

42 The Systemic Review

Because systemic problems demand systemic solutions

Between the Boxes

A systemic reading of why bounded technical failures inside UK air-traffic control keeps producing nationwide aviation disruption

Theoretical basis: the laws and principles of systems theory, applied here as a diagnostic framework rather than as a source of opinion.

Prepared September 2026

On 8 September 2026, a technical fault inside NATS's Swanwick control centre halted normal air-traffic processing over southern England for roughly three hours. The fault itself was fixed within the afternoon. The disruption it set off was not: by the following day close to 2,000 flights had been cancelled or delayed at airports from Heathrow to Aberdeen, tens of thousands of passengers were stranded, and the Transport Secretary had summoned NATS's chief executive, Martin Rolfe, for the second time in just over a year. Reporting on the episode, including in the Guardian, noted that it was the third major NATS stoppage in three years, and that it came eleven weeks after the regulator had formally closed the last recommendation from its review of the previous one. This paper does not investigate what caused the technical fault; that is the proper business of NATS's own inquiry and the Civil Aviation Authority's review. It asks the systemic question those inquiries are not designed to answer: why does a bounded, hours-long failure in one part of UK air-traffic control keep producing disruption so far out of proportion to it, and where, if anywhere, does responsibility for that gap actually sit?

Executive Summary

This paper asks a bounded question: what do the laws and principles of systems theory reveal about why a several-hour technical failure inside UK air-traffic control repeatedly produces multi-day, nationwide aviation disruption, and about where the gap between cause and consequence actually sits? It is not a finding on what caused the 8 September 2026 fault, nor a verdict on any individual's conduct — those are matters for NATS's own investigation and the Civil Aviation Authority's review. It applies the framework consistently, naming the governing principle before stating what follows from it.

CENTRAL FINDING

UK aviation operates as one tightly interdependent system but is governed for resilience as several separate ones. The efficiency that makes each part of the network affordable to run is the same efficiency that has removed most of the spare capacity a large disturbance needs to be absorbed, and the agency, information and risk involved in a network-wide recovery sit predominantly inside individual organisations rather than at the level where the disruption actually occurs. Three major stoppages in three years, each followed by an investigation that improved the component involved, indicate that this structure — not any one organisation's reliability — is what the recurring pattern is actually diagnosing.

The findings that follow are set out in full in the next section. In summary: the disproportionate scale of the disruption is a property of the network's structure, not primarily of NATS's reliability; the efficiency that keeps aviation affordable

has removed the spare capacity that recovery depends on; nobody is currently positioned, resourced or authorised to manage a recovery at the level where it actually has to happen; the reform cycle triggered by each failure runs slower than the interval between failures, so the network is repeatedly disturbed again before it has settled; and because the structure generating the pattern has not itself changed, each investigation is well placed to fix the component that failed this time and poorly placed to prevent the network from amplifying the next one.

Findings

These findings are the headline output of the diagnosis. The technical analysis that follows names the specific systems principle behind each one and traces how it plays out in the current situation.

Finding 1. A relationship problem, not chiefly a NATS problem.

The disruption that followed the 8 September fault was not an output of NATS acting alone; it emerged from the interaction of air-traffic capacity, airline scheduling, airport slot constraints, and the positions of aircraft and crews once the network had been perturbed. British Airways alone reported close to 550 cancellations and disruption to more than 85,000 customers — most of it after the technical fault had already been fixed, because aircraft and crews had been left in the wrong places. Investigating NATS in isolation can establish what broke; it cannot by itself explain why the coupled network produced consequences several orders of magnitude larger than a three-hour fault.

Finding 2. Efficiency and resilience are in structural tension, and efficiency has been winning.

Aircraft, crews, airport slots and airspace are run at high utilisation because spare capacity looks like waste in normal operation. The same absence of slack is what limited how quickly the network could reconfigure once the fault hit — displaced aircraft and time-limited crews had few alternative options to draw on.

Finding 3. Agency sits at the level of the organisation, not at the level of the network.

NATS can manage air traffic, airlines can manage their aircraft and crews, and airports can manage their own capacity — and each does. None of them is individually positioned to manage the recovery of the network as a whole once

a disturbance has crossed organisational lines, and no separate function currently holds that role either.

Finding 4. The reform cycle is slower than the interval between failures — and the interval is shrinking.

The Civil Aviation Authority's independent review of the August 2023 failure took from October 2023 until 19 June 2026 to close its final recommendation — nearly two years and nine months. A materially similar failure occurred eleven weeks later. The interval between the first and second stoppages was about 23 months; between the second and third, about 13 months. The reform cycle is not just slower than the disturbance interval, it is losing ground to it.

Finding 5. Individually sound risk management has produced a high-risk whole.

NATS, airlines, airports and ground handlers each manage risk within their own domain, and each does so competently by its own measure. Nobody has assessed the combined effect of those separate choices on the network as a whole. The clearest instance is structural: the August 2023 failure occurred because NATS's primary and backup systems, built to the same specification, failed on the same fault simultaneously; reporting on the September 2026 event has raised the same structural question about the current system without yet confirming the answer.

Finding 6. The pattern is recurring because the structure that produces it has not changed.

Three major stoppages in three years have each been followed by investigation, findings and component-level correction. None of the three reviews is known to have treated the network's capacity to amplify a bounded failure as the object of the inquiry in its own right, which is one systemic reason a materially similar pattern of disruption has now recurred three times against three different technical triggers.

Technical Analysis

The remainder of this paper sets out the evidence and the systems-theory mechanism behind each finding above: the boundary of the system under examination, the situation on the ground, and the principle-by-principle diagnosis and its cross-cutting dynamics.

1. System Framing and Method

1.1 System Framing

Systems laws describe the behaviour of a defined system in relation to its environment and are only meaningful once that boundary is set explicitly.

System	UK air-traffic control and the wider operational aviation network it sits inside — NATS's air-traffic services together with the airlines, airports, ground handlers, aircraft and crews whose schedules and positions those services immediately affect
Key components	NATS (en-route and approach air-traffic services, including the Swanwick and Prestwick control centres); UK airlines and their aircraft, crews and schedules; UK airports (Heathrow, Gatwick, Stansted, Luton, City and the wider national network) and connecting European airports; ground handlers; the Civil Aviation Authority; the Department for Transport
Environment	Passenger demand and schedule density; commercial pressure on airlines and airports for high aircraft, crew and slot utilisation; the technical architecture of flight-data processing, including third-party supplier systems and European flight-plan data formats NATS must interoperate with; the reform cycle following each previous failure
Key dynamic	A gap between how tightly the operational network is coupled during normal running and how separately its constituent organisations are governed, resourced and held accountable for resilience — a gap that converts a bounded technical fault into network-wide, multi-day disruption

1.2 Method

Each section below identifies the governing principle first, states the mechanism by which it operates, and only then derives the implication for UK aviation, so that findings remain traceable to a named principle rather than appearing as free-standing opinion. Principles are drawn from the full set of systems laws, thinking patterns and miscellaneous systems ideas, checked directly against the reference material rather than relied on from memory, and reported only where they materially apply to this situation — not as an exhaustive checklist. Interaction between principles (“splicing”) is noted at the point each connection is found, rather than deferred to a single later section.

2. Conventional Model and Systemic View Compared

Before turning to the situation itself, it is worth setting the two approaches side by side, since the contrast is the organising question of this paper.

Conventional assumption	What the systemic view reveals
A NATS failure is a NATS problem: establish the technical cause, fix it, and hold NATS accountable.	The scale of the disruption is not an output of NATS acting alone; it emerges from the interaction of air-traffic capacity, airline scheduling, airport constraints and the positions of aircraft and crews (Principle of Emergence). Investigating NATS in isolation answers what broke, not why the coupled network behaved as it did.
The organisational boundary between NATS, airlines, airports and the regulator is the natural unit for managing resilience.	Legal and organisational boundaries need not correspond to the operational system that actually produces the emergent disruption (Law of Calling); the operational network is interdependent, the governance structure is partitioned by organisation, and that mismatch has not been the direct object of any review to date.
Unused aircraft, crew, airspace or slot capacity is inefficiency to be designed out.	The same spare capacity that looks wasteful in normal operation is the adaptive variety the network draws on to recover from a large disturbance (Viability Principle); the design question is not how much spare capacity is affordable but how much latent variety the system needs to remain viable under plausible disruption.

Every organisation managing its own operational risk competently adds up to a resilient network.	A component optimised for its own reliability has not thereby been optimised for its systemic context; a system built entirely from such components can still be high-risk as a whole (Adams' 3rd Law) — visible directly in the shared primary/backup design that failed simultaneously in 2023.
An investigation that assigns responsibility and closes recommendations is the same thing as regulating the network's resilience.	Accountability asks who should answer for what happened; regulation asks what mechanism keeps the system within acceptable states going forward. A system can have strong accountability and weak regulation at once (Law of Sufficient Complexity), and the recurrence of this pattern suggests that is close to the current position.

3. Situation Brief

NATS, the UK's monopoly provider of en-route air-traffic control, has now had three major system stoppages in just over three years. On 28 August 2023, a flight-planning system failure — triggered when two waypoints in different locations shared an identical name, a fault that caused both the primary and backup systems to enter failsafe mode simultaneously — disrupted the August bank holiday weekend, affecting more than 700,000 passengers (around 300,000 through cancellations, a further 95,000 through delays exceeding three hours, and 300,000 more through shorter delays) and costing airlines an estimated £75–100 million in compensation and recovery costs. The Civil Aviation Authority commissioned an independent review, chaired by Jeff Halliwell, which published an interim report in March 2024 and, in November 2024, 34 recommendations covering contingency planning, engineering staff allocation and disruption notification. The CAA formally closed the final recommendation on 19 June 2026.

On 30 July 2025, a radar-related fault at Swanwick halted departures for more than four hours and led to over 150 cancellations; the Transport Secretary, Heidi Alexander, summoned Martin Rolfe to account for it.

On 8 September 2026, eleven weeks after the CAA closed the last of its 2023-review recommendations, a fault in the flight-data processing system that converts European flight-plan data into UK airspace format — reportedly the

FPRSA-R system supplied by Frequentis AG — surfaced at Swanwick at around 13:46 local time; a fix was applied by roughly 16:40, and NATS declared the issue resolved at approximately 19:34. NATS ruled out a cyberattack, and Rolfe described the fault as “something different that we’ve never seen in 50 years of operation.” The disruption continued into a second day: across 8 and 9 September, close to 2,000 flights were affected — around 30% of scheduled departures on the first day and roughly 5% on the second — with cancellations and delays reported at Heathrow, Gatwick, Stansted, Luton, City, Edinburgh, Bristol, Manchester, Birmingham, Belfast, Glasgow and Aberdeen. British Airways alone reported close to 550 cancellations and disruption to more than 85,000 customers; Ryanair reported over 65,000 affected passengers.

Heidi Alexander summoned Rolfe for the second time in just over a year; NATS opened its own investigation, and the CAA is assessing what further action its own review should take once NATS reports. Ryanair’s chief executive, Michael O’Leary, publicly called for Rolfe’s removal, describing NATS as having “failed to invest in a resilient primary system” and “failed to invest in an effective backup system,” and characterised the episode as the fourth major NATS failure of 2026 alone — a count this review has not been able to independently verify against the “third major stoppage in three years” figure reported elsewhere, and flags as disputed rather than resolving in either direction.

Reporting on the September 2026 event has noted, without yet confirming it as the specific cause, that NATS’s primary and backup systems are understood to run substantially the same software logic — the same structural characteristic that allowed a single fault to disable both systems simultaneously in August 2023.

What is unverified or unknown: the specific technical trigger for the 8 September 2026 fault, which NATS had not disclosed at the time of writing; whether the primary/backup shared-logic vulnerability identified in 2023 is confirmed as a cause of the September 2026 event specifically, as opposed to a structural similarity between the two; the true count of “major” NATS failures in 2026, given the discrepancy between Ryanair’s public figure and the “third stoppage in three years” reported elsewhere; what NATS’s within-days internal

investigation and the CAA’s further review will conclude; and, by the nature of the question this paper asks, whether either inquiry will treat the network’s amplification of a bounded failure — rather than the failure itself — as part of its remit.

4. Systemic Diagnosis

The table below sets out the principles that materially apply to this situation, the mechanism by which each operates, and the effect it implies for UK aviation. Discussion of each, and the connections between them, follows.

Principle & Core Statement	Predicted / Observed Systemic Effect	Aviation-System Impact
Principle of Emergence <i>The whole is more than the sum of its parts; emergence is a property of a system, not a property of any part of it taken alone.</i>	Thousands of cancellations are not an output of NATS’s technical fault by itself; they emerge from the interaction of air-traffic capacity, airline schedules, airport slot constraints, and the positions of aircraft and crews once the network has been disturbed.	Investigating NATS in isolation can explain what broke; it cannot, by itself, explain why the coupled network produced disruption several orders of magnitude larger than the three-hour fault that triggered it.
Law of Calling + Law of Crossing <i>Drawing a boundary is a judgement, not the discovery of a natural line; crossing a boundary is itself a change of state.</i>	The operational aviation network — NATS, airlines, airports, aircraft and crews interacting continuously — is treated, for governance and accountability purposes, as a set of separately bounded organisations, each with its own regulator, board and remit.	The governance boundary is drawn at exactly the places the operational system is most connected; every hand-off between organisations during a disturbance is a boundary crossing that can lose information, authority or time, and none of those crossings is currently anyone’s named responsibility.
Viability Principle A <i>A system’s viability depends on balancing autonomy with cohesion and stability with change; because efficiency is output divided by resource, a highly efficient system has little spare resource, and adaptability depends on having some.</i>	Tightly rostered crews, high aircraft utilisation and fully allocated airport slots and airspace keep the network cheap to run in normal conditions and leave it with little capacity to reconfigure when conditions change sharply.	What looks like inefficiency in normal operation — an unused aircraft, a spare crew, an uncommitted slot — is the adaptive variety the network draws on to recover from a large disturbance; the design question is not how much spare capacity is

		affordable but how much latent variety the network needs to remain viable, and that question does not currently have an owner.
Law of Requisite Variety <i>How well a system manages depends on how well it matches the variety it faces; only variety can destroy variety.</i>	Normal operation deliberately narrows the network into a small range of states. A major disruption multiplies that range almost instantly — aircraft in the wrong places, crews against duty-time limits, airport capacity constrained, passenger connections broken — while the network's capacity to respond stays roughly fixed.	If recovery capability does not expand when disruption variety expands, control deteriorates by construction, independent of how competently any single organisation manages its own part of the response — consistent with disruption continuing for a full day after the technical fault itself was fixed.
Power Structuration Theorem <i>A nested system has optimal agency when the need for agency at the system level is balanced against the need for agency of its constituent sub-systems.</i>	During normal operations, NATS, airlines and airports each hold substantial agency over their own domain. Once a disturbance crosses organisational lines, an emergent problem exists at the level of the network — but agency remains distributed among the same organisations.	Nobody is positioned to optimise the recovery of the whole; each actor optimises its own recovery, and those separately rational choices can pull against each other — an airline's fastest recovery option can add demand to an airport, and an airport's capacity constraint can limit what an airline can do.
Conant–Ashby Theorem + Redundancy of Potential Command Principle <i>Every good regulator of a system must be a model of that system; and effective action in a complex situation depends on bringing together an adequate concatenation of information, not simply on having information available somewhere in the network.</i>	NATS models air traffic; airlines model their own fleets, crews and passengers; airports model their own gates and ground state. No party is known to hold a working model of how a loss of air-traffic capacity propagates through aircraft positioning, crew legality, airport capacity and airline recovery decisions over the following hours and days.	Recovery decisions are made by parties who each see part of the picture and none of whom sees the whole; that is a structural description of why plausible, locally rational decisions during recovery can still add up to a slower, more prolonged network-wide recovery.
Adams' 3rd Law <i>A system composed entirely of components selected for their own low risk will itself be a high-risk system, because components optimised for internal reliability have not thereby been</i>	NATS, airlines, airports, ground handlers and their regulator each manage risk within their own remit, competently by their own measure. The clearest evidence that this does not guarantee low risk for the network as a	A resilience assessment confined to each organisation's own risk register will not surface this class of failure, because the risk exists in the combination, not in any individual component; it requires a whole-

<i>optimised for their systemic context.</i>	whole is structural: NATS's primary and backup systems, each independently specified to be reliable, are understood to run substantially the same software logic — which is why a single fault disabled both simultaneously in August 2023, and why the same structural question has been raised about the September 2026 event.	network risk assessment run alongside, not instead of, each organisation's own.
Law of Reciprocity of Connections <i>If A connects to B, B also connects to A, even where institutional accounting recognises only one direction; relationships are loops, not one-way chains.</i>	A loss of air-traffic capacity affects an airline's schedule; the airline's response affects an airport; the airport's state affects other airlines; those airlines' responses alter demand back on NATS; the changed network state alters the next round of recovery decisions.	Accident investigation that proceeds organisation by organisation — what did NATS do, what did the airline do, what did the airport do — can answer each question correctly and still miss why the coupled system behaved as it did, because the causal structure is a loop rather than a chain running one way from NATS outward.
Law of Sufficient Complexity <i>A complex system constitutes its own simplest behavioural description; if its structure remains unchanged, its characteristic behaviour will tend to reassert itself, whatever adjustments are made to its inputs.</i>	Three investigations — into the August 2023, July 2025 and September 2026 events — have each been followed by findings and component-level correction, without the underlying structure connecting NATS, airlines and airports having been treated as the object of the inquiry.	Repeated investigation can improve the specific component that failed each time while leaving the network's capacity to amplify a bounded failure intact; that is consistent with a materially similar pattern of disruption recurring three times against three different technical triggers.
Relaxation Time Principle + Structural Viability Theorem <i>A system repeatedly shocked at intervals shorter than its recovery time may never fully stabilise; and where a system's levels need to change at different rates, viability depends on that differential being kept in balance rather than allowed to diverge.</i>	The Civil Aviation Authority's review of the August 2023 failure took from October 2023 to 19 June 2026 — nearly two years and nine months — to close its final recommendation. The interval between the three major stoppages on record has fallen from about 23 months to about 13 months. The institutional level of the system, where reform happens, is moving markedly slower than the operational level, where disturbances recur — and the gap is widening, not narrowing.	A reform cycle that takes longer to complete than the interval between the disturbances it responds to cannot, by construction, keep the network ahead of the next one; each review is still in progress, or barely closed, when the next disturbance arrives.

Root Structuration Theorem <i>Complexity is minimised when a system is structured so the number of its sub-systems approaches the square root of the number of its elements, rather than allowing every element to interact directly with every other.</i>	The current architecture has grown a high density of direct interdependencies — aircraft, crews, slots and airport capacity all coupled to NATS’s air-traffic state and to each other — without a corresponding discipline of clustering related functions and limiting the interfaces between them.	The available lever is not simply more redundancy inside NATS, but redesigning which dependencies are allowed to propagate disruption directly across organisational lines and which are deliberately contained, so that fewer elements of the network need to interact with everything else during a disturbance.
Double Bind (miscellany) <i>A situation constrained by two mutually contradictory imperatives, where satisfying either one means failing the other.</i>	Each organisation faces a subsystem imperative to use expensive aviation resources as efficiently as possible, while the network as a whole needs enough unused adaptive capacity to absorb rare, large disturbances. The organisation that carries the cost of holding that capacity is not the one that mainly receives the benefit when it is needed.	Each actor can rationally reduce its own spare capacity; the aggregate effect can be a less resilient network, without any single decision along the way looking unreasonable from inside the organisation that made it.
Darkness Principle <i>No system can be known completely by any observer or sub-system; irreducible unknowns are structural, not incidental.</i>	Even a well-designed whole-network recovery function would not achieve complete knowledge of how a given disturbance will propagate; the technical trigger of the September 2026 fault itself remained undisclosed at the time of writing, and the discrepancy between “third stoppage in three years” and Ryanair’s “fourth failure in 2026” has not been resolved by this review.	This bounds what any structural fix can promise. The realistic aim of a whole-network capability is not to make disruption predictable in advance, but to hold enough of a model, enough information and enough authority at the network level to respond well when the unpredictable happens — a different, more achievable standard than eliminating the uncertainty itself.

4.1 An emergent property, not a component failure

The Principle of Emergence holds that the whole is more than the sum of its parts, and that emergence is a property of the system rather than of any part taken alone. The 8 September fault lasted about three hours before NATS’s own fix took effect. The disruption did not track that timeline: it continued into a second day and touched nearly every major UK airport, because the aircraft, crews, slots and connections displaced during those three hours had already

begun interacting with the wider network before the fault was fixed. British Airways' own figures are direct evidence of this: the bulk of its cancellations and the disruption to more than 85,000 customers accumulated after the technical cause had already been resolved, driven by aircraft and crews left in the wrong places. The systemic question this paper asks — why a bounded failure produces disproportionate consequence — only becomes visible once the object of analysis is widened from NATS to the network NATS sits inside; it is invisible to an inquiry scoped to NATS alone, however thorough that inquiry is.

4.2 Boundaries drawn where the organisation chart is, not where the disruption is

The Law of Calling holds that drawing a boundary is a judgement, not the discovery of a natural line, and warns specifically against accepting a legal or organisational boundary as though it were the true boundary of the system generating a problem; its partner, the Law of Crossing, adds that moving across a boundary is itself a change of state. Applied here, there are two different maps of the same network. Operationally, NATS, airlines, airports, aircraft and crews are continuously interdependent. Institutionally, responsibility is partitioned into NATS, individual airlines, individual airports, a regulator and a government department, each separately accountable. This connects directly to §4.1: an emergent, network-level property is being investigated through a set of boundaries drawn at the organisational rather than the operational level, so each hand-off between organisations during a disturbance — a delayed aircraft crossing into another airline's schedule, a displaced crew crossing into another airport's ground operation — is a boundary crossing that can lose time, information or authority, and none of those crossings is currently anyone's named responsibility.

4.3 Efficiency and resilience in structural tension

The Viability Principle states that a system's viability depends on balancing autonomy with cohesion and stability with change, and adds a sharper mechanism: because efficiency is defined as output divided by resource, a

highly efficient system has very little spare resource, and the capacity to adapt depends on having some. UK aviation is under continuous commercial pressure to run aircraft, crews, airport slots and airspace at high utilisation. That pressure is entirely rational for each organisation — it is the same autonomy the Law of Calling in §4.2 shows is organised around individual organisations — but its system-level effect is to leave the network with little latent capacity when a disturbance arrives. What ordinary managerial language classifies as inefficiency — an unused aircraft, a spare crew, an uncommitted slot — is, in this framing, adaptive variety, and the relevant design question is not how much spare capacity is affordable but how much latent variety the network needs to remain viable under plausible disturbance.

4.4 A network whose recovery variety has not grown with its disruption variety

The Law of Requisite Variety states that control of a system requires variety at least equal to the variety the system presents; only variety can destroy variety. Normal aviation operation is deliberately narrow — scheduled routes, allocated slots, rostered crews. A major disruption multiplies the number of states the network can be in almost instantly: aircraft in the wrong places, crews against duty-time limits, airport capacity constrained, passenger connections broken, and each airline forming its own recovery priorities. This compounds directly on §4.3: it is precisely the spare capacity that efficiency has removed that would have supplied the additional response variety the network now needs. If the network's capacity to respond does not expand when disruption variety expands, Ashby's law says that control must deteriorate — independent of how competently any single organisation manages its own part of the response, and consistent with disruption continuing for a full day after the technical fault itself was fixed.

4.5 Agency at the level of the organisation, not the network

The Power Structuration Theorem holds that a nested system has optimal agency when the need for agency at the system level is balanced against the need for agency of its constituent sub-systems — in plain terms, that decision-

making power should exist at the level where the complexity requiring action actually exists. During ordinary operations this balance broadly holds: NATS manages air traffic, airlines manage aircraft and crews, airports manage their own capacity, and each has substantial agency over its own domain. Once a disturbance crosses organisational lines — the boundary crossings identified in §4.2 — an emergent problem exists at the level of the network, but agency remains where it was, distributed among the same organisations. Nobody is positioned to optimise the recovery of the whole; each actor optimises its own recovery, and those separately rational choices can conflict — an airline’s fastest recovery option can add demand to an airport, an airport’s capacity constraint can limit what an airline can do, and NATS’s safest traffic-management response can prolong airline disruption. None of this requires an incompetent actor anywhere in the network; it is a description of where authority sits relative to where the problem sits.

4.6 The model-and-information gap, and the regulator’s own regulator problem

The Conant–Ashby Theorem — the good-regulator theorem — states that every effective regulator of a system must hold an adequate model of that system, where “regulator” means anything trying to guide, control or direct a system, not only a statutory regulator. NATS models air traffic; airlines model their own fleets, crews and passengers; airports model their own gates and ground state. No party is known to hold a working model of how a loss of air-traffic capacity propagates through aircraft positioning, crew legality, airport capacity and airline recovery decisions over the following hours and days. The Redundancy of Potential Command Principle sharpens why that matters in practice: effective action in a complex decision network depends on bringing together an adequate concatenation of information, not simply on the information existing somewhere in the network — and, in the framework’s own corollary, organisations tend to take the decisions their available information enables them to take. Each actor in §4.5 is therefore not just short of authority for network-level recovery; it is also short of the combined information a network-level decision would need, which is a second, independent reason recovery

decisions can be locally sound and collectively slow. A further point follows from applying the same theorem recursively: any coordinating function created to close this gap is itself a system, and Ashby's law applies to it too — a coordinating body with authority but no current information cannot act effectively, and one with information but the wrong model of how disruption propagates will still regulate badly, however senior it is.

4.7 Individually rational risk management, collectively high-risk

Adams' 3rd Law states that a system composed entirely of components selected for their own low risk will itself be a high-risk system, because components optimised for internal reliability have not thereby been optimised for their systemic context. NATS, airlines, airports, ground handlers and their regulator each manage risk competently within their own remit. The clearest evidence that this does not guarantee a low-risk network is structural, not hypothetical: the August 2023 failure occurred because NATS's primary and backup systems, each independently specified to be reliable, ran substantially the same software logic and failed on the same fault simultaneously. Reporting on the September 2026 event has raised the same structural question about the current system without yet confirming the answer — this review has not independently verified whether shared logic caused the second event, and states that as an open question rather than a finding. Either way, a resilience assessment confined to each organisation's own risk register, of the kind each party in §4.5 and §4.6 is well placed to produce, will not surface this class of risk, because the risk exists in the combination rather than in any individual component.

4.8 Loops, not chains

The Law of Reciprocity of Connections holds that if A connects to B, B also connects to A, even where institutional accounting recognises only one direction — relationships are loops, not one-way chains. A loss of air-traffic capacity affects an airline's schedule; the airline's response affects an airport; the airport's state affects other airlines; those airlines' responses alter demand back on NATS; the changed network state alters the next round of recovery

decisions. This is the causal structure that §4.1's emergent disruption actually runs on. Accident investigation that proceeds organisation by organisation — what did NATS do, what did the airline do, what did the airport do — can answer each of those questions correctly, and still miss why the coupled system behaved as it did, because the questions assume a chain running one way from NATS outward when the actual structure is closer to a loop with no single starting point.

4.9 Why component-level fixes keep producing the same pattern

The Law of Sufficient Complexity states that a complex system constitutes its own simplest behavioural description: if the system's structure remains unchanged, its characteristic behaviour will tend to reassert itself, whatever adjustments are made to its inputs. Three investigations — into the August 2023, July 2025 and September 2026 events — have each been followed by findings and component-level correction: the Halliwell review's 34 recommendations addressed contingency planning, engineering staff allocation and disruption notification, all sound measures, all aimed at NATS and its immediate procedures. None is known to have treated the structure connecting NATS, airlines and airports — the boundary mismatch of §4.2, the efficiency–variety tension of §4.3 and §4.4, the agency gap of §4.5 — as the object of the inquiry in its own right. Repeated investigation can therefore improve the specific component that failed each time while leaving the network's capacity to amplify a bounded failure intact, which is consistent with a materially similar pattern of disruption recurring three times against three different technical triggers.

4.10 A reform cycle slower than the disturbance interval — and losing ground

The Relaxation Time Principle warns that a system repeatedly shocked at intervals shorter than its recovery time may never fully stabilise. The Structural Viability Theorem adds that where a system's different levels need to change at different rates, viability depends on keeping that differential in balance rather

than letting it diverge. The Civil Aviation Authority's review of the August 2023 failure took from October 2023 to 19 June 2026 — nearly two years and nine months — to close its final recommendation. Measured against the record of actual disturbances, that reform cycle is not merely slow: the interval between the first and second major stoppages was about 23 months, and between the second and third about 13 months. The institutional level of the system, where reform happens, is moving markedly slower than the operational level, where disturbances recur, and the gap between the two rates is widening rather than closing. This is the Structural-Viability reading of §4.9: it is not only that the structure has not changed, but that the level responsible for changing it is structurally unable to move at the pace the level generating disturbances now requires.

4.11 Too many direct interdependencies, not enough clustering

The Root Structuration Theorem states that complexity is minimised when a system is structured so the number of its sub-systems approaches the square root of the number of its elements, rather than allowing every element to interact directly with every other — it is in effect the antidote to the kind of complexity-driven instability described in §4.4. The current architecture has grown a high density of direct interdependencies — aircraft, crews, slots and airport capacity all coupled to NATS's air-traffic state and to each other — without a corresponding discipline of clustering related functions and limiting the interfaces between them. The available lever this principle points to is not simply more redundancy inside NATS — which, per §4.7, can itself replicate rather than reduce risk if built to the same specification twice — but redesigning which dependencies are allowed to propagate disruption directly across organisational lines and which are deliberately contained, so that fewer elements of the network need to interact with everything else during a disturbance.

4.12 A double bind nobody is pricing

The concept of the double bind describes a situation constrained by two mutually contradictory imperatives, where satisfying either one means failing the other, often operating at two different structural levels. Each organisation in UK aviation faces a subsystem imperative to use expensive resources — aircraft, crews, slots, airspace — as efficiently as possible, which is the pressure identified in §4.3. The network as a whole needs enough unused adaptive capacity to absorb rare, large disturbances, which is the variety identified in §4.4. The organisation that would carry the cost of holding that capacity is not the one that mainly receives the benefit when it is needed — the benefit accrues to the network, including competitors, while the cost sits on whichever organisation's balance sheet holds the spare aircraft or crew. Each actor can therefore rationally reduce its own spare capacity; the aggregate effect can be a less resilient network, without any single decision along the way looking unreasonable from inside the organisation that made it.

4.13 What can't be known, and what that does and doesn't excuse

The Darkness Principle holds that no system can be known completely by any observer or sub-system, and that this is a structural limit, not merely a current gap in effort. It applies here in two directions at once. First, it bounds this review's own claims: the specific technical trigger for the 8 September fault had not been disclosed at the time of writing, and this paper has been unable to resolve the discrepancy between the widely reported “third major stoppage in three years” and Ryanair's public claim of a “fourth major failure in 2026” — both are reported here as stated by their sources, not as independently verified facts. Second, and more importantly, it bounds what any structural fix to the problem diagnosed in §4.1 through §4.12 can honestly promise. A whole-network recovery function, however well designed, will not make disruption fully predictable in advance — the network is too complex and too dynamic for that. Its realistic purpose is not to eliminate uncertainty but to hold enough of a model (§4.6), enough information (§4.6), enough authority (§4.5) and enough response variety (§4.4) at the level where a disturbance actually occurs to

respond well when the unpredictable happens — a different, and more achievable, standard than the one conventional accountability exercises implicitly promise each time they name a cause and a responsible party.

5. Cross-Cutting Dynamics

Read together, the principles above converge on a single structural account rather than a list of independent findings. An emergent, network-level disruption (Principle of Emergence) is being managed through boundaries drawn at the organisational rather than the operational level (Law of Calling, Law of Crossing), inside a network whose efficiency has removed the spare capacity that adaptation depends on (Viability Principle) precisely as the variety it needs to manage a large disturbance multiplies (Law of Requisite Variety). Agency to respond to that network-level problem remains distributed at the organisational level where it was always held (Power Structuration Theorem), compounded by the fact that no party holds either the model or the combined information a network-level decision would need (Conant–Ashby Theorem, Redundancy of Potential Command Principle) — a gap that would recur even inside any new coordinating body created to close it, since that body would itself be a system subject to the same requirements. Each organisation's individually rational risk management has not been assessed in combination (Adams' 3rd Law), the causal structure connecting them runs in loops rather than one-way chains (Law of Reciprocity of Connections), and because none of the three post-incident reviews has treated this structure as its object, each investigation has improved the failed component while leaving the amplifying structure intact (Law of Sufficient Complexity) — a structure whose own reform cycle is now moving more slowly than the interval between the disturbances it responds to, and losing ground (Relaxation Time Principle, Structural Viability Theorem). Underneath all of this sits an architecture with more direct interdependencies than deliberate clustering (Root Structuration Theorem) and an incentive structure in which the cost of resilience and the benefit of it fall on different actors (Double Bind). None of this is offered as a complete or certain account — the Darkness Principle applies to this diagnosis as much as

to the system it describes — but it is a coherent one: the pattern recurring across three stoppages in three years is best read not as three unrelated technical failures but as one structural gap expressing itself three times.

6. Indicative Interventions

The diagnosis above points to specific design moves, not just a diagnosis of what is wrong. Each is traceable to a principle already named, treated as a systemic design move rather than an implementation instruction — the Darkness Principle (§4.13) means none of them can be offered as a guarantee, only as a way of holding more model, information, authority and variety at the level where disruption actually occurs.

1. Build and maintain a whole-network disruption model.

Not a model of NATS, or of any one airline or airport, but one spanning air-traffic capacity, airline scheduling, aircraft and crew positioning, and airport ground state, tested against scenarios before a live incident rather than assembled only after one (Conant–Ashby Theorem, §4.6).

2. Create a standing cross-organisational recovery function with a defined activation threshold.

Not a permanent central command — the Viability Principle warns against that — but a function whose authority expands once a disturbance crosses a defined scale, so that agency exists at the network level precisely when the complexity requiring it moves there (Power Structuration Theorem, §4.5).

3. Design the information architecture backwards from the recovery decisions that actually have to be made.

Starting from “what does a network-level recovery decision need to know” rather than from what each organisation is currently willing to share, so the relevant concatenation of information can be assembled quickly when it is needed (Redundancy of Potential Command Principle, §4.6).

4. Protect some current “inefficiency” as adaptive slack rather than treating it as a cost to be designed out.

Spare aircraft, crew duty hours or slot capacity held explicitly for disturbance absorption, distinguished from genuine waste rather than collapsed into the same efficiency metric (Viability Principle, §4.3; Law of Requisite Variety, §4.4).

5. Run a whole-network risk assessment alongside, not instead of, each organisation’s own.

Testing explicitly for correlated failure in shared-specification components — the kind of primary/backup design that failed simultaneously in 2023 — rather than relying on each organisation’s risk register to add up to network-level safety (Adams’ 3rd Law, §4.7).

6. Reduce unnecessary direct interdependencies rather than adding redundancy inside NATS alone.

Clustering related functions and limiting the interfaces between them, so fewer elements of the network need to interact directly with everything else during a disturbance (Root Structuration Theorem, §4.11).

7. Track network recovery time, not technical fault duration, as the headline resilience metric.

Reporting publicly how long the whole network took to return to normal operation after each incident, not only how quickly the originating technical fault was fixed (Relaxation Time Principle, §4.10).

8. Give the reform cycle a pace target tied to the empirical interval between disturbances.

The Civil Aviation Authority’s review of the 2023 failure took nearly two years and nine months to close; the interval between disturbances has fallen to about 13 months. A reform cycle that does not commit to closing faster than that interval will keep losing ground to it (Structural Viability Theorem, §4.10).

9. Price the double bind explicitly.

Have whoever mainly benefits from network-wide resilience — passengers, the wider network, government — share the cost of the spare capacity that produces it, rather than leaving that cost on individual organisations' balance sheets while the benefit accrues elsewhere (Double Bind, §4.12).

Conclusion

Applied consistently, the laws and principles of systems theory suggest that the live debate prompted by each NATS failure — what went wrong this time, and who is responsible — is aimed at the wrong level to prevent the next one. The more consequential questions are whether the network's efficiency has removed capacity that its resilience depends on, whether any function is positioned to model, sense and act on disruption at the level where it actually occurs, and whether the reform cycle triggered by each failure can move faster than the interval between failures rather than slower. None of these are answered by establishing what broke inside NATS on 8 September 2026, however necessary that investigation is on its own terms, and all of them recur across every stoppage examined here — from the shared primary/backup logic of August 2023 to the multi-day recovery of September 2026. The practical implication is that an investigation confined to the organisation where a fault happened to occur, however thorough, is unlikely to resolve a vulnerability this framework shows to sit between the organisations rather than inside any one of them.

Appendix A — Glossary of Systems Principles Referenced

Principle	Definition
Principle of Emergence	The whole is more than the sum of its parts; emergence is a property of a system that is not a property of any of its parts on their own.
Law of Calling	The act of making a distinction is the act of drawing a boundary and defining a system; difference creates boundaries and boundaries create difference.
Law of Crossing	Crossing a system boundary is itself a change of state, for both what crosses and the observer.
Viability Principle	A system's viability depends on balancing the autonomy of its sub-systems against the cohesion of the whole, and stability against adaptation, over time.
Law of Requisite Variety	The control achievable over a system is limited by the variety of the regulator relative to the variety of the system being regulated; only variety can destroy variety.
Power Structuration Theorem	A nested system has optimal agency when the need for agency at the system level is balanced against the need for agency of its constituent sub-systems.
Conant–Ashby Theorem	Every good regulator of a system must be a model of that system; the ability to act effectively on any situation depends on the adequacy of your model of it.
Redundancy of Potential Command Principle	Effective action in a complex decision network depends on bringing together an adequate concatenation of information.
Adams' 3rd Law	A system composed entirely of the lowest-risk components available will itself be a high-risk system.
Law of Reciprocity of Connections	If A connects to B, then B also connects to A, even where institutional accounting recognises only one direction.
Law of Sufficient Complexity	A complex system constitutes its own simplest behavioural description; if its structure remains unchanged, its characteristic behaviour will tend to reassert itself, whatever adjustments are made to its inputs.

Relaxation Time Principle	A system can only remain stable if its relaxation time is shorter than the average interval between disturbances.
Structural Viability Theorem	A system has optimal viability when its rate of change relative to its environment is similar to that of its sub-systems.
Root Structuration Theorem	Complexity is minimised when a system is structured so the number of its sub-systems approaches the square root of the number of its elements.
Double Bind (miscellany)	A situation constrained by two mutually contradictory imperatives, where satisfying either one means failing the other.
Darkness Principle	No system can be known completely by any observer or sub-system; irreducible unknowns are structural, not incidental.

References

Sources consulted for the factual grounding of the Situation Brief (Technical Analysis §3). Theoretical statements throughout the paper are drawn from the systems-theory reference framework and are not separately footnoted.

1. The Guardian, “UK flight cancellations: Nats, Martin Rolfe, Heidi Alexander”, 9 September 2026. <https://www.theguardian.com/world/2026/sep/09/uk-flight-cancellations-nats-martin-rolfe-heidi-alexander>
2. Tech Times, “NATS Grounds Nearly 1,000 Flights in Third Swanwick Failure Since 2023 Reforms”. <https://www.techtimes.com/articles/327036/20260908/nats-grounds-nearly-1000-flights-third-swanwick-failure-since-2023-reforms.htm>
3. IBTimes UK, “NATS’ Glitch Disrupts More Than 1,000 Flights as Heathrow, Gatwick Chaos Echoes 2023’s £100M Disaster”. <https://www.ibtimes.co.uk/uk-air-traffic-control-glitch-disrupts-flights-1818653>
4. HNGN, “Third NATS failure in 3 years disrupts 2,000 UK flights”. <https://www.hngn.com/articles/273154/20260909/uk-air-traffic-failure-hits-2000-flights-minister-summons-nats-chief.htm>
5. UK Civil Aviation Authority, “Regulator publishes progress report on Independent Review into August 2023 NATS Flight Planning System Failure”. <https://www.caa.co.uk/newsroom/news/regulator-publishes-progress-report-on-independent-review-into-august-2023-nats-flight-planning-system-failure/>
6. Investigate / Ryanair Holdings, “Ryanair Calls For NATS CEO Rolfe To Be Sacked”, 9 September 2026. <https://www.investigate.co.uk/announcement/rns/ryanair-holdings-cdi---orya/ryanair-calls-for-nats-ceo-rolfe-to-be-sacked/9763403>
7. Travel PR News, “British Airways cancellations near 550 after NATS outage disrupts more than 85,000 customers”. <https://travelprnews.com/british-airways-cancellations-near-550-after-nats-outage-disrupts-more-than-85000-customers/travel-press-release/2026/09/09/>