conceded ground, by the futures this line has already written down (Coil 2026c).

End state. A system whose most load-bearing assumption is watched instead of argued about — where "is the surge still there?" has a public answer updated on a schedule, where the answer's arrival was decided before anyone needed it, and where the morning the order comes back empty above Pine Bench is legible for what it is, because somebody wrote down, years earlier and in the calm, what that morning would look like.

Exhibit A — The Mutual Aid Assumption Register (v1)

Discipline: assumption → indicator (unique, valid, observable, collectable) → pre-decided response. Conditioning rule: all rates reported by season position and national/GACC preparedness level, never annualized. Pairing rule: paired indicators are read together; divergence between pair members is itself a signal. Collectability: [A] = assemblable now from published products; [B] = requires a reporting decision — the Tier-B column is itself a finding.

A1 — Local initial-attack capacity exists. Indicators: federal fire workforce vacancy rate by unit type [B]; modules statused unavailable-due-to-staffing, peak window [B]; initial-attack success rate trend, conditioned on fuel/weather severity [A/B]. Response: Tier 2 — home-unit reliance assumptions in unit plans re-priced; Tier 3 — cooperators formally notified of reduced federal IA reliability by zone.

A2 — The geographic area can backfill. Indicators: intra-GACC UTF rate at GACC PL4+ [B], paired with intra-GACC orders-placed trend [B]; count of days per season with 3+ dispatch zones in one GACC at high activity simultaneously [A, derivable]. Response: Tier 2 — pre-positioning triggers rewritten against correlated-demand days rather than single-fire forecasts.

A3 — The nation can backfill the geographic area. Indicators: days per year at national PL5, and earliest-PL5 date, ten-year trend [A]; days with 3+ GACCs simultaneously at PL4+ [A]; inter-GACC UTF rate at national PL4+ [B] paired with orders-placed [B]. Response: Tier 2 — national strategy-default guidance adjusts at defined simultaneity counts; Tier 3 — declared void for planning: downstream plans notified.

A4 — Command capacity exists. Indicators: IMTs available vs. committed at peak, and date rotation depth is first exhausted [A, sit report]; roster fill shortfalls at mobilization [B]; command-and-general-staff trainee pipeline throughput vs. projected retirements [B]. Response: Tier 2 — team configuration and rotation policy adjusted preseason; bench-depth reporting elevated to leadership dashboard with the same cadence as acres.

A5 — Aviation arrives contracted, crewed, current. Indicators: exclusive-use contract counts by platform, trend [A]; crewed-and-available days as percent of contract days [B], paired with contract counts [A]; UTF rate on aviation orders at PL4+ [B]. Response: Tier 2 — strategy defaults de-weight aviation-dependent tactics at defined availability thresholds; contract-structure review triggered.

A6 — Surge beyond the system exists. Indicators: Guard/military augmentation requested vs. received, with response-time trend [B]; count of days per season with concurrent national-level emergencies drawing the same augmentation pool [A, derivable]. Response: Tier 2 — surge assumptions in national plans conditioned on concurrent-emergency probability, not availability-on-paper.

A7 — International augmentation is available. Indicators: partner-nation season overlap days with US peak, ten-year trend [A, derivable]; international resources requested vs. received, trend [A/B]; instances of direct competition for the same contract airframes [B]. Response: Tier 2 — international augmentation reclassified from assumed capacity to opportunistic capacity in national planning documents.

A8 — Demand is a given (it is not). Indicators: share of national Type 1 crew and large-airframe commitments on incidents below defined values-at-risk thresholds, peak window [B]; strategy-mix trend (full suppression vs. other) on ignitions in comparable fuel/weather classes [A/B, derivable from incident records]. Response: Tier 2 — strategy-default review with the same standing as the supply-side responses; the register reports both ledgers in every reading.

References

- Calkin, David E., Matthew P. Thompson, and Mark A. Finney. 2015. "Negative Consequences of Positive Feedbacks in US Wildfire Management." *Forest Ecosystems* 2 (9).
- Coil, Jayson. 2026a. *Operational Leader's Field Guide*. Flagstaff, AZ: Cogplexiti.
- Coil, Jayson. 2026b. *The Second Skill: Why the Qualification System Trains Recognition and Assumes the Rest*. Cogplexiti Systems & Decision Analysis. Flagstaff, AZ: Cogplexiti.
- Coil, Jayson. 2026c. *The System Tier*. Cogplexiti Systems & Decision Analysis Series No. 1. Flagstaff, AZ: Cogplexiti.

- Coil, Jayson. 2026d. *Nobody's Villain*. Cogplexiti Systems & Decision Analysis. Flagstaff, AZ: Cogplexiti.
- Heuer, Richards J., Jr., and Randolph H. Pherson. 2015. *Structured Analytic Techniques for Intelligence Analysis*. 2nd ed. Thousand Oaks, CA: CQ Press.
- University of Foreign Military and Cultural Studies (UFMCS). 2019. *The Red Team Handbook: The Army's Guide to Making Better Decisions*. Version 9.0. Fort Leavenworth, KS: TRADOC G-2.

The assumption fails silently on the mornings the wind lies down — and the record keeps no ledger of them.

The national suppression architecture rests on one assumption so load-bearing it is almost never stated: resources scale with demand. Every incident strategy, evacuation plan, cooperator agreement, and community bargain prices it in. Rising simultaneity is eroding it at the margin, worst-day-first, with the degradation absorbed by the field as workload and reported nowhere as system state — so every autumn the same argument restarts from zero.

“Is the surge still there?” should be checkable, not arguable. A system that cannot check it is making promises — to communities and to itself — that it cannot know it can keep.

The analysis decomposes the master assumption into an eight-assumption stack — local capacity, geographic and national backfill, command capacity, aviation, external surge, international augmentation, and the demand-side assumption underneath them all — and interrogates each with the Key Assumptions Check. It names the failure signature that defeats ordinary reporting: silent, marginal, laundered into workload by the system’s own adaptive competence.

The deliverable is the register: assumption, indicator, pre-decided response — indicators conditioned on season position and preparedness level, paired where gaming is cheap, tiered by collectability, with the reporting gaps published as findings in their own right. Tier 3 is a preseason act of honesty with everyone who built plans on the system’s promise.

Part of the Cogplexiti Systems & Decision Analysis line. Operationalizes the line’s alternative-futures and premortem analyses; its Tier 3 communication plugs into the convening architecture designed in Nobody’s Villain.

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.

