When "Green" Systems Fail and Human Sovereignty Must Prevail

Case Study: LaGuardia Runway Incursion

Prepared by: Jose A. Cedeño, Principal Advisor
Organization: Cofresí Consulting Services
Date: March 26, 2026
Classification: Confidential — Advisory Distribution Only

I. Executive Summary

On the night of March 22, 2026, two pilots reported for duty, completed their checklists, and trusted the system around them to keep them safe. Antoine Forest was 30 years old. MacKenzie Gunther was 38. At 11:37 PM, Air Canada Express Flight 8646 collided with a Port Authority fire rescue vehicle on Runway 4 at LaGuardia Airport. Both pilots were killed. Forty-one others were hospitalized. Families were changed forever in a matter of seconds.

What makes this tragedy particularly profound is that no single person that night set out to cause harm. The controllers were doing their jobs under conditions the system placed them in. The pilots were doing everything right. The fire crew was responding to an emergency. Every human in that operational environment was trying to do the right thing — and the architecture around them failed to give any of them the information they needed, at the moment they needed it, to prevent what happened next.

This case study is written in their memory and in the memory of every person who has been failed by a system that reported green when the reality was anything but. It applies the Sovereign Human Layer framework to the LaGuardia incident across three dimensions: the convergence of failure conditions that made the accident possible; how VStrike deployed as a prevention architecture would have intervened at each failure point; and the advanced vision for an AI-powered Cognitive Fusion Engine that transforms situational awareness from a human memory burden into a real-time, synthesized, anticipatory capability. Its purpose is not to assign blame. It is to ensure that what was preventable here becomes structurally impossible to repeat.

> *VStrike does not merely witness the accident that almost happened. It architecturally removes the conditions that made it possible.*

II. The Incident: A Convergence of Compounding Failures

The collision was not the result of a single catastrophic error. It was the predictable outcome of multiple independent safety layers each degrading simultaneously — a pattern the NTSB has identified repeatedly in its most significant investigations.

Failure Layer	NTSB Finding
Sensor Blindness	The fire truck lacked a transponder. The ASDE-X ground radar failed to generate a collision alert because merging vehicle tracks near the

Failure Layer	NTSB Finding
	runway prevented creation of a high-confidence track. The truck was physically present but digitally invisible.
Human Cognitive Overload	Both air traffic controllers were simultaneously managing two positions each during a midnight shift. Under reduced staffing, one controller held responsibility for active runway operations, ground traffic, and departure information concurrently.
Environmental Degradation	Heavy rain and visibility reduced to 3 miles compounded cognitive load and limited direct visual observation of the airfield at the moment of highest operational risk.
Communication Failure	A critical radio transmission was stepped on by a simultaneous transmission in the final seconds. The stop directive may never have been received by the fire truck crew.
Purely Reactive Architecture	Every deployed safety tool — ASDE-X, radio protocols, visual observation — was designed to detect a conflict already developing. None were designed to prevent one from becoming possible.

> **In fewer than three minutes from landing clearance to collision, five independent failure conditions converged without triggering a single preventive intervention.**

III. The Three Root Conditions

A. The Poisoned Operational Baseline

The ASDE-X system reported runway 4 as clear. From the perspective of every digital instrument available to the controllers that night, that assessment was accurate — because the system had no way to know what it could not see. The truck was invisible not due to malfunction, but by design: a vehicle operating in a safety-critical environment without the digital identifier required for the monitoring system to register its existence.

This is the poisoned baseline in its most literal form. The system's internal truth — the runway is clear — had diverged completely from the physical truth — a 25-tonne fire vehicle is crossing the landing path. The system was not wrong because it malfunctioned. It was wrong because its sensor requirements were not met, and it had no mechanism to recognize they had not been met.

> *A system that cannot detect its own blind spots will report those blind spots as confirmed safe zones.*

B. Degraded Human Cognitive Readiness

The NTSB Chair explicitly named the midnight shift as a concern the board has raised repeatedly — including as a contributing factor in the Reagan National collision that killed 67 people in January 2025. That night, both controllers were managing two positions simultaneously. The same configuration. The same systemic vulnerability. A different airport. The same outcome.

The doctrine of Out-of-Band Human Sovereignty assumes that the designated human authority is cognitively capable of exercising that authority when called upon. A controller managing four concurrent responsibilities during a midnight shift in degraded weather has had their sovereign

override capacity systematically eroded before the incident begins. Human override authority is not a binary state. It exists on a continuum governed by cognitive load, fatigue, environmental stress, and task saturation.

> *Cognitive Readiness is the fifth pillar of the Sovereign Human Layer framework. Fatigue is not a personal failing — it is a system vulnerability that must be monitored and managed with the same rigor as a firewall rule.*

C. A Purely Reactive Safety Architecture

Every safety layer deployed at LaGuardia that night was designed to respond to a conflict already in progress. Each intervention required a human to observe, process, and act — under time pressure, under cognitive load, with incomplete information. The predictive conflict prevention model shifts this fundamentally: the system removes the possibility of a conflict before it can develop, rather than racing to catch one already unfolding.

IV. The VStrike Intervention: Prevention Architecture

VStrike's role is not to add another alarm to an already saturated environment. It serves as the Out-of-Band Sovereign Visualization platform through which human operators maintain independent situational awareness — and through which predictive conflict prevention is enforced before conflicts enter the physical world.

Failure Condition	VStrike Prevention Architecture
Truck invisible to ASDE-X due to missing transponder	VStrike ingests asynchronous sensor streams including thermal cameras and ground radar independent of transponder data. The truck is rendered in the 3D canvas regardless of digital identifier status.
No predictive conflict detection — system only detects converging tracks	VStrike's Predictive Runway Sovereignty Layer ingests the live flight schedule as a primary data feed. Any active landing window activates a temporal exclusion zone. Ground crossing requests trigger an automatic hold state.
Single-channel voice confirmation only — radio transmission stepped on	Hold state requires dual confirmation: voice clearance AND electronic acknowledgment through VStrike. If either channel fails, the hold remains. The truck cannot be cleared by radio alone.
Controller cognitive load at four concurrent positions during midnight shift	VStrike behavioral analytics monitor operator interaction patterns for response time degradation and task saturation. A human readiness alert is generated when cognitive load indicators breach defined thresholds.
No independent view of converging flight and ground paths	VStrike renders a live 3D reconstruction of converging vectors the moment a crossing is requested against an active runway window — before the conflict enters its terminal phase.

> *The controller's role shifts from constant vigilance over every possible conflict to deliberate confirmation of pre-flagged crossing requests. That is a cognitive load a fatigued human can reliably carry.*

V. The Doctrine of the Witness

The LaGuardia tragedy proves that automation cannot certify its own integrity, and humans cannot exercise sovereignty over systems that have consumed their capacity to observe independently. Three lessons emerge that apply across aviation, cybersecurity, critical infrastructure, and any complex environment governed by automated systems.

Lesson 1: Digital Telemetry is a Suggestion, Not a Verdict

The ASDE-X report that the runway was clear was an accurate report of what the system could see. The gap between what the system could see and what was physically true is not a bug — it is a structural feature of any monitoring architecture that depends on digital identifiers to register physical reality. VStrike closes this gap by treating physical presence as sovereign over digital registration.

Lesson 2: Human Readiness is a System Variable, Not a Human Virtue

The controllers were not negligent. They were operating under a workload configuration the NTSB had already identified as dangerous. Cognitive readiness is a measurable, manageable variable. It must be treated as a security control with defined thresholds, active monitoring, and automatic escalation protocols when those thresholds are breached.

Lesson 3: Reactive Architectures Guarantee Reactive Outcomes

A system designed only to detect conflicts in progress will always be racing the clock. The predictive conflict prevention model — removing the possibility of a conflict before it can develop — is not a feature enhancement. It is a philosophical shift in safety architecture: not what do we do when this goes wrong, but how do we make this structurally unable to go wrong.

VI. The Cognitive Fusion Engine: AI-Powered Situational Awareness

The Distinction That Matters

Situational awareness is not an information display. A radar screen displays information. A radio delivers audio. A flight schedule is a document. None of those individually, or even together on separate screens, constitute situational awareness. They constitute data. Situational awareness is what happens when a human mind successfully integrates all of that into a coherent picture of what is actually happening right now — and what is about to happen next.

That integration is precisely what breaks down under fatigue, cognitive overload, and time pressure. It is the highest cognitive function and the first to degrade. It is Level 3 Situational Awareness: not merely perceiving what is happening, not merely understanding what it means, but projecting what is about to happen. That is the level that failed at LaGuardia. And that is the level the Cognitive Fusion Engine is designed to restore.

> *The Cognitive Fusion Engine does not replace human integration. It performs the integration so the human receives comprehension, not raw data.*

Architecture: The Four Layers

The Conversation Layer

The AI monitors every active voice channel simultaneously — approach control, tower, ground control, gate operations, emergency frequencies — in real time. It is not transcribed. It is parsing intent and cross-referencing each communication against the live operational picture in VStrike.

The moment approach control clears a flight for runway 4 and ground control simultaneously clears a vehicle to cross, the AI detects the conversational conflict — two clearances that are physically incompatible — before either the plane or the truck has moved into the danger zone. At LaGuardia, the collision was authorized verbally before it became physical. The Conversation Layer catches it at the verbal layer.

- All voice channels ingested simultaneously as live semantic data streams.
- Each clearance, instruction, and acknowledgment parsed for operational intent.
- Cross-referenced in real time against active runway states, ground positions, and flight telemetry.
- Conflicting commitments flagged the moment they are spoken — not the moment they converge physically.

Layer 2: The Commitment Fusion Layer

VStrike already ingests sensor data: radar returns, thermal imaging, transponder signals, flight telemetry. The AI adds the semantic dimension — what the humans have said, what they have committed to, and what instructions are in flight that have not yet been physically executed.

The result is a real-time commitment map: not just where things are, but where things are going because a human said so. This is the precise gap that existed at LaGuardia — the truck had been cleared verbally, but no system held that commitment as an operational variable against the aircraft's approach vector. The Commitment Fusion Layer closes that gap structurally.

- Live flight schedule fused with real-time ATC clearances as a unified operational data model.
- Ground vehicle positions tracked against their authorized routes and committed crossing windows.
- Every active instruction is treated as a predicted future state, not just a logged past event.
- VStrike 3D canvas updated in real time to reflect committed future positions, not only current ones.

Layer 3: The Cognitive Offload Layer

The fatigued controller's fundamental problem is that he cannot hold the entire operational picture in working memory simultaneously under the load he was carrying that night. The Cognitive Offload Layer addresses this directly: the AI holds the complete operational picture so the human does not have to.

The controller still makes every decision. He still holds all authority. But he makes those decisions with a complete, synthesized, validated picture that the AI has assembled and surfaced — rather than a fragmented picture he is constructing himself from disparate inputs while managing four concurrent responsibilities at midnight.

- Operator interaction patterns monitored continuously for cognitive load indicators.
- Task saturation alerts generated when a controller's active responsibilities exceed defined safety thresholds.

- Automated supervisor notification triggered when human readiness falls below operational minimums.
- Non-critical information filtered from display during high-load moments — only what matters now.

Layer 4: The Anticipatory Alert Layer

This is the most critical departure from every system deployed at LaGuardia that night. Every existing alert fires when a conflict is detected. The Anticipatory Alert Layer fires when the trajectory of commitments is converging toward conflict — seconds earlier, before the terminal phase, when a go-around is still possible, when a stop is still possible, when the outcome is still a choice rather than a consequence.

The AI does not wait to see the truck and the plane converge. It calculates the intersection of their committed paths the moment both clearances exist in the system and projects the conflict forward in time. The alert is not "conflict detected." It is "conflict will occur in 47 seconds on current vectors. Recommend go-around for Flight 8646 and hold for Truck 1."

- Predictive conflict modeling runs continuously against all active commitments and live positions.
- Anticipatory alerts generated with projected time-to-conflict and recommended resolution actions.
- VStrike renders the predicted conflict path visually on the 3D canvas before it becomes physical.
- Resolution options displayed with downstream consequence modeling — the controller sees the ripple effect of each choice before executing it.

CFE Layer	What It Does	How It Changes the Outcome
Conversation	Monitors all voice channels simultaneously, parsing intent and cross-referencing against live operational state.	Conflicting clearances detected at the moment they are spoken — before either vehicle moves.
Commitment Fusion	Fuses spoken instructions with sensor data to build a real-time map of committed future positions.	The truck's authorized crossing becomes a predicted future state that the AI holds against the aircraft's approach vector.
Cognitive Offload	Monitors operator cognitive load and filters display to what matters now. Holds the complete picture the fatigued controller cannot.	The controller receives comprehension, not data. Override authority is preserved because cognitive capacity is preserved.
Anticipatory Alert	Project trajectory of all active commitments forward in time. Alerts on predicted conflicts before they enter the physical domain.	The alert fires with 47 seconds to impact, not 3. The go-around is commanded. The collision does not occur.

The Hierarchy That Must Hold

The Cognitive Fusion Engine operates under one inviolable constraint: it never removes the human from the decision chain. This is not a limitation of the technology. It is the governing doctrine.

The AI monitors. The AI fuses. The AI projects. The AI alerts. The human decides. The human commands. The human holds authority over the system, and the system exists to ensure that authority is exercised with the best possible information at the moment it is needed most.

> *VStrike provides the sovereign visual canvas. The Cognitive Fusion Engine provides the semantic and predictive fusion layer. The human provides the decision and the authority. None of the three can do the other's job — and none should try.*

VII. Conclusion: The Hierarchy Must Hold

Antoine Forest was 30 years old. MacKenzie Gunther was 38. They had completed their landing checklist. The plane had been cleared. From every perspective available to every person in that system that night, the approach was proceeding normally.

The runway was not clear.

The Sovereign Human Layer framework exists to address exactly this gap — the space between what the system reports and what is physically true. The Cognitive Fusion Engine closes the awareness gap. The Predictive Runway Sovereignty Layer closes the prevention gap. The doctrine of Cognitive Readiness as a Security Control closes the human capacity gap. And Out-of-Band Human Sovereignty ensures that when all of those layers are working, the human at the center of the system is equipped, informed, and empowered to exercise the authority that only they can hold.

We build these frameworks after accidents because that is when institutions are willing to listen. The names of two pilots are on this page because they deserve to be — and because the most respectful thing we can do for them now is make certain that the system that failed them cannot fail the next crew in the same way.

> **The purpose of this case study is not to assign blame. It is to make the case that we need not wait for the next accident to build what we already know how to build.**

About the Author

Jose A. Cedeño is Principal Advisor at Cofresí Consulting Services, specializing in Zero Trust architecture, Critical Infrastructure protection, and Operational Technology security. The Sovereign Human Layer framework — including Out-of-Band Human Sovereignty, Cognitive Readiness as a Security Control, Predictive Conflict Prevention, and the Cognitive Fusion Engine — was developed through applied advisory work across government and enterprise environments.

Companion Document: Governing the Human Layer — A Dynamic Zero Trust Framework with Out-of-Band Human Sovereignty (March 26, 2026)

References

[1] National Transportation Safety Board (NTSB). Press Briefings: Air Canada Express Flight 8646 / Port Authority Ground Vehicle Collision, LaGuardia Airport, March 22–24, 2026. Statement by NTSB Chair Jennifer Homendy. Washington, D.C.: NTSB, 2026. Available at: https://www.ntsb.gov/investigations/Pages/DCA26MA161.aspx

[2] Aviation Law Group, P.S. "Air Canada Express Flight 8646 Collision at LaGuardia: Early Facts and All Too Familiar Recurring Runway Safety Issues." March 2026. Available at: https://www.aviationlawgroup.com/air-canada-express-flight-8646-collision-at-laguardia-early-facts-and-all-too-familiar-recurring-runway-safety-issues/

[3] Brazy, Doug (NTSB Investigator-in-Charge) and Homendy, Jennifer (NTSB Chair). Cockpit Voice Recorder and Flight Data Recorder timeline briefing. NTSB Second Public Press Conference, LaGuardia Airport, March 24, 2026. Reported by Aerospace Global News. Available at: https://aerospaceglobalnews.com/news/ntsb-timeline-air-canada-crash-laguardia/

[4] CNN. "Fire truck in LaGuardia runway collision had no transponder, limiting tower's ability to track it, NTSB says." March 24, 2026. Available at: https://www.cnn.com/2026/03/24/us/laguardia-plane-collision-air-canada-wwk

[5] Wikipedia. "Air Canada Express Flight 8646." Wikimedia Foundation, March 2026. Available at: https://en.wikipedia.org/wiki/Air_Canada_Express_Flight_8646

[6] Endsley, Mica R. "Toward a Theory of Situation Awareness in Dynamic Systems." Human Factors: The Journal of the Human Factors and Ergonomics Society, 37(1), 32–64, 1995. The foundational source for the three-level situational awareness model (Perception, Comprehension, Projection) applied in Section VI of this case study.

[7] NPR. "NTSB shares early findings in deadly LaGuardia Airport plane crash." March 25, 2026. Available at: https://www.npr.org/2026/03/25/nx-s1-5759965/ntsb-shares-early-findings-in-deadly-laguardia-airport-plane-crash

[8] Cedeño, Jose A., Cofresí Consulting Services. "Governing the Human Layer: A Dynamic Zero Trust Framework with Out-of-Band Human Sovereignty." Cofresí Consulting Services White Paper, March 26, 2026. Companion document establishing the doctrinal framework applied in this case study.

[9] Homendy, Jennifer (NTSB Chair). Statement on systemic aviation safety factors including controller staffing and combined-position midnight shift operations. NTSB Press Conference, LaGuardia Airport, March 23–24, 2026. Cited in CNN: "What's next in the investigation into the deadly Air Canada collision at LaGuardia." March 25, 2026. Available at: https://www.cnn.com/2026/03/25/us/laguardia-collision-investigation

Cofresí Consulting Services
Principal Advisory Practice | Elkridge, MD
March 2026