Infrastructure as a Determinant of Human Control: Modern Data and Network Architectures in Command and Control Systems

Muhamad Aria Armada Djojosugito[1]*
[1] *Information Systems Programme, Universitas Sains Indonesia, Bekasi, Indonesia*

ABSTRACT
Command and Control (C2) systems increasingly rest on Big Data infrastructure and automated decision-support tools, yet their design is rarely examined as a determinant of Meaningful Human Control (MHC), a concept central to International Humanitarian Law (IHL) debates on autonomous weapon systems. This paper argues that infrastructure is not a neutral substrate but an active variable shaping whether human oversight remains substantive or becomes procedural. Combining a review of MHC and IHL literature with an original interview-based design study conducted with Indonesian Air Force C2 personnel, the paper identifies three infrastructure-level determinants of human control: data latency and fusion speed, the degree of automation in decision-support layers, and human-machine interface architecture. These are synthesised into a framework for evaluating C2 infrastructure against MHC criteria. Responsible C2 modernisation, the paper concludes, requires evaluating infrastructure decisions for their effect on human accountability and control, not performance alone.

ABSTRAK
Sistem Komando dan Kendali (C2) semakin dibangun di atas infrastruktur Big Data dan perkakas pendukung keputusan otomatis, tetapi desainnya jarang ditelaah sebagai penentu Kendali Manusia yang Bermakna (Meaningful Human Control/MHC), konsep sentral dalam wacana Hukum Humaniter Internasional (HHI) tentang sistem senjata otonom. Artikel ini berargumen bahwa infrastruktur bukan substrat yang netral, melainkan variabel aktif yang menentukan apakah pengawasan manusia tetap substantif atau sekadar prosedural. Dengan memadukan kajian literatur MHC dan HHI dengan studi desain orisinal berbasis wawancara bersama personel C2 TNI Angkatan Udara, artikel ini mengidentifikasi tiga penentu kendali manusia di tingkat infrastruktur: latensi data dan kecepatan fusi, derajat otomasi pada lapisan pendukung keputusan, dan arsitektur antarmuka manusia-mesin. Ketiganya disintesis menjadi kerangka untuk mengevaluasi desain infrastruktur C2 terhadap kriteria MHC. Artikel ini menyimpulkan bahwa modernisasi C2 yang bertanggung jawab menuntut evaluasi keputusan infrastruktur terhadap dampaknya pada akuntabilitas dan kendali manusia, bukan kinerja semata.

CONTACT
muhammad.aria@lecturer.sains.ac.id

KEYWORDS
Meaningful Human Control, Command and Control Systems, Big Data Infrastructure

INTRODUCTION

Contemporary command and control (C2) modernisation is, at its core, a data-infrastructure enterprise. The capabilities that militaries now describe as "AI-enabled" — accelerated targeting, multi-source sensor fusion, and predictive situational awareness — are delivered not by weapons in isolation but by the pipelines, storage layers, network fabrics, and interface software that ingest, transform, and present battlefield data. A modern intelligence,

surveillance, and reconnaissance (ISR) chain, for example, moves raw optical or signals data through deep-learning detection, geolocation, and real-time geospatial mapping before any human ever sees a candidate target (Chougule et al., 2025), while hierarchical architectures distribute decision functions across cloud, edge, and terminal layers to manage volume and tempo (Jia et al., 2026). Yet debates over the legality and ethics of these capabilities have concentrated overwhelmingly on the weapon or the algorithm at the point of engagement, treating the infrastructure beneath as a neutral technical substrate.

This paper argues the opposite: infrastructure is an active determinant of Meaningful Human Control (MHC). MHC (the requirement that humans retain substantive, responsibility-bearing control over the use of force, and the organising concept of IHL discourse on autonomous and semi-autonomous weapon systems) is typically operationalised at the level of the human-machine decision (Wood, 2026; Abbasi & He, 2026). But the conditions under which a human decision is substantive rather than procedural are set upstream, by architectural choices about how fast data arrives, how much interpretation is automated before it reaches an operator, and how information is rendered for human review. An operator who is nominally "on-the-loop" but who is handed a ranked, confidence-scored recommendation with seconds to respond is exercising a very different form of control from one who is handed contestable evidence and time to weigh it, even though both sit at the same point in the decision chain.

The two literatures that bear on this problem rarely meet. The MHC and IHL literature is rich on doctrine and on the algorithm, but tends to abstract away the engineering of the systems that produce the information on which judgement depends. The C2 and data-infrastructure literature is rich on performance (latency, throughput, fusion accuracy, resilience) but seldom treats those same properties as legally or ethically loaded. This paper bridges them, taking as its unit of analysis the architecture rather than the algorithm or the weapon.

Concretely, the paper isolates three infrastructure-level factors that influence whether human control is preserved or eroded: data latency and fusion speed, which set the window available for human judgement; the degree of automation embedded in decision-support and target-recommendation layers, which positions the operator either to reason or to defer; and human-machine interface (HMI) architecture, which determines whether operators receive interpretable, contestable information or opaque outputs. The analysis draws on a synthesis of the MHC and IHL literature and on an original interview-based design study of Big Data C2 infrastructure conducted for an Indonesian Air Force context, and it distils the result into a proposed framework for evaluating C2 infrastructure against MHC criteria. The framework is offered as a starting point for infrastructure-conscious policy and system design, not as a metric or a finished legal test.

The remainder of the paper is organised as follows. The Methods section sets out the review design and the applied case. The Results and Discussion section synthesises the literature, develops the three infrastructure determinants, reads them against an applied Indonesian Air Force C2 design, consolidates them into an evaluation framework, and draws out the implications and limitations. The Conclusion restates the argument and its consequences for responsible C2 modernisation.

METHODS

This article employs a qualitative, integrative research design that pairs a narrative review of the scholarly literature with an original mixed-methods field study. This design suits the aim of synthesising dispersed, multidisciplinary work into a coherent analytic contribution; it is not a systematic review and makes no claim to statistical generalisability. Three bodies of source material were consulted: (i) scholarship on Meaningful Human Control (MHC) and International Humanitarian Law (IHL), including foundational philosophical statements and institutional position papers; (ii) the engineering literature on Big Data command and control (C2) architectures, decision-support systems, and intelligence, surveillance, and reconnaissance (ISR) pipelines; and (iii) reported accounts of AI-enabled targeting in contemporary conflict. Sources were identified through targeted searches of major scholarly indexes and full-text collections (including Scopus, IEEE Xplore, SpringerLink, and Google Scholar) using combinations of terms such as "meaningful human control", "autonomous weapon systems", "AI decision support", "command and control", and "Big Data military infrastructure", supplemented by citation chaining from key works. Inclusion required that a source be peer-reviewed or a primary institutional document, published in English, and bear directly on the relationship between infrastructure-level design choices and the substance of human control. Works published principally between 2018 and 2026 were prioritised, with earlier foundational works retained where they originate concepts used in the analysis. Investigative and reported accounts were treated as indicative rather than definitive, and are flagged as such where they are used.

The field-study component consisted of multi-echelon, semi-structured interviews conducted by the author with four participants at the strategic, operational, and technical levels of Indonesian Air Force C2 modernisation

(coded P1–P4), using a combined instrument of 50 qualitative and 30 quantitative (Likert-scale) items per echelon. Qualitative responses were analysed thematically and triangulated across echelons; Likert responses were tallied as descriptive counts and treated as indicative. All participants gave written informed consent and are reported under anonymised codes.

Sources were synthesised thematically. The resulting distinctions were then read against an applied Big Data C2 infrastructure design developed for an Indonesian Air Force context, which serves as an anchoring case. Analytically, the study proceeds in four steps: it consolidates the literature into a synthesis matrix; isolates three infrastructure determinants of human control; reads each determinant against the applied design; and consolidates the determinants into an evaluation framework whose dimensions are then checked against the applied findings. Because the interview sample is small and purposive, the applied findings are treated as illustrative and are bounded by the limits of a single national and service context.

RESULTS AND DISCUSSIONS

This section synthesises the reviewed literature, develops the three infrastructure determinants, reads each against the applied Indonesian Air Force design, and consolidates the result into an evaluation framework. Table 1 sets out how the three bodies of work that bear on the problem — the normative MHC and IHL scholarship, the critique of AI-enabled decision-support systems, and the engineering of Big Data C2 architectures — each illuminate part of the question while leaving the infrastructure determinants of human control unaddressed, which is the gap this paper occupies.

Table 1. Literature synthesis: contributions and gaps across the strands bearing on infrastructure and human control.

Literature strand	Core contribution	Treatment of infrastructure	Gap this paper addresses
MHC theory and doctrine (Horowitz & Scharre, 2015; Santoni de Sio & van den Hoven, 2018; Ekelhof, 2019; Amoroso & Tamburrini, 2020; Boulanin et al., 2020; ICRC, 2021; Wood, 2026)	Defines MHC and the will-extending / will-offloading and policy- / firing-loop distinctions; locates control in reason-responsive human judgement rather than mere human presence.	Control is theorised at the level of the decision and the weapon; infrastructure enters only implicitly, through ex ante constraints.	Does not analyse how data and network architecture set the conditions under which that judgement is exercised.
AI-DSS and proportionality (Abbasi & He, 2026; Lushenko & Sparrow, 2026)	Shows AI-DSS reshape human judgement through salience, bias, opacity, automation bias, and tempo; user trust is conditional and not self-justifying.	Identifies the mechanisms but frames them at the tool and decision level, not as properties of the surrounding architecture.	Does not translate the mechanisms into design levers of the data infrastructure.
C2 and Big Data engineering (Liu et al., 2024; Chougule et al., 2025; Jia et al., 2026)	Specifies layered C2 architectures (cloud-edge-terminal H-OODA; ISR fusion pipelines) and their tempo and automation behaviour.	Rich on performance (latency, throughput, fusion) and even names the autonomy / oversight balance, but treats it as an engineering problem.	Does not evaluate those same properties against a legal or ethical human-control standard.
Resilience, governance, and cross-domain (Boulanin et al., 2020; Cole et al., 2026; Tofan et al., 2025; Madjid & Yuliani, 2026)	Documents infrastructure fragility, practical elements of human control, and the recurrence of the control problem across domains and governance.	Establishes that fragility and institutional setting condition control, and that practical control elements can be specified.	Leaves unspecified a design and evaluation framework linking infrastructure choices to the preservation of control.
This paper	Treats infrastructure as an active determinant of MHC and isolates three determinants: tempo, automation locus, and interface contestability.	Architecture is the unit of analysis rather than the algorithm or the weapon.	Bridges the strands in an infrastructure-conscious MHC evaluation framework, read against an Indonesian Air Force C2 design.

Meaningful Human Control, AI-DSS, and the C2 Architectural Stack

Meaningful Human Control in IHL Discourse

The concept of MHC originated in the regulatory debate over autonomous weapon systems (AWS) and has since become central to how scholars and states reason about the human role in the use of force (Wood, 2026). It gained early definitional form in policy-oriented work identifying informed decision, sufficient information, and effective control as its essential components (Horowitz & Scharre, 2015), and it now anchors institutional positions, most prominently the ICRC's call for legally binding limits on autonomous weapon systems built around requirements for human control and judgement (ICRC, 2021). The concept was given systematic philosophical form by Santoni de Sio and van den Hoven (2018), whose "tracking" and "tracing" conditions require that a system remain responsive to the relevant human reasons and that its states be traceable to a human moral understanding, an account on which much later work, Wood's included, builds. Parallel efforts pushed the debate beyond semantics toward the operational level (Ekelhof, 2019), sought to specify the practical elements of human control in fielded systems (Boulanin et al., 2020), and mapped its ethical and legal dimensions (Amoroso & Tamburrini, 2020). A common taxonomy distinguishes systems in which a human plays a direct role in the decision ("in-the-loop"), systems that operate but remain subject to persistent human override ("on-the-loop"), and systems that a human cannot guide or override once deployed ("off-the-loop"). The same tripartite scheme recurs in the engineering literature on intelligent command and control, which classifies command modes as participatory ("human-in-the-loop"), supervisory ("human-on-the-loop"), and autonomous ("human-outside-the-loop") decision-making (Liu et al., 2024). Much of the critical literature has treated contemporaneous human oversight as a precondition for permissible use, such that off-the-loop systems appear ruled out almost by definition.

Wood resists this strong interpretation. He maintains that meaningful control does not depend on contemporaneous oversight; rather, it can be established ex ante, through the selection of constrained hardware, the imposition of operational boundaries (such as geographic and temporal limits), and the specification of targeting parameters that restrict the spectrum of possible outcomes (Wood, 2026). Drawing on Welsh's distinction (as cited in Wood, 2026) between the "policy loop", the mechanical rules the system follows, and the "firing loop", the execution of those rules, Wood argues that MHC turns on the policy loop. A system may therefore be entirely absent from the firing loop at the moment of engagement yet remain a genuine extension of the deployer's will, provided that the deployer has competently circumscribed its permissible conduct. This reasoning yields his core dichotomy: *will-extending* systems, which amplify human intent and remain within the bounds of MHC, versus *will-offloading* systems, into which humans effectively surrender their judgement.

Central to the present argument is that Wood identifies AI-DSS as the archetypal will-offloading case. Because such systems present more-informed-seeming assessments first, and because humans are prone to defer to them, they can lead an operator to subvert their own judgement even when every strike is nominally human-approved. The operator becomes, in the metaphor associated with Sparrow and Henschke (2023) and adopted by Wood, a "minotaur" rather than a "centaur". The same conclusion is reached from the doctrinal side. Proportionality — the prohibition on attacks whose expected incidental civilian harm would be excessive relative to the concrete and direct military advantage anticipated — is not a calculation but a context-sensitive human judgement, assessed on what a reasonable commander knew or should have known at the time of the attack (Abbasi & He, 2026). It cannot be reduced to numerical thresholds or automated outputs. Both literatures thus locate MHC in the quality of human reasoning, not merely in the presence of a human. (For completeness, the term has partly migrated in official fora toward "context-appropriate human judgement and control", but the underlying intuition is continuous with MHC (Wood, 2026).)

How AI-DSS Reconfigure Human Judgement

AI-DSS do not select or engage targets autonomously; they shape the informational conditions under which humans decide (Abbasi & He, 2026), which is what makes them consequential for MHC. Abbasi and He identify five mechanisms through which such systems reshape proportionality reasoning even where formal human control persists: informational salience, whereby the system frames what appears relevant or urgent; embedded bias, inherited from training data and model design; opacity, the black-box problem of outputs that cannot be understood or reconstructed; automation bias, the tendency to over-trust machine outputs under time pressure, a phenomenon documented in the human-factors literature well before the current AI debate (Parasuraman & Riley, 1997); and accelerated decision-making, which narrows the time available to question, contextualise, or override.

Reported practice illustrates the stakes. In Gaza, systems such as Lavender, The Gospel, and Where's Daddy? have been described as scoring individuals for suspected affiliation, generating target lists at machine speed, and cueing strikes on residences; investigative accounts report probability-based lists on the order of tens of thousands of individuals, review reduced in some accounts to a few seconds per case, and standardised collateral-harm thresholds applied rather than case-by-case assessment (Abraham, 2024; Abbasi & He, 2026). These accounts are investigative rather than adjudicated findings, and Israeli authorities have disputed aspects of them; they are used here, consistent with the source-handling approach set out in the Methods, as indicative illustrations of how recommendation-layer systems structure human review. In Ukraine, data-integration platforms fuse satellite imagery, drone footage, and sensor data into a near-real-time common picture that structures what commanders treat as important (Abbasi & He, 2026). In each case, the system does not pull the trigger; it reshapes the judgement of the person who does.

User attitudes reinforce the point that human presence is not a proxy for human control. In a survey experiment among U.S. military cadets, Lushenko and Sparrow (2026) find that a younger cohort often assumed to be uncritical "digital natives" is in some respects more trusting of AI than senior officers, yet more conservative about its appropriate use, preferring configurations with human oversight and strategic-level, rather than machine-controlled tactical, decision-making. They caution, further, that trust in a system is a distinct question from whether that trust is justified, and that automation bias may mean humans nominally "on-the-loop" do not in practice exercise meaningful control over the systems they are supposed to supervise. The presence of a human, in short, guarantees neither substantive control nor sound judgement.

The C2 Architectural Stack

These decision dynamics do not float free; they are produced by an architecture. Jia et al. (2026) formalise a representative recent design, a hierarchical Observe-Orient-Decide-Act (H-OODA) framework layered across cloud, edge, and terminal tiers and implemented with network function virtualisation. Their own analysis makes the infrastructure-tempo relationship explicit: batch processing of edge-layer data inherently introduces latency, which predictive caching can mitigate, and their experiments show that deeper decision loops improve service quality and lower error rates. They also list the "balance of autonomy and human supervision" as a core open challenge, with explainable AI, human-machine collaboration, and ethical decision frameworks named as the required directions, an admission from the engineering literature itself that architectural choices carry governance weight. Liu et al. (2024), surveying envisioned "intelligent" C2 modes, make the same division explicit from the other direction: they cast machines as "advisors" that shoulder computationally intensive fusion and situational analysis while commanders retain strategic judgement and decision authority, an engineering account that already separates the automation of analysis from the automation of the decision, prefiguring the plumbing-versus-recommendation distinction developed below.

Chougule et al. (2025) supply a concrete pipeline at the terminal end of such a stack: optical sensors feed deep-learning object detection (using architectures such as YOLO and Faster R-CNN), whose pixel outputs are transformed into real-world coordinates via photogrammetry, homography, triangulation, and distortion correction, and then plotted on a dynamic geospatial map that updates in near real time and can be fed into C2 systems and predictive-analytics tools. The design goal is explicitly speed and situational awareness. The same pipeline, however, is a chain of automated interpretation steps — detect, classify, localise, prioritise — each of which fixes what a downstream human will treat as a candidate object of force.

The fragility of these architectures under real conditions is itself a control problem. Cole et al. (2026), writing on military medicine, catalogue how AI-dependent systems presuppose connectivity, cloud compute, and data integrity that are often unavailable in austere or contested environments; how degraded communications, jamming, spoofing, and cyber-attack can corrupt or disable exactly the tools most needed under stress; and how centralised platforms create single points of failure. Their recurring mitigation is human-in-the-loop oversight, together with hardened, locally operable fallbacks and "cold-start" procedures, and their recurring warning is that automation bias and opaque outputs can produce confusion, mistrust, or moral injury when machine recommendations contradict a professional's own judgement. Taken together, these three works show that the C2 stack is real, layered, and consequential, and that its salient properties (tempo, degree of automation, opacity, and resilience) are exactly the variables on which MHC turns.

Nor is this a military-only concern. The requirement that humans retain meaningful control over consequential machine-mediated judgement recurs across domains: in military medicine, where the same infrastructure fragilities and automation-bias risks apply (Cole et al., 2026), and in civilian professions such as tax consulting, where AI increasingly performs technical work while practitioners are urged to retain advisory

judgement, representation, and accountability, with black-box opacity flagged as a core risk (Tofan et al., 2025). This cross-domain recurrence supports treating infrastructure-level control as a general property of responsible system design rather than a concern unique to the law of war.

Three Infrastructure-Level Determinants of Human Control

The literature reviewed above converges on a claim rarely stated as an infrastructure claim: the substance of human control is set upstream of the decision, by architecture. This section develops that claim through three determinants — tempo, automation locus, and interface contestability — arguing in each case that a design property ordinarily evaluated for performance is simultaneously a control variable, and that Wood's (2026) will-extending / will-offloading distinction supplies the criterion for evaluating it.

Data Latency and Fusion Speed

The first determinant is temporal. The window available for human judgement is not a fixed human constant; it is an output of the architecture. Abbasi and He (2026) argue that accelerated decision-making narrows the time to question, reinterpret, or contextualise a system's output, and that when there is insufficient time to do so, reliance on the output can bypass the reasonability standard on which proportionality depends, a dynamic vividly reflected in reports of seconds-per-case review at scale. Jia et al. (2026), from the engineering side, show that the cloud-edge-terminal architecture itself sets tempo: it introduces latency where global integration is needed and compresses it where predictive caching or edge processing is available. Tempo, in other words, is a design parameter.

Wood's policy-loop framing sharpens what is at stake. If humans have genuinely set the rules ex ante, control can persist even off-the-loop. But that argument presupposes that the policy loop is actually exercised, that humans have the standing and the time to constrain the system meaningfully. Where fusion speed outpaces an operator's capacity to interrogate outputs at decision time, the policy loop degrades into a rubber stamp on the firing loop, which is exactly the will-offloading condition Wood warns against. Latency and fusion speed are therefore not neutral efficiency properties: they determine whether the policy loop is substantive or nominal.

The design corollary is that fusion speed should be bounded by the human review budget for decisions that bear on the use of force, not maximised for its own sake. Whether fusion is streaming or batched, how alerting cadence is tuned, and whether enforced review windows are built into force-bearing workflows are all control-relevant architectural levers, and they trade directly against the throughput metrics by which such systems are usually optimised.

Degree of Automation in Decision-Support and Target-Recommendation Layers

The second determinant concerns where in the stack automated decision-content is produced. It is useful to distinguish two kinds of automation that are often conflated. The first automates data plumbing: ingesting, cleaning, correlating, geolocating, and fusing sensor data. The second automates the recommendation itself: scoring, ranking, prioritising, or nominating candidate targets. These differ in kind for MHC, not merely in degree.

Automation of the first kind can be will-extending. Wood (2026) notes that richer multi-modal sensing, fed through well-specified target parameters, can make a system more predictable and more tightly bound to the deployer's intent, an amplification of will rather than a surrender of it. Automation of the second kind tends toward will-offloading. A ranked target list or a confidence-scored nomination positions the human to defer, engaging precisely the automation bias and salience effects that Abbasi and He (2026) identify and that Lushenko and Sparrow (2026) find can hollow out nominal on-the-loop oversight. The reported Lavender pattern illustrates the limiting case: probability scores generated across a very large population, with human review in some accounts reduced to a cursory check (Abraham, 2024; Abbasi & He, 2026).

The design corollary is to keep the recommendation layer evidentiary rather than conclusory. Automating fusion differs from automating judgement: a system can surface the basis for a candidate assessment, including the underlying evidence, its calibrated uncertainty, and any disconfirming data, rather than a rank or score presented as a verdict. Where automated decision-content is produced, and how authoritatively it is framed, is thus a first-order architectural choice with direct MHC consequences.

Human-Machine Interface Architecture

The third determinant is the interface: whether operators receive interpretable, contestable information or opaque outputs. Abbasi and He (2026) decompose opacity into two failures, a lack of interpretability, where the

operator cannot understand or predict why an output was produced, and a lack of traceability, where the output cannot be reconstructed after the fact. The second failure is doubly damaging: it disables both real-time scrutiny and subsequent accountability, complicating investigation into alleged violations. Cole et al. (2026) show the human cost of interfaces that fail on interpretability: confusion, mistrust, and moral injury when opaque outputs contradict a professional's own judgement; and Jia et al. (2026) themselves list explainable AI as a necessary condition for balancing autonomy against human supervision.

Contestability, on this account, is the operational substance of MHC. The ability to question, override, and audit an output is what converts nominal supervisory status into genuine control. An interface that renders a confidence score without provenance, without an expression of uncertainty, and without a low-friction override path turns "on-the-loop" oversight into procedural compliance, whatever the doctrine on paper. Conversely, an interface designed so that its outputs are questionable, exposing provenance, calibrated uncertainty, and disconfirming evidence, with an accessible override and immutable decision logging, preserves both real-time contestability and the ex post traceability that accountability requires. Interface architecture, ordinarily treated as a usability concern, is in this light a determinant of whether human control is real.

Applied Perspective: A Big Data C2 Infrastructure Design in an Indonesian Air Force Context

This section reads the three determinants against the design of a Big Data C2 infrastructure developed for an Indonesian Air Force context. The aim is not to reconstruct operational procedures but to show how the determinants surface as concrete design tensions in a real modernisation effort, and to locate where the general argument meets specific engineering choices.

The field study develops a Big Data C2 information-system model for the Indonesian Air Force from the interviews described in the Methods section. Two of its findings bear directly on the present argument. First, it identifies human involvement as a structural rather than transitional feature of the current architecture: because systems procured from multiple Western and Eastern suppliers do not interoperate, human operators function as "active middleware", manually translating formats and reconciling sensor outputs across incompatible systems, with technical-level respondents estimating existing decision support to be roughly one-half to three-quarters manual. Human control here is a technical necessity produced by the infrastructure rather than a doctrinal preference, a direct instance of this paper's claim that architecture sets the substance of control. Second, national data sovereignty emerges as a non-negotiable, overarching design constraint, rated at the maximum (5/5) on the strategic-level sovereignty assessment and extending, in respondents' accounts, to a requirement that AI components be built domestically for confidentiality reasons; sovereignty therefore conditions the automation and interface choices examined below.

A canonical C2 Big Data stack for this setting comprises four stages: ingestion of heterogeneous ISR and sensor feeds; fusion, where feeds are correlated, geolocated, and resolved into tracks and a common operating picture; analytics and decision-support, where objects are classified, scored, and surfaced for attention; and a human-machine interface through which operators act, feeding the C2 decision cycle. The operational environment is distinctive: a wide-area, archipelagic theatre with heterogeneous and intermittent sensor coverage and connectivity, which mirrors closely the latency-prone, austere, and contested conditions that Cole et al. (2026) and Jia et al. (2026) treat as the hard case for AI-dependent systems.

Tempo: Where the Review Window Sits

The first design tension is between fusion speed and review time. A streaming fusion pipeline maximises currency of the common operating picture but compresses the interval before an operator must respond; a batched or tiered pipeline (as in the cloud-edge-terminal pattern) trades currency for the latency headroom in which human judgement can occur. The control-relevant decisions are where the human review window is placed in the workflow, whether it is enforced for force-bearing outputs rather than advisory ones, and how alerting cadence is tuned to comprehension rather than throughput. In the field study, the proposed fusion engine targets sub-second common-operating-picture updates, yet the binding tempo constraint in the archipelagic setting is connectivity rather than compute: respondents repeatedly located the bottleneck in bandwidth and data-service infrastructure at remote sites, and operational-level respondents reported that, under time pressure, for example within a scrambled response to an airspace-violation scenario, informal channels were often faster and more effective than the formal C2 path. That operators route around the designed workflow when it is too slow is an empirical instance of tempo overriding the review window, and it underscores that where the human review interval sits, and whether it is enforced for force-bearing rather than advisory outputs, is a design choice rather than an artefact.

Automation Locus: Plumbing Versus Recommendation

The second tension is where automation is sited. Automating ingestion, correlation, and geolocation is defensible as will-extending: it widens the evidence an operator can weigh. Automating the recommendation, producing a ranked or scored candidate list, raises the will-offloading risk. A defensible design keeps the recommendation layer evidentiary: surfacing the basis, uncertainty, and disconfirming data for any candidate rather than a bare verdict. The field study concentrates automation in the data-plumbing layers, a protocol adapter/gateway that normalises legacy Western and Eastern formats into a common protocol, and an ETL-based fusion engine that resolves heterogeneous feeds into a single common operating picture, while its decision-support layer is explicitly specified to keep the human as final decision authority, presenting AI output as probabilistic situational awareness rather than as an executable verdict. Its stated design sequencing makes the priority plain: build the Big Data infrastructure that supports the human integration role first, and extend automation to well-defined decision nodes only as interoperability matures. In the vocabulary of the automation-locus determinant, this is a deliberate choice to automate plumbing before recommendation and to keep the recommendation layer evidentiary.

Interface: Contestability and Logging

The third tension is interface design. The determinant predicts that dashboards should expose provenance and calibrated uncertainty, provide a low-friction override, and log decisions immutably for traceability. The field study's interface findings foreground usability under stress: technical respondents prioritised a simple interface, real-time synchronisation, offline capability in remote areas, and high security, and the proposed decision-support layer accordingly specifies a simplified emergency-mode interface, mobile and portable-device integration, and role-based (need-to-share) access. The study does not yet specify a provenance-and-uncertainty display, an override mechanism, or an immutable decision-logging scheme; the interface determinant developed above marks these as the additions that would convert a usable dashboard into a contestable and auditable one. Its central finding sharpens the stakes: where the interface and the underlying protocols fail to integrate, operators absorb the gap as manual "middleware", which is precisely the interpretability failure — humans performing the reconciliation the architecture does not — that a contestable interface is meant to remove.

Finally, a governance caveat frames the applied claim. Even a well-designed infrastructure operates within institutional and political arrangements that can diffuse or constrain accountability. Madjid and Yuliani (2026) document how, at the level of international governance, political prerogative can override collective safeguards and how enforcement of humanitarian norms can become contingent on will rather than rule. Infrastructure design cannot substitute for institutional accountability. What infrastructure can do — and what the determinants above are meant to secure — is to make human control legible, contestable, and auditable, so that responsibility remains attributable to identifiable humans even as systems accelerate.

An Infrastructure-Conscious MHC Evaluation Framework

The three determinants can be consolidated into a framework for evaluating C2 infrastructure design against MHC criteria, at both design time and review time. Its organising criterion is Wood's (2026) will-extending / will-offloading distinction: for each infrastructure dimension, the framework contrasts the pattern that pushes a system toward offloading human will with the design that keeps it extending human will, pairs each with a diagnostic question, and anchors that question in the relevant IHL standard. Because the interface determinant comprises two operationally distinct facets, interpretability in real time and traceability after the fact, the framework separates them into distinct evaluation dimensions (interface and auditability), yielding four dimensions from the three determinants. Table 2 presents the framework.

Table 2. An infrastructure-conscious framework for evaluating C2 systems against Meaningful Human Control criteria.

Infrastructure dimension	Erosion pattern (tends toward will-offloading)	Preservation design (tends toward will-extending)	Diagnostic question and IHL anchor
Tempo: data latency and fusion speed	Fusion outpaces the operator's capacity to interrogate outputs; review collapses to seconds; the policy loop is nominal.	Fusion speed bounded by an explicit human review budget; enforced review windows for force-bearing decisions; alert cadence tuned to comprehension, not throughput.	Does the architecture guarantee enough time for a reasoned, good-faith judgement? (Proportionality as a foresight-based judgement; Art. 57 precautions.)

Infrastructure dimension	Erosion pattern (tends toward will-offloading)	Preservation design (tends toward will-extending)	Diagnostic question and IHL anchor
Automation locus: where decision-content is produced	Automated scoring, ranking, or target nomination presented as conclusory; the operator is positioned to defer.	Automation concentrated in data plumbing/fusion; the recommendation layer stays evidentiary—surfacing basis, uncertainty, and disconfirming data rather than a verdict.	Is the human extending will or offloading it? Is the output a tool or a decision? (Will-extending vs. will-offloading systems.)
Interface: interpretability and contestability	Opaque outputs—confidence scores without provenance, uncertainty, or a low-friction override path.	HMI exposes provenance, calibrated uncertainty, and disconfirming evidence; override is accessible and low-friction; outputs are questionable by design.	Can the operator understand, question, and override the output? (Interpretability; effective human oversight.)
Auditability: traceability and accountability	Non-traceable outputs; decisions cannot be reconstructed for after-action or investigation; accountability diffuses.	Immutable, decision-linked logging of inputs, model version, and human action; traceability preserved across the system life cycle.	Can a decision be reconstructed and attributed to a responsible human? (Art. 36 review scope; state responsibility.)

The framework is intended to be applied at two moments. At design time, each dimension is treated as a first-class requirement alongside performance: the human review budget is set first and fusion speed bounded to it; automation is concentrated in data plumbing while the recommendation layer is kept evidentiary; interfaces are built to be contestable; and decisions are logged for traceability. At review time, the framework maps onto existing legal machinery. Abbasi and He (2026) argue that AI-DSS fall within the scope of the Article 36 obligation to review new means and methods of warfare, and that because machine-learning systems can change after deployment, such reviews must be living rather than one-off. Extending Article 36 review to fusion and interface architecture, not merely to the terminal weapon, follows directly, and is consistent with the ICRC's position that legal review must account for whether a system's effects can be sufficiently understood, predicted, and explained (ICRC, 2021). The framework also connects to the Article 57 duty to take precautions in attack (to verify the target, minimise civilian harm, and suspend an attack that becomes excessive), the feasibility of which is, in an AI-mediated workflow, partly a property of the infrastructure: an operator cannot take a precaution the interface and the tempo do not permit.

The framework is deliberately qualitative. It is a structure for infrastructure-conscious deliberation, not a scoring metric, and it is offered as a starting point to be refined and, ultimately, validated empirically.

Implications, Limitations, and Future Directions

For system designers and acquirers, the central implication is that performance-only evaluation of C2 infrastructure is insufficient. Tempo, automation locus, interface contestability, and traceability are control variables, and evaluating them requires the kind of structured attention the proposed framework is meant to provide. The point is not to slow C2 systems as an end in itself; it is to recognise that some architectural choices which improve throughput simultaneously erode the human control the system is legally required to preserve, and to make that trade explicit rather than incidental.

For IHL and policy, the implication is that MHC assessment should look upstream of the weapon. If AI-DSS influence proportionality by shaping what commanders know, should know, and have time to consider (Abbasi & He, 2026; Wood, 2026), then weapons review, procurement, and certification should reach the decision-support infrastructure itself—its fusion tempo, its automation locus, and its interface—and should require the traceability without which after-action accountability is impossible. This reframes a set of questions usually treated as engineering detail as questions of legal compliance.

At the same time, infrastructure design has limits that it cannot transcend. Governance can be politically constrained, and the enforcement of humanitarian norms can be overridden by prerogative rather than settled by rule (Madjid & Yuliani, 2026). Infrastructure that makes control legible and auditable supports accountability but does not guarantee it; the framework is a necessary, not a sufficient, condition for responsible modernisation. The cross-domain recurrence of the same logic in military medicine (Cole et al., 2026) and in civilian professions

where AI mediates high-stakes judgement (Tofan et al., 2025) suggests, however, that the underlying design principle generalises well beyond the law of armed conflict.

Several limitations bound these claims. The applied dimension rests on a small, purposive interview sample within a single national and service context, and the framework itself has not yet been empirically validated, and its dimensions are analytic distinctions that require testing against operational data. The account of fielded systems relies, as in the underlying literature, on public and reported sources whose technical and operational detail is partial and in places contested (Abraham, 2024; Abbasi & He, 2026). Future work should instrument the human review window directly, conduct human-factors studies of contestable C2 interfaces, and validate the framework across services and in coalition or multinational settings, where interoperability and divergent attitudes to human oversight introduce further variation (Lushenko & Sparrow, 2026).

CONCLUSIONS

This paper has argued that infrastructure is a determinant of human control, not a neutral substrate beneath it. Whether the oversight exercised over a C2 system is substantive or merely procedural is decided upstream of the human-machine decision, by architectural choices about tempo, about where automated decision-content is produced, about whether interfaces are interpretable and contestable, and about whether decisions are traceable. Reading the MHC and IHL literature alongside the engineering of modern C2 stacks shows that these are the same properties: the variables engineers optimize for performance are the variables on which meaningful human control depends. Responsible C2 modernisation therefore requires that infrastructure decisions be evaluated not only for what they do to speed and situational awareness, but also for their downstream effect on human accountability and control. The framework proposed here is offered as a starting point for that evaluation, an infrastructure-conscious bridge between data-infrastructure engineering and IHL and security studies, to be refined, contested, and validated in the work that follows.

REFERENCES

Abbasi, A. Y., & He, J. (2026). AI-enabled decision support systems and the reconfiguration of proportionality in international humanitarian law. *Computer Law & Security Review, 61*, 106319. https://doi.org/10.1016/j.clsr.2026.106319

Abraham, Y. (2024, April 3). 'Lavender': The AI machine directing Israel's bombing spree in Gaza. *+972 Magazine*. https://www.972mag.com/lavender-ai-israeli-army-gaza/

Amoroso, D., & Tamburrini, G. (2020). Autonomous weapons systems and meaningful human control: Ethical and legal issues. *Current Robotics Reports, 1*(4), 187–194. https://doi.org/10.1007/s43154-020-00024-3

Boulanin, V., Davison, N., Goussac, N., & Peldán Carlsson, M. (2020). *Limits on autonomy in weapon systems: Identifying practical elements of human control.* Stockholm International Peace Research Institute & International Committee of the Red Cross. https://www.sipri.org/sites/default/files/2020-06/2006_limits_of_autonomy.pdf

Chougule, R., Patil, V., Rane, A., Mangla, M., Hole, P., & Sankhe, M. (2025). AI-driven optical intelligence for modern warfare command and control. In *2025 9th International Conference on Computing, Communication, Control and Automation (ICCUBEA)* (pp. 1–6). IEEE. https://doi.org/10.1109/ICCUBEA65967.2025.11283928

Cole, R., Simmons, S., Duncan, J., Cole, J., & Peacock, J. G. (2026). Readying military medicine for AI-enabled warfare. *Military Medicine, 191*(7–8), 122–126. https://doi.org/10.1093/milmed/usaf460

Ekelhof, M. (2019). Moving beyond semantics on autonomous weapons: Meaningful human control in operation. *Global Policy, 10*(3), 343–348. https://doi.org/10.1111/1758-5899.12665

Horowitz, M. C., & Scharre, P. (2015). *Meaningful human control in weapon systems: A primer.* Center for a New American Security. https://www.cnas.org/publications/reports/meaningful-human-control-in-weapon-systems-a-primer

International Committee of the Red Cross. (2021). *ICRC position on autonomous weapon systems: ICRC position and background paper*. https://www.icrc.org/en/document/icrc-position-autonomous-weapon-systems

Jia, Z., Wu, Y., Wu, Q., He, L., Zhu, Q., Zhou, F., & Han, Z. (2026). Hierarchical observe-orient-decide-act enabled UAV swarms in uncertain environments: Frameworks, potentials, and challenges. *IEEE Network*. https://doi.org/10.1109/MNET.2026.3671364

Liu, S., Hu, P., Du, G., Li, G., Zhang, H., & Zhao, X. (2024). Intelligent command and control and its envisioned future typical modes. In Chinese Institute of Command and Control (Ed.), *C2 2023* (Lecture Notes in Electrical Engineering, Vol. 1124, pp. 540–554). Springer. https://doi.org/10.1007/978-981-99-9021-4_50

Lushenko, P., & Sparrow, R. (2026). Artificial intelligence and U.S. military cadets' attitudes about future war. *Armed Forces & Society, 52*(3), 1027–1050. https://doi.org/10.1177/0095327X241284264

Madjid, M. A., & Yuliani, A. N. (2026). War without end: Structural limitations and the United Nations' position on the Israel–Palestine conflict. *Ianna Journal of Interdisciplinary Studies, 8*(2), 273–285. https://doi.org/10.5281/zenodo.20229927

Parasuraman, R., & Riley, V. (1997). Humans and automation: Use, misuse, disuse, abuse. *Human Factors, 39*(2), 230–253. https://doi.org/10.1518/001872097778543886

Santoni de Sio, F., & van den Hoven, J. (2018). Meaningful human control over autonomous systems: A philosophical account. *Frontiers in Robotics and AI, 5*, 15. https://doi.org/10.3389/frobt.2018.00015

Sparrow, R., & Henschke, A. (2023). Minotaurs, not centaurs: The future of manned-unmanned teaming. *Parameters, 53*(1), 115–130. https://doi.org/10.55540/0031-1723.3207

Tofan, A., Vitriya, R., & Tiaramurti, G. (2025). AI integration in tax consulting services: Technology innovation or a threat to the profession? *Risenologi, 10*(2), 205–211. https://doi.org/10.21009/risenologi.102.10

Wood, N. (2026). Autonomous and AI-enabled systems: Extensions or replacements of human will and control? *Ethics and Information Technology, 28*, Article 3. https://doi.org/10.1007/s10676-025-09876-9