The Autonomous Flight Deck: Safety Implications of Single-Pilot and Reduced-Crew Commercial Operations

Albert N. Clark

Independent Author

Published: August 30, 2026

ASX Research Journal and Database

ISSN 3068-3351 (Online)

Place of Publication: Cadiz City, Philippines

Publisher: ASXResearch.org

Author Note

Albert N. Clark

Department of Aerospace Sciences, ASXResearch.org

ORCID iD: https://orcid.org/0009-0002-7348-4395

The author reports no conflicts of interest.

Correspondence concerning this article should be addressed to Albert N. Clark,

Email: aclark@asxresearch.org

Abstract

The commercial flight deck has evolved toward greater automation and smaller crews, yet removing the second pilot presents a fundamentally different safety challenge. This article examines single-pilot and reduced-crew operations, emphasizing abnormal situations, pilot incapacitation, workload spikes, automation failure, cybersecurity, human–machine teaming, and the loss of independent human cross-checking. Although manufacturers, regulators, and military programs have demonstrated substantial advances in autonomous flight, emergency diversion, automated landing, and intelligent decision support, these technologies have not yet demonstrated the resilience, contextual reasoning, and adaptive problem solving provided by a second qualified pilot. Legal liability, certification, operational authority, and public acceptance further complicate implementation. The article concludes that autonomous systems should first be extensively validated as safety partners within two-pilot operations. Removing the second pilot should occur only when operational evidence demonstrates equivalent or superior safety during compound failures, incapacitation, cyber events, and unforeseen emergencies.

Keywords: autonomous flight, single-pilot operations, aviation safety

The Autonomous Flight Deck: Safety Implications of Single-Pilot and Reduced-Crew Commercial Operations

Figure 1

The Autonomous Flight Deck: Human–AI Teaming and Safety Challenges in Reduced-Crew Commercial Operations.

Note. Click image for full-size image or click here.

Commercial aviation has been quietly moving toward the single-pilot question for decades, long before artificial intelligence made the idea sound technologically fashionable. The progression from five-person flight decks to four, three, and eventually two crewmembers followed improvements in engines, avionics, navigation, flight management systems, and automatic flight control. Flight engineers and navigators disappeared because machines could perform sufficiently bounded functions with demonstrable reliability. Removing the second pilot is fundamentally different. That crewmember is not merely another operator of controls; the pilot monitoring provides an independent cognitive channel capable of questioning assumptions, detecting errors, interpreting ambiguous situations, communicating while the other pilot flies, and assuming command if the pilot flying becomes impaired. Myers and Starr (2021) observed that economic pressure, pilot availability, automation, and artificial intelligence have made single-pilot operations increasingly attractive, but their analysis also illustrates why eliminating the second pilot cannot simply be treated as the next historical step in cockpit automation. The problem is not whether automation can fly an airplane from departure to destination under normal circumstances. Modern aircraft have been capable of doing much of that for decades. The real question is whether an autonomous system can reproduce the safety resilience created when two qualified humans encounter something neither expected.

That distinction explains why contemporary research increasingly separates Extended Minimum-Crew Operations (eMCO) from true Single-Pilot Operations (SiPO). Under eMCO, two pilots remain aboard, but only one occupies the active flight deck during portions of cruise while the other rests; SiPO ultimately envisions one pilot conducting the entire operation with technological and potentially ground-based assistance. Schmid and Stanton (2020), after systematically reviewing 75 publications, concluded that workload allocation, pilot incapacitation, communications, health monitoring, data links, and certification remained unresolved areas requiring integration rather than isolated technological fixes. Vu et al. (2018) similarly found that single pilots could resolve some off-nominal scenarios when provided either advanced cockpit automation or ground support, but emphasized that the technologies remained prototypes and that no single operational concept had established itself as superior. That historical research has now encountered regulatory reality. EASA's completed eMCO-SiPO safety research concluded that, using the current cockpit as the baseline, equivalent safety between eMCO and conventional two-pilot operations could not be sufficiently demonstrated. Particularly troublesome areas included incapacitation, fatigue and drowsiness, sleep inertia, physiological needs, and—critically—the loss of human cross-checking. EASA therefore shifted the near-term emphasis toward developing and proving “Smart Cockpit” technologies first within conventional two-pilot operations rather than simply authorizing reduced crews and hoping automation closes the gap.

Abnormal operations expose the heart of the problem. Airline cockpits are designed around task sharing because emergencies rarely arrive as clean, isolated failures. An engine failure may coincide with weather, terrain, air traffic control instructions, checklist execution, passenger or cabin problems, diversion planning, fuel calculations, and degraded automation. One pilot can fly while the other diagnoses, communicates, verifies switches, retrieves procedures, challenges decisions, and maintains the broader operational picture. With only one pilot, these tasks converge on the same cognitive system precisely when its capacity is most stressed. Li et al. (2024) examined neural activity and visual behavior during abnormal and emergency single-pilot scenarios and found measurable changes in pilots' cognitive and visual responses, reinforcing the importance of defining a safe human-performance envelope rather than assuming that automation automatically compensates for the missing crewmember. Miranda (2025) went further: when unexpected automation failure was introduced into simulated single-pilot operations, pilots exhibited evidence of cognitive overload including attentional tunneling, degraded situation awareness, and inattentional deafness. This produces an uncomfortable paradox. The automation intended to make single-pilot flight possible can itself become the event that overwhelms the single remaining human. In a two-pilot cockpit, one crewmember can become cognitively saturated while the other retains enough capacity to recognize the deterioration. In SiPO, there may be nobody left to notice that the human component of the system is failing.

Physical incapacitation is even less forgiving because it converts a reduced-crew aircraft into a zero-pilot aircraft instantaneously. A heart attack, seizure, stroke, severe gastrointestinal illness, hypoxia, medication reaction, or loss of consciousness does not negotiate with the flight plan. Simons et al. (2025) concluded that eMCO introduces substantial aeromedical uncertainties because reliable systems for detecting physical and cognitive incapacitation are not yet sufficiently mature, while fatigue, boredom, sleep inertia, and ordinary physiological requirements create additional complications. A viable autonomous flight deck therefore requires far more than an autopilot. It needs continuous and extremely reliable assessment of pilot state, discrimination between sleep, distraction, cognitive impairment, and genuine incapacitation, an escalation protocol, secure communication with ground personnel, automatic stabilization of the aircraft, diversion selection, weather and runway assessment, communications with ATC, approach configuration, landing, runway evacuation or stopping, and ultimately coordination with emergency responders. False negatives could leave an incapacitated pilot nominally in command; false positives could cause automation to seize authority from a perfectly capable pilot. Puca and Guglieri (2025) consequently describe civil SPO as a system-of-systems challenge involving automation, communications, ground support, human factors, certification, and operational architecture rather than merely an avionics upgrade.

The second pilot also provides something considerably harder to engineer than another pair of hands: independent skepticism. Crew Resource Management institutionalized cross-checking because aviation learned, often through accidents, that humans make errors of perception, interpretation, memory, and judgment. Two pilots can misunderstand the same situation, but they do not necessarily misunderstand it identically. A first officer can challenge a captain's unstable approach, notice an incorrect altitude, question a checklist response, identify an automation-mode error, or recognize that the other pilot is becoming disoriented. Pechlivanis and Harris (2025), in developing and assessing a single-pilot concept of operations, identify hazards and mitigation requirements across the system rather than treating automation as a direct substitute for a crewmember. This distinction is essential. An AI trained on the same sensor inputs and operating assumptions as the aircraft's automation may not constitute truly independent redundancy. If the machine shares a corrupted data source, faulty model, erroneous database, compromised software component, or incorrect contextual assumption, the "second opinion" may simply reproduce the first error faster. Human-machine teaming therefore needs deliberate diversity: the autonomous partner must be capable not merely of confirming what the pilot is doing, but of detecting disagreement, explaining why it disagrees, escalating uncertainty, and—under carefully bounded circumstances—challenging the human.

That requirement transforms the concept from automation into genuine human-autonomy teaming. Tokadlı and Dorneich (2023) argue that humans and autonomous systems possess asymmetric capabilities and that present automated systems still lack many characteristics necessary to function as true teammates. Humans remain unusually capable at contextual reasoning, improvisation, moral judgment, interpreting weak signals, and recognizing that "something isn't right" before the problem fits a predefined category. Machines excel at persistent monitoring, computational precision, rapid database retrieval, and simultaneously observing quantities of information that exceed human attention. The ideal autonomous copilot therefore should not be designed as an electronic imitation of a first officer. It should be designed as a complementary cognitive agent: one that remembers everything, monitors continuously, never becomes fatigued, but understands its own uncertainty and knows when the human's contextual judgment should dominate. This also means the human must understand what the machine knows, why it recommends an action, what information it is using, and how confident it is. Otherwise automation bias merely changes form: instead of two humans cross-checking each other, one overloaded human may accept an authoritative-looking machine recommendation because there is no second person available to question it.

Cybersecurity makes this architecture considerably more dangerous than the traditional autopilot. A conventional flight-control computer can be protected through partitioning, deterministic software, redundancy, and tightly controlled interfaces. A future autonomous copilot may require access to communications, aircraft health information, navigation databases, weather, airline operational data, biometric pilot monitoring, and perhaps ground-based human assistance. Every additional information path expands both capability and attack surface. Schmid and Stanton (2020) specifically identified data-link and certification issues among the unresolved problems in reduced-crew operations, while EASA now explicitly includes prevention of security threats among the capabilities required before reduced-crew concepts could be reconsidered. The cybersecurity problem becomes especially severe if a ground operator is expected to replace some functions of the absent pilot. Loss, latency, spoofing, corruption, denial of service, or malicious takeover of that link cannot be allowed to leave the onboard pilot suddenly without the support upon which the safety case depends. Nor can an autonomous system be permitted to interpret untrusted external information as authoritative flight guidance. A credible SiPO architecture therefore requires graceful degradation: the aircraft must remain safe when connectivity disappears, when the AI is unavailable, when sensor information conflicts, and even when the system suspects that it itself may have been compromised. In safety engineering terms, autonomy cannot merely be reliable when healthy; it must fail safely when unhealthy.

Industry is already building pieces of that future, with Airbus providing perhaps the clearest large-commercial-aircraft example. Its Autonomous Taxi, Take-Off and Landing (ATTOL) program demonstrated vision-based autonomous taxiing, takeoff, and landing after more than 500 test flights, and Airbus UpNext's DragonFly demonstrator subsequently tested automated emergency diversion, automatic landing, and taxi assistance on an A350-1000. These are striking capabilities because emergency diversion following pilot incapacitation is precisely one of the functions a reduced-crew safety case would require. Yet Airbus's work should not be confused with certification of an autonomous passenger airliner. Gao et al. (2025) found that public acceptance of an AI pilot remains strongly influenced by negative emotion, trust, and perceived risk, with participants preferring dual-pilot operations even when experimental scenarios stipulated equivalent safety. Kioulepoglou and Makris (2023) similarly found generally negative attitudes toward single-pilot airline operations, although participants preferred a combination of sophisticated onboard automation and a permanent ground operator when asked how the second pilot might be replaced. Technological capability, therefore, is only one certification hurdle. An airline may eventually demonstrate mathematically that SiPO meets an equivalent level of safety and still discover that passengers refuse to buy the ticket.

The military is considerably further ahead because its risk equation is different, and DARPA has directly attacked the problem. The Aircrew Labor In-Cockpit Automation System (ALIAS) program was created to develop a portable, adaptable automated assistant capable of reducing onboard crew while supporting entire missions from takeoff through landing and responding to contingency events. DARPA partnered with Sikorsky, a Lockheed Martin company, and ALIAS-derived MATRIX autonomy ultimately flew a UH-60A Black Hawk with nobody aboard in 2022. More importantly, the technology did not remain a laboratory curiosity: in March 2026, DARPA transferred an experimental H-60Mx Black Hawk equipped with the MATRIX autonomy suite to the U.S. Army for operational testing. The military value is obvious. Autonomy can reduce exposure of pilots to hostile fire, operate in degraded visual environments, permit optional crewing, and allow human operators to concentrate on mission decisions rather than basic aircraft manipulation. But military success does not automatically establish civil acceptability. Combat aviation explicitly accepts operational risks that passenger transportation does not, and a Black Hawk autonomously executing a military mission is not carrying hundreds of fare-paying civilians whose legal expectation is transportation under an extraordinarily mature safety regime. DARPA nevertheless matters enormously because ALIAS demonstrates that high-level autonomy can be retrofitted into complex aircraft and can manage far more than straight-and-level flight.

Elon Musk, despite his prominence in artificial intelligence, autonomous vehicles, reusable spacecraft, and transportation technology, does not appear to be a substantive participant in the present commercial reduced-crew or single-pilot airline research ecosystem. There is no credible evidence that Musk, Tesla, SpaceX, xAI, or another Musk-controlled enterprise is presently developing a certified autonomous flight-deck replacement for the second pilot of Part 121 passenger aircraft. SpaceX unquestionably demonstrates extremely sophisticated autonomous guidance, navigation, rendezvous, docking, launch-abort, and landing capabilities, but spacecraft automation operates under a radically different certification, operational, and liability environment. The more relevant civil actors are Airbus, aviation avionics and autonomy developers, NASA researchers, EASA, aircraft manufacturers, universities, and human-factors laboratories; militarily, DARPA, Sikorsky/Lockheed Martin, and the U.S. Army provide concrete examples. This distinction is important because technological celebrity can distort the discussion. Commercial aviation does not certify charisma, disruption, or a spectacular demonstration. It certifies systems against defined failure probabilities, operational hazards, human-performance limitations, maintenance requirements, software behavior, and foreseeable misuse. An autonomous airliner must therefore prove itself not once on a dramatic demonstration flight, but millions of times statistically—including on the day when several improbable things go wrong together.

The legal implications may ultimately prove almost as difficult as the engineering. In the United States, 14 C.F.R. § 121.385 currently establishes two pilots as the minimum pilot crew for Part 121 operations and specifically rejects satisfying multiple simultaneously required certificated-airman functions through one person. A transition to single-pilot airline operations would therefore require more than an aircraft manufacturer demonstrating technical feasibility; it would require regulatory change, new certification assumptions, revised operational rules, training standards, dispatch and ground-support requirements, cybersecurity obligations, and an explicit allocation of authority between pilot, automation, airline, manufacturer, and possibly a remote operator. Civil liability becomes especially fascinating after an accident. If an AI rejects the pilot's correct command, is responsibility assigned to the pilot, airline, aircraft manufacturer, software developer, training organization, or certification authority? If the pilot follows erroneous AI advice, was that reasonable reliance on certified automation or negligent surrender of command? If a ground operator simultaneously supervises several aircraft and two experience emergencies, which aircraft receives the human? These are not philosophical decorations around the engineering problem. They determine how authority, duty of care, product liability, negligence, evidence preservation, insurance, and accident investigation must be structured before the system enters revenue service. A future cockpit data recorder may consequently need to preserve not only control inputs and spoken words but the autonomous system's recommendations, confidence states, data sources, internal mode transitions, disagreements with the pilot, and reasons for assuming or relinquishing authority.

The most defensible path forward is therefore not “replace the first officer with AI,” but progressively create an autonomous safety partner and force it to prove its value while two humans are still present. Lopes et al. (2026) frame future single-pilot operations as a socio-technical problem spanning onboard technology, ground operations, human factors, certification, economics, and public acceptance; that integrated perspective is considerably more convincing than treating crew reduction as a simple consequence of better autopilots. The next generation of flight decks could introduce intelligent monitoring, abnormal-situation diagnosis, adaptive checklist management, incapacitation detection, automated diversion and landing, cyber-resilient ground assistance, and explainable decision support while retaining both pilots. Billions of operational hours could then reveal whether these systems actually reduce errors, identify incapacitation reliably, manage workload spikes, resist cyberattack, and challenge pilots appropriately. Only after that evidence exists should aviation ask whether one human can safely be removed. The decisive certification question should not be whether one pilot plus AI can perform as well as two pilots on an ordinary Tuesday afternoon. It should be whether that combination is at least as resilient when the airplane is damaged, the weather is deteriorating, the automation is confused, communications are failing, the pilot is overloaded or incapacitated, and the event unfolding at 35,000 feet is one that nobody thought to program into the simulator. Until that question can be answered empirically rather than optimistically, the empty right seat remains not wasted capacity, but one of commercial aviation's oldest and most remarkably effective layers of redundancy.

References

Gao, S., Lu, Z., Luan, H., Yin, M., & Wang, L. (2025). *AI pilot in the cockpit: An investigation of public acceptance.* International Journal of Human–Computer Interaction, 41(1), 543–556. https://doi.org/10.1080/10447318.2024.2301856

Kioulepoglou, P., & Makris, I. (2023). *Single pilot operations and public acceptance: A mixed methods study conducted in Greece*. International Journal of Aviation, Aeronautics, and Aerospace, 10(3). https://doi.org/10.58940/2374-6793.1822

Li, Q., Chen, C.-H., Ng, K. K. H., Yuan, X., & Yiu, C. Y. (2024). *Single-pilot operations in commercial flight: Effects on neural activity and visual behaviour under abnormalities and emergencies.* Chinese Journal of Aeronautics, 37(8), 277–292. https://doi.org/10.1016/j.cja.2024.04.007

Lopes, N. M., Aparicio, M., Costa, C., Neves, F. T., & Bernardino, C. (2026). *Designing the future of flight: A socio-technical framework for single-pilot operations in commercial aviation.* Journal of Air Transport Management, 132, 102939. https://doi.org/10.1016/j.jairtraman.2025.102939

Miranda, N. (2025). *Single pilot operations: Performance in normal and abnormal scenarios*. Transportation Research Procedia, 88, 185–192. https://doi.org/10.1016/j.trpro.2025.05.023

Myers, P. L., III, & Starr, A. W., Jr. (2021). *Single pilot operations in commercial cockpits: Background, challenges, and options*. Journal of Intelligent & Robotic Systems, 102, Article 19. https://doi.org/10.1007/s10846-021-01371-9

Pechlivanis, K., & Harris, D. (2025). *Single pilot concept of operations: Hazard identification and mitigation measures*. The International Journal of Aerospace Psychology, 35(4), 161–184. https://doi.org/10.1080/24721840.2025.2481880

Puca, N., & Guglieri, G. (2025). *Enabling civil single-pilot operations: A state-of-the-art review.* Aerotecnica Missili & Spazio, 104, 187–212. https://doi.org/10.1007/s42496-024-00223-7

Schmid, D., & Stanton, N. A. (2020). *Progressing toward airliners' reduced-crew operations: A systematic literature review*. The International Journal of Aerospace Psychology, 30(1–2), 1–24. https://doi.org/10.1080/24721840.2019.1696196

Simons, R., Maher, D., Vermeiren, R., & Wagstaff, A. S. (2025). *Aeromedical concerns about extended minimum crew operations.* Aerospace Medicine and Human Performance, 96(7), 590–592. https://doi.org/10.3357/AMHP.6671.2025

Tokadlı, G., & Dorneich, M. C. (2023). *Comparison and synthesis of two aerospace case studies to develop human-autonomy teaming requirements*. Frontiers in Aerospace Engineering, 2, 1214115. https://doi.org/10.3389/fpace.2023.1214115

Vu, K.-P. L., Lachter, J., Battiste, V., & Strybel, T. Z. (2018). *Single pilot operations in domestic commercial aviation.* Human Factors, 60(6), 755–762. https://doi.org/10.1177/0018720818791372