

## Bombing Clouds:

# When Military AI Depends on Commercial Data Centres

By Anne-Marie Buzatu

March 2026

**Abstract:** March 2026 fighting in the Gulf brought into view two linked but distinct episodes: Iranian strikes on Amazon’s AWS facilities in the region, amid claims that commercial cloud infrastructure may be supporting AI-enabled military targeting; and U.S. operations in Iran in which AI-assisted systems reportedly helped identify, locate, and prioritise roughly 1,000 targets in 24 hours. This article argues that together they expose a double challenge for international humanitarian law. First, commercial data centres supporting AI-enabled military functions may, in some circumstances, qualify as military objectives (though the public record does not establish the lawfulness of the AWS strikes and any such attack would remain subject to proportionality and precautions in attack). Second, AI-assisted targeting can compress and obscure decision-making, leaving humans formally involved while eroding the judgment the law requires. As military and civilian functions become more deeply entwined in shared cloud systems, both distinction and sound attack decision-making become harder to preserve. This article does not address the separate legal questions raised by striking facilities in other countries or by treating data itself as a target.

**Introduction**

On 1 March 2026, the cloud became a target. In the preceding 24 hours, public reporting indicated that U.S. forces had used AI-enabled systems to help identify, locate, and prioritise roughly 1,000 targets in Iran. The following day, Iranian drones struck Amazon Web Services facilities in the Gulf, setting off fires, knocking out power, and causing outages far beyond the data centres themselves. In the UAE and Bahrain, the effects were felt not only in enterprise and government systems, but across banking platforms, payment services, food-delivery apps, transport systems, logistics networks, and other civilian digital services, with material effects on civilian life. (1)(2) Taken together, these two episodes briefly made visible what is usually hidden: ordinary civilian life and modern military power can now depend on the same commercial cloud infrastructure. (1)(2)

The episodes also brought two questions into sharp relief. First, can commercial data centres that support AI-enabled military targeting qualify as military objectives? Second, does the speed of those systems leave enough room for meaningful human judgment before attacks are carried out?

The core problem is twofold: embedding military data, models, and algorithmic support in shared commercial cloud makes civilian infrastructure more plausibly targetable and civilian harm harder to contain; embedding AI-assisted recommendations in compressed workflows makes lawful human judgment thinner, faster, and harder to defend.

For diplomats and policy-makers, these strikes raise immediate questions about how to protect the essential civilian digital services on which modern life now depends, how to manage escalation risks around cloud infrastructure, and what precautions states and providers must now build into AI-enabled operations.

This article argues that the AWS strikes exposed a dangerous double shift in AI-enabled warfare: shared commercial cloud infrastructure may become more plausibly targetable as it supports military AI, even as the speed and opacity of those same systems may erode the soundness of the human judgment on which lawful attack decisions depend. The first problem is one of classification. The second concerns decision-making. Both grow out of the same trend: militaries are increasingly relying on AI-enabled systems to support operations and attack decisions, while running those systems on infrastructure that also sustains important civilian digital life. (2)(4)(5)(6)

## I. The AWS strikes made a legal question impossible to ignore

A commercial data centre is not an obviously military object. But the legal question raised by the AWS strikes is a serious one. Under the law of armed conflict, data centres are classified as objects. (2)(3) Objects may qualify as military objectives if, by their use or purpose, they make an effective contribution to military action and if their destruction, capture, or neutralisation offers a definite military advantage in the circumstances ruling at the time. (3) If a data centre is materially sustaining military intelligence processing, operational planning, or AI-enabled targeting support, the case for treating it as a military objective becomes a serious one, subject always to proportionality and precautions in attack. (2)(3)

That said, this does **not** mean these particular strikes were lawful. The public record does not establish that the specific AWS facilities struck in the UAE and Bahrain were in fact supporting military workloads. And that qualification matters. The law does not allow an attacker to hit first and determine later whether a facility was serving military functions; Article 57(2)(a)(i) AP I requires attackers to do everything feasible in the circumstances to verify that a target is a military objective before attack. (2)(3) Nor does it permit the claimed military advantage to remain vague or speculative. (3) If sufficient information was not available, in the circumstances ruling at the time, to conclude that a specific facility was making an effective contribution to military action or was intended for such use, it has to be treated as civilian. (2)(3)

Even if a data centre qualifies as a military objective because part of its infrastructure supports military action, that does not end the legal analysis. Mixed-use facilities remain subject to the rules of proportionality and precautions in attack. The more civilians depend on the same infrastructure, the more seriously an attacker must assess foreseeable incidental civilian harm and the availability of feasible alternatives or more discriminating means.

This is also unlikely to be an isolated problem. Recent public partnerships involving Anthropic, Palantir, and AWS, together with Anthropic's own announcements on defence operations and expanded government access, point in a clear direction of travel: commercial AI, commercial cloud, and military or intelligence use are becoming more tightly integrated, not less. In addition, other commercial AI providers including Google, Meta and OpenAI have quietly and recently changed their policies to provide military support services. (See Box) The question raised by the AWS strikes is therefore not confined to one incident in one conflict. (19) It is part of a broader structural shift in the organization of military capability. (4)(5)(6)

**Box: Major AI Firms Quietly Re-Open the Door to Military Uses**

**Google rolled back its ban on weapons and surveillance AI (2025).**

In February 2025, Google quietly removed key “Responsible AI” commitments that barred use of its AI in weapons and certain surveillance, prompting civil-society groups to warn it had “dropped its pledge” and “rolled back” its responsible-AI principles to allow military applications again. (12)

**OpenAI removed its explicit prohibition on “military and warfare” (2024).**

In January 2024, OpenAI revised its use-case policy to delete a categorical ban on “military and warfare,” while retaining narrower prohibitions on uses like developing weapons or conducting communications surveillance; commentators note this “opened the door” to military customers and positive-spin “national security” use cases. (13)

**Meta enables access to its AI by defense and security agencies.**

Meta maintains public policy language against “military, warfare, nuclear industries or applications,” but from late 2024 has allowed government and defense-sector agencies to access its LLaMA-based models under special arrangements, effectively authorizing military-linked uses despite the nominal restrictions. (14)

**These shifts are part of a wider militarization trend.** A 2025 investigation in The New York Times groups Google, Meta and OpenAI together as leading examples of Silicon Valley firms that have weakened earlier red-line bans on military uses and are increasingly pursuing defense and war-related contracts. (15)

*The following developments are drawn from public reporting by outlets such as Wired, TechCrunch, AIHub and The New York Times.*

That dependence is not accidental. Contemporary militaries increasingly buy rather than build because leading commercial providers can often offer more advanced, scalable, or more rapidly deployable cloud, compute, models, and engineering support than governments can reproduce internally at comparable speed. U.S. defence policy documents increasingly reflect an adopt-buy-create logic for AI and cloud, under which commercial solutions can provide best-in-class capabilities for many dual-use applications. (16) But that operational advantage is also the legal problem: the more military capability is built on shared commercial systems because they are faster and better, the less plausible it becomes to treat the resulting commingling with civilian digital life as incidental or temporary. Outsourcing the analytic layer does not privatize the legal burden of compliance: a privately owned facility may qualify as a military objective because of the military use it supports, while responsibility for wrongful conduct turns on whether the relevant conduct remains attributable to the state.

## II. The opaqueness problem

The deepest difficulty is not the legal test itself. It is the architecture to which the test is now being applied. Cloud systems are distributed, virtualised, and built to route workloads across infrastructure rather than tie them neatly to one physical location. (2) In highly time-sensitive military contexts, latency matters, which creates incentives to move compute and algorithmic support close to the point of use. (2) Commercial cloud providers can reroute or open processing pathways in different locations according to demand and speed. But from the outside, those pathways are hard to see. Neither attackers nor outside observers can easily follow the path of individual data packets, much less know which data or algorithms are being processed at a particular site at a particular moment. (2) For obvious cybersecurity and commercial reasons, cloud systems are meant to be opaque to outsiders. Operationally, that opacity is a feature. Legally, it is a complication.

That matters because international humanitarian law (IHL), also referred to as the law of armed conflict, depends on factual judgments before force is used. If a commercial data centre supports both civilian services and militarily relevant functions, its legal status becomes difficult to assess in a grounded way. (2)(3) Yet simply ignoring the military support function because civilian services also run on the same infrastructure is not satisfactory either. Private commercial data centres can provide real military value to belligerents. (2)(4)(5)(6) If that reality is treated as legally irrelevant whenever civilian uses are present, an important contribution to military effectiveness goes unacknowledged.

The opposite move is no better. The distributed and opaque nature of cloud infrastructure makes it tempting to say that once a commercial data centre serves any militarily relevant function, it should simply be treated as military. (2)(3) But that response cuts against the law's requirement to presume civilian status where doubt remains. (2)(3) The problem, then, is not that the civilian presumption is wrong. It is that cloud architecture makes it much harder to know what the presumption is resting on. The law expects a distinction between civilian objects and military objectives. (3) Shared commercial cloud systems make that distinction harder to verify without collapsing it.

This is the real gap the AWS strikes exposed. Commercial cloud infrastructure may now be providing operationally significant military support, especially for AI-enabled functions, (2)(4)(5)(6)(7) while its distributed design makes that support difficult to see, verify, or localise. That creates a growing friction between IHL's categories and the architecture of contemporary digital infrastructure. And the harder it becomes to know what a data centre is really supporting, the greater the burden placed on the human beings asked to make lawful decisions about attacking it.

### III. Faster targeting, thinner judgment

The second problem exposed by the AWS strikes is not about the status of infrastructure, but about the quality of decision-making. *The Washington Post* reported that U.S. forces struck roughly 1,000 targets in 24 hours in Iran, supported by AI-enabled systems that helped identify, locate, and prioritise very large numbers of targets in extremely compressed timeframes.<sup>(7)</sup> If that reporting is even broadly correct, the relevant question is no longer simply whether a human being appears somewhere in the chain of decision. It is whether that human still has enough time, information, and cognitive space to exercise meaningful judgment before force is used.

The legal problem is not simply whether a human remains somewhere in the system. It is whether those responsible for attack decisions retain enough **time, information, and authority** to do what IHL requires of them: verify the target, assess anticipated military advantage and expected civilian harm, consider feasible precautions and alternatives, and cancel or suspend an attack where the legal conditions for lawful attack are not met.

This is where the language of “human-in-the-loop” becomes too weak. A human can remain formally present in a workflow while performing little more than nominal ratification of machine-generated outputs. Meaningful human judgment requires more than a final click of approval; it requires access to the basis of the recommendation, visibility into uncertainty, practical time to assess anticipated military advantage and expected civilian harm, authority to question or refuse the recommendation, and a traceable record of who reviewed what and on what basis.

U.S. policy and recent public statements have likewise emphasized that current AI systems should not be used to kill without human involvement. <sup>(17)</sup> But that baseline does not settle the harder question addressed here: whether the human role that remains is sufficiently informed, independent, and uncompressed to amount to meaningful human judgment rather than formal signoff. The legal question, in other words, is not whether a human is somewhere in the workflow, but whether the workflow preserves a form of human judgment capable of satisfying IHL’s requirements.

As the scale and tempo of targeting increase, even highly trained analysts and commanders may struggle to meaningfully evaluate hundreds of machine-assisted recommendations under severe time pressure unless the standard of review becomes thinner, more routinized, or more deferential to the system. The problem, in other words, is not only automation. It is compression. Even halving the reported scale to 500 recommendations in 24 hours leaves the

compression problem intact. At 15–30 minutes of individualized review per recommendation, the process would require roughly 125–250 reviewer-hours: about 16–32 personnel across three eight-hour shifts before adding commanders, legal advisers, intelligence support, technical staff, or any dual-review requirement. Public reporting does not describe a review architecture on anything like that scale. That does not prove such a structure was absent. But it does make it harder to infer meaningful individualized judgment from the mere fact that a human remained formally somewhere in the chain. (18)

AI systems can accelerate target generation, target correlation, and prioritisation well beyond unaided human speed, (7) but international humanitarian law does not relax because the pace of operations increases. The obligations of distinction, proportionality, and precautions in attack remain grounded in human judgment. (8)(9)

This is why the issue should not be framed as whether humans have been “removed” from the decision. They may not have been. The more serious concern is that AI-assisted targeting can leave humans in place while eroding the soundness of the judgment they are supposed to exercise. If that judgment is compressed into rapid endorsement of outputs whose basis, confidence level, and downstream effects are only partially understood, the lawfulness of attack decisions becomes harder to defend even when formal human involvement remains intact.

#### **IV. Build precautions upstream**

The AWS strikes suggest that the response cannot be limited to debating after the fact whether one facility was lawfully targetable. If shared commercial infrastructure is increasingly being drawn into military AI workflows, (4)(5)(6) precautions need to be built earlier—into infrastructure design, procurement, and targeting practice. Those precautions cannot be left to vendor promises or procurement clauses alone. If governments increasingly govern military AI through contracts, those arrangements should complement rather than substitute for public-law safeguards, auditable system design, and operational review standards.

**States and militaries should separate military from civilian cloud infrastructure as far as feasible—or, where full separation is not possible, strictly compartmentalize military data, models, inference pipelines, and associated orchestration layers from civilian cloud infrastructure, as far as feasible.** States should stop treating hyperscale commercial cloud as the default home for sensitive military AI and intelligence workloads when the same facilities sustain important civilian digital services. (16) South Korea’s Defense Integrated Data Center and associated defense cloud—an example also highlighted by Klonowska and Schmitt in their

discussion of passive precautions—offer a useful directional example of dedicated military infrastructure rather than routine commingling with civilian systems. (2)(10)(11)

**States and militaries should define meaningful human judgment in operational terms.** States should adopt concrete procedures for AI-assisted target recommendations that require documented review of the basis for the recommendation, the degree of uncertainty, the anticipated military advantage, expected civilian harm, and feasible alternatives. A human should have both the authority and the practical time to question, modify, or refuse the recommendation.

**States, militaries, companies and civil society should launch a Montreux Document-style process for AI-enabled military operations.** A focused multi-stakeholder initiative could help clarify how IHL applies when military AI runs on private cloud systems, including questions of mixed-use infrastructure, auditability, passive precautions, provider responsibilities, and operational standards for meaningful human judgment of AI-supported recommendations. A useful starting point is that meaningful human judgment requires at least five things: access to the basis of the recommendation, visibility into uncertainty, practical time to assess military advantage and expected civilian harm, authority to question or refuse the recommendation, and a traceable record.

None of these steps would remove legal uncertainty altogether. But they would make it harder to hide militarily relevant functions inside civilian infrastructure, and harder to reduce human review to a ritual.

## **Conclusion**

The AWS strikes brought into view two linked shifts in AI-enabled warfare. Commercial cloud infrastructure may increasingly support militarily relevant functions in ways that make it more plausibly targetable under the law of armed conflict.(2)(3) But the distributed and opaque architecture of shared cloud systems makes those functions difficult to see, verify, and localise in practice.(2) The result is a growing friction between IHL's legal categories and the digital infrastructure on which both military operations and civilian life now depend.

At the same time, the speed at which AI-enabled systems can generate, correlate, and prioritise targets raises a different concern. Even where humans remain formally present in the chain of decision, the pace of operations may erode the quality of the judgment they are supposed to exercise.(7) The central question is no longer only whether a person is “in the loop,” but whether that person still has the time, information, and authority needed to assess uncertainty,



evaluate the anticipated military advantage, weigh expected civilian harm, and stop or change an attack recommendation before force is used.(8)(9)

This is why the AWS strikes should not be treated as a novel but isolated incident. They are an early warning about what happens when military AI is built on commercial cloud systems that also sustain important civilian digital services.(1)(2) The more military functions are embedded in shared private infrastructure, the harder it becomes to preserve distinction at the level of objects and sound legal judgment at the level of decisions. Because that dependence is driven by structural capability advantages in the private sector rather than by temporary expedience alone, the problem is likely to deepen rather than recede. If military AI continues to run on shared commercial cloud, the law will not fail all at once. It will erode in practice—through blurred objects, compressed decisions, and civilian systems drawn ever deeper into the battlespace.

## Endnotes

1. Reuters, "Amazon cloud unit flags issues at Bahrain, UAE data centers amid Iran strikes," 2 March 2026. Available at: <https://www.reuters.com/world/middle-east/amazon-cloud-unit-flags-issues-bahrain-uae-data-centers-amid-iran-strikes-2026-03-02/> (last accessed 16 March 2026).
2. Klaudia Klonowska and Michael Schmitt, "Iranian Attacks on the Amazon Data Centers: A Legal Analysis," *Just Security*, 12 March 2026. Available at: <https://www.justsecurity.org/133685/iranian-attacks-amazon-data-centers-legal-analysis/> (last accessed 16 March 2026).
3. International Committee of the Red Cross, *Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I), Article 52*. Available at: <https://ihl-databases.icrc.org/en/ihl-treaties/api-1977/article-52>; International Committee of the Red Cross, *Commentary of 1987: Article 52 – General protection of civilian objects*. Available at: <https://ihl-databases.icrc.org/en/ihl-treaties/api-1977/article-52/commentary/1987> (last accessed 16 March 2026).
4. TechCrunch, "Anthropic teams up with Palantir and AWS to sell its AI to defense customers", 6 November 2024, available at: <https://techcrunch.com/2024/11/07/anthropic-teams-up-with-palantir-and-aws-to-sell-its-ai-to-defense-customers/> (last accessed 16 March 2026).
5. Anthropic, "Anthropic and the Department of Defense to Advance Responsible AI in Defense Operations," 14 July 2025. Available at: <https://www.anthropic.com/news/anthropic-and-the-department-of-defense-to-advance-responsible-ai-in-defense-operations> (last accessed 16 March 2026).
6. Anthropic, "Offering Expanded Claude Access Across All Three Branches of Government," 12 August 2025. Available at: <https://www.anthropic.com/news/offering-expanded-claude-access-across-all-three-branches-of-government> (last accessed 16 March 2026).
7. *The Washington Post*, "Anthropic's AI tool Claude central to U.S. campaign in Iran, amid a bitter feud," 4 March 2026. Available at: <https://www.washingtonpost.com/technology/2026/03/04/anthropic-ai-iran-campaign/> (last accessed 16 March 2026).
8. International Committee of the Red Cross, *Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I), Article 48*. Available at: <https://ihl-databases.icrc.org/en/ihl-treaties/api-1977/article-48> (last accessed 16 March 2026).
9. International Committee of the Red Cross, *Protocol Additional to the Geneva Conventions of 12 August 1949, and relating to the Protection of Victims of International Armed Conflicts (Protocol I), Article 57*. Available at: <https://ihl-databases.icrc.org/en/ihl-treaties/api-1977/article-57> (last accessed 16 March 2026).

10. “고가용성 보장형 국방 클라우드 시스템 도입 전략” [Strategy for Introducing a High-Availability Defense Cloud System]. Available at: <http://syscore.sejong.ac.kr/~woongbak/publications/DJ30.pdf>
11. NATO Cooperative Cyber Defence Centre of Excellence, *National Cybersecurity Organisation: Republic of Korea*(2022). Available at: <https://ccdcoe.org/uploads/2022/12/ROK-Country-report.pdf> (last accessed 16 March 2026).
12. Wired, “Google Lifts a Ban on Using Its AI for Weapons and Surveillance”, 4 February 2025, available at: <https://www.wired.com/story/google-responsible-ai-principles/> (last accessed 16 March 2026).
13. TechCrunch, “OpenAI changes policy to allow military applications”, 11 January 2024, available at: <https://techcrunch.com/2024/01/12/openai-changes-policy-to-allow-military-applications/> (last accessed 16 March 2026).
14. AIHub, “Meta now allows military agencies to access its AI software. It poses a moral dilemma for everybody”, 24 November 2024, available at: <https://aihub.org/2024/11/25/meta-now-allows-military-agencies-to-access-its-ai-software-it-poses-a-moral-dilemma-for-everybody/> (last accessed 16 March 2026).
15. The New York Times, “The Militarization of Silicon Valley”, 4 August 2025, available at: <https://www.nytimes.com/2025/08/04/technology/google-meta-openai-military-war.html> (last accessed 16 March 2026).
16. Department of Defense, *Data, Analytics, and Artificial Intelligence Adoption Strategy* (Washington, D.C., 2023), available at: [https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD\\_DATA\\_ANALYTICS\\_AI\\_ADOPTION\\_STRATEGY.PDF](https://media.defense.gov/2023/Nov/02/2003333300/-1/-1/1/DOD_DATA_ANALYTICS_AI_ADOPTION_STRATEGY.PDF). (accessed 17 March 2026).  
Department of Defense, *DoD Cloud Strategy* (Washington, D.C., 2018), available at: <https://media.defense.gov/2019/feb/04/2002085866/-1/-1/1/dod-cloud-strategy.pdf> (accessed 17 March 2026). Segregation, however, is not a complete answer: where dedicated military cloud remains privately owned or operated, it may reduce ambiguity at the level of infrastructure while raising separate questions about civilian contractor personnel engaged in operationally proximate support, including whether civilian personnel employed by commercial providers may, depending on the functions they perform, lose protection against direct attack for such time as they directly participate in hostilities. That doctrine applies to persons, not objects, and is therefore beyond the scope of this article.
17. United States Department of Defense, *DoD Directive 3000.09, Autonomy in Weapon Systems*, 25 January 2023, available at: <https://www.esd.whs.mil/portals/54/documents/dd/issuances/dodd/300009p.pdf> (accessed 17 March 2026) (requiring autonomous and semi-autonomous weapon systems to be designed to allow commanders and operators to exercise “appropriate levels of human judgment over the use of force”); Covington & Burling LLP, “White House Issues National Security Memorandum on Artificial Intelligence,” 24 October 2024, available at: <https://www.cov.com/en/news-and-insights/insights/2024/11/white-house-issues->

[national-security-memorandum-on-artificial-intelligence-ai](#) (accessed 17 March 2026); White House, “Initial Rescissions of Harmful Executive Orders and Actions,” 20 January 2025, available at: <https://www.whitehouse.gov/presidential-actions/2025/01/initial-rescissions-of-harmful-executive-orders-and-actions/>(accessed 17 March 2026).

18. For a reported example of how human review may collapse into near-instant ratification in another context, see Yuval Abraham, “‘Lavender’: The AI machine directing Israel’s bombing spree in Gaza,” *+972 Magazine / Local Call*, 3 April 2024, available at: <https://www.972mag.com/lavender-ai-israeli-army-gaza/> (accessed 17 March 2026); see also *Longreads*, “‘Lavender’: The AI Machine Directing Israel’s Bombing Spree in Gaza,” 3 April 2024, available at: <https://longreads.com/2024/04/03/lavender-the-ai-machine-directing-israels-bombing-sprees-in-gaza/> (accessed 17 March 2026). The investigation reports one officer spending roughly 20 seconds per target and acting largely as a “stamp of approval.”
19. Some subsequent reporting also indicated that at least two data centres in Tehran were struck in subsequent U.S.–Israeli operations, alongside broader attacks affecting Iranian digital and telecommunications infrastructure. If accurate, that reporting would further underscore the extent to which data centres are being drawn into contemporary conflict. See Anadolu Agency, “Iran war shows data centers emerging as critical targets,” 2026, available at: <https://www.aa.com.tr/en/middle-east/iran-war-shows-data-centers-emerging-as-critical-targets/3852984> (accessed 17 March 2026); *Inside Towers*, “U.S., Israel Hit Tehran Data Centers After Iran Targets Amazon,” 2026, available at: <https://insidetowers.com/u-s-israel-hit-tehran-data-centers-after-iran-targets-amazon/>(accessed 17 March 2026).