

A Design Science Research Architecture for Specialized Data Centers Supporting Real-Time C5I Operations in Contested, Hyper-Connected Environments: An EWNOC-Integrated Design with Simulation-Based Evaluation

Original article

Received: 2026-09-26

Revised: 2026-10-12

Accepted: 2026-10-13

Final review: 2026-10-13

Peer review: 2026-10-01

Double blind

Keywords:

C5I; EWNOC; data-centre architecture; D-DIL; Zero Trust; tactical edge computing; design science research; availability modelling; fault injection; electromagnetic spectrum operations

This work is licensed under the Creative Commons

Attribution-NonCommercial-NoDerivatives 4.0 License

M. Rizwan Yasin^{1, A-F}

[ORCID ID 0009-0004-5499-495X](https://orcid.org/0009-0004-5499-495X)

A - Research concept and design, B - Collection and/or assembly of data, C - Data analysis and interpretation, D - Writing the article, E - Critical revision of the article, F - Final approval of the article

¹Emiri Signal Corps, Qatar Armed Forces, Qatar

Abstract

Objectives: This paper proposes, specifies and evaluates at design stage a five-layer data-centre architecture integrated with an Electronic Warfare Network Operations Center (EWNOC) for real-time Command, Control, Communications, Computers, Cyber and Intelligence (C5I) operations in contested environments subject to Denied, Disconnected, Intermittent and Limited (D-DIL) connectivity. No physical system has been built; the contribution is a design artefact together with a reproducible model-based evaluation.

Methods: Following Design Science Research (DSR), the study states three research questions, six design objectives and six testable design propositions. It specifies system boundaries, data flows, interfaces, trust boundaries and Zero Trust enforcement points. It then evaluates the artefact against a centralized baseline and an intermediate edge configuration using (i) a parameterised decision-cycle latency model with Monte Carlo uncertainty and Sobol' sensitivity analysis; (ii) an alternating-renewal Monte Carlo availability model with failover coverage, common-cause failures, thermal ride-through, electromagnetic-spectrum (EMS) denial and physical attrition across four operating scenarios; and (iii) a model-based fault-injection campaign of eleven fault classes. All parameters, units, ranges and their basis are tabulated, and the simulation code is released as supplementary material.

Results: The proposed configuration has a modelled machine-side decision-cycle latency of 0.96 ms at idle and 4.60 ms at peak load (90% interval 4.20–5.01 ms), against 8.21 ms for the edge configuration and 64.8 ms for the centralized baseline. Under mid-range planning parameters the modelled mission availability is 99.907% with a system-level MTTR of 14.2 min (baseline 99.40%, 89 min). The previously stated targets of 99.985% availability and a 12-min MTTR are therefore not met under planning assumptions. They are met (99.993%, 2.2 min) only within an explicit requirement envelope of failover coverage, common-cause separation and repair logistics. Under combined EMS, thermal and physical stress, availability falls to 99.979%. Repair time, failover coverage and common-cause fraction dominate the sensitivity.

Conclusions: The architecture is a standards-traceable design artefact with transparent and falsifiable performance claims. Its availability and recovery targets are conditional design requirements rather than intrinsic properties. Prototype-based evaluation and structured expert evaluation remain necessary before any operational claim can be made.

1. Introduction

Hyper-connected C5I environments couple sensing, communications, computing and command decision support into a single cyber-physical system. The engineering problem addressed here is narrow and facility-centred: how to bound the machine-side delay that physical transport, storage, integration, reach-back dependency and decision-support processing add to the decision cycle, while keeping that capability available when external connectivity is degraded or denied and when the facility itself is stressed. U.S. joint concepts frame the operational need through Combined Joint All-Domain Command and Control (CJADC2), which seeks to connect sensors and effectors across services into a shared, interoperable data environment (Congressional Research Service, 2022; U.S. Department of Defense, 2022b). The present paper is not concerned with weapon engagement, targeting or fire control. It treats the data centre that hosts C5I services as the object of design.

Contested multi-domain settings add a further constraint. Cyber, electronic warfare and information activities increasingly overlap and are often conducted together (Marsili, 2023). Adversarial influence and hybrid operations target information infrastructures and the confidence that commanders place in them (Marsili, 2021). The proliferation of networked sensors and devices on the battlefield, the "Internet of Battle Things", multiplies the data sources a forward facility must ingest and the attack surface it must defend (Kott et al., 2016). A facility for this environment therefore has to combine dependability engineering, edge computing and Zero Trust governance rather than optimise any one of them in isolation.

1.1 Research questions, design objectives and propositions

The first version of this work reported availability, recovery and latency figures without a reproducible basis and described the artefact as "implemented". This revision reframes the contribution as a design-stage DSR artefact and states the following research questions (RQ):

- **RQ1.** How can the machine-side decision-cycle latency of a forward C5I facility be decomposed, and which contributors dominate under load and under EMS degradation?
- **RQ2.** Under which component-reliability, failover and repair conditions can the proposed five-layer architecture meet availability and recovery targets, and how does it compare with a centralized and an intermediate edge configuration?
- **RQ3.** How do EMS denial, environmental (thermal and power) stress and physical attrition degrade those metrics, and which parameters drive the degradation?

Table 1 lists the design objectives (DO) and the testable propositions (P) derived from them. DO1–DO3 carry over the quantitative targets of the original manuscript, which are now treated explicitly as requirements to be tested. Each proposition states its threshold so that readers can re-test it with the supplementary code.

Table 1. Design objectives and testable design propositions.

ID	Design objective / proposition	Test (Section 7)
DO1	Machine-side decision-cycle latency $L_{\text{total}}(w = 1) \leq 5 \text{ ms}$	Latency model, Monte Carlo
DO2	Mission availability $A \geq 99.985\%$ ($\approx 79 \text{ min}$ downtime per year)	Availability Monte Carlo
DO3	System-level mean time to restore mission capability (MTTR) $\leq 12 \text{ min}$	Availability Monte Carlo
DO4	Mission capability is independent of reach-back connectivity (D-DIL survivability)	Scenario S1, fault F11
DO5	No implicit trust: every trust-zone crossing is authenticated, authorised and logged	Architecture specification (Section 4)
DO6	Consequential actions pass a human-authorisation gate; every layer is traceable to a standard	Architecture specification (Section 4)
P1	$P(L_{\text{total}}(1) \leq 5 \text{ ms}) \geq 0.90$ under the stated parameter uncertainty	Section 7.1
P2	Nominal $L_{\text{total}}(1)$ is $\geq 40\%$ lower than the edge configuration and $\geq 90\%$ lower than the centralized baseline	Section 7.1
P3	Within the requirement envelope, DO2 and DO3 are met in scenarios S0–S2	Section 7.2
P4	Under mid-range planning parameters, DO2 and DO3 are not met; attainment is conditional on coverage, common-cause and repair parameters	Sections 7.2–7.3
P5	Reach-back denial (S1) changes the proposed configuration's mission availability by < 0.01 percentage points, but reduces both comparators by > 0.25 points	Section 7.2
P6	For every single fault in a redundant group, expected mission outage is at least ten times lower than for the same fault in a non-redundant comparator	Section 7.4

Source: Author's own elaboration.

1.2 Contributions and scope

The paper makes four contributions: (i) a specification of the architecture that makes its system boundary, data flows, interfaces, trust boundaries and failure model explicit; (ii) a fully parameterised latency and availability model with stated units, ranges and the basis for each value; (iii) a simulation-based artefact evaluation, combining Monte Carlo uncertainty propagation, global sensitivity analysis and model-based fault injection, against two defined comparators; and (iv) a derived requirement envelope stating the conditions under which the targets can be met. The evaluation is *ex ante* and artificial in the sense defined by Venable et al. (2016). It provides model-based evidence, not empirical evidence. Operational-security considerations limit disclosure of site-specific detail, and all parameters are generic planning values.

2. Related Work and Positioning

2.1 C5I, multi-domain operations and cyber–EW convergence

CJADC2 policy documents describe the intent to federate service networks into a shared data environment, and congressional analysis has repeatedly noted that its success depends on common data and integration layers rather than on individual service programmes (Congressional Research Service, 2022; U.S. Department of Defense, 2022b). In scholarly work, Marsili analyses how cyber, information and cognitive operations differ, overlap and intersect, and argues that loosely used "warfare" labels obscure the legal and operational distinctions between them. That analysis motivates treating cyber and electromagnetic effects

as distinct but interacting threats to the same infrastructure (Marsili, 2023). His study of influence strategy in contested neighbourhoods shows how hybrid campaigns combine informational and technical means against a target's information space (Marsili, 2021). Kott et al. describe the battlefield as a dense network of heterogeneous, intermittently connected devices whose data must be fused under adversarial conditions (Kott et al., 2016). These works set the threat context for this paper: a forward facility must assume that external links are degraded or denied and that its own inputs may be manipulated. The architecture responds with D-DIL-tolerant local execution (DO4) and Zero Trust governance (DO5).

2.2 Tactical edge computing and digital twins

Edge computing places compute and storage close to data sources to reduce latency and dependence on backhaul (Satyanarayanan, 2017; Shi et al., 2016). Tactical cloudlets apply the idea to military settings by provisioning forward, discoverable compute that keeps operating when reach-back to enterprise clouds is lost (Lewis et al., 2014). Digital twins combine live telemetry with models of the physical asset to support prediction and what-if analysis (Fuller et al., 2020). The proposed Layer 2 adapts tactical-edge execution to a hardened facility, and Layer 4 consumes digital-twin predictive alerts as one of three fused data streams.

2.3 Dependability, resilience and fault injection

Avizienis et al. (2004) provide the dependability terminology used in this study, including availability, reliability, fault, error, failure, and fault tolerance. Kott and Linkov distinguish resilience, the capacity to absorb, recover from and adapt to disruption, from robustness (Kott & Linkov, 2019). NIST SP 800-160 Vol. 2 operationalises cyber resiliency through goals of anticipating, withstanding, recovering and adapting (Ross et al., 2021). Reliability block diagrams and beta-factor common-cause modelling are standardised in IEC 61078 and IEC 61508-6 (International Electrotechnical Commission, 2010; International Electrotechnical Commission, n.d.). Fault injection is an established way to evaluate fault-tolerance mechanisms before or alongside hardware testing (Hsueh et al., 1997). This study applies these methods at model level.

2.4 Zero Trust

NIST SP 800-207 defines Zero Trust in terms of policy decision points (PDP), policy enforcement points (PEP) and per-request, identity-based access decisions (Rose et al., 2020). SP 800-207A extends the model to cloud-native, multi-location applications with identity-tier and network-tier policies (Chandramouli & Butcher, 2023). The DoD Zero Trust Reference Architecture v2.0, the DoD Zero Trust Strategy and DTM 25-003 set the defence implementation baseline (Defense Information Systems Agency & National Security Agency, 2022; U.S. Department of Defense, 2022a; U.S. Department of Defense, 2025). Section 4.3 maps these concepts to concrete enforcement points.

2.5 Design science evaluation and positioning

The study follows the DSR guidelines proposed by Hevner et al. (2004) and the process model developed by Peffers et al. (2007), with the presentation conventions of Gregor and Hevner (2013). Larsen et al. caution against conflating an artefact's internal consistency with its empirical validity (Larsen et al., 2025). This paper follows that caution: every quantitative

result is labelled as model-derived. The author's companion working manuscripts informed the design but are not used as evidence for any claim (Yasin, 2026). Table 2 positions the contribution.

Table 2. Positioning against related work strands.

Strand	Representative work	What this paper adds
Cyber–EW / multi-domain threat context	(Kott et al., 2016; Marsili, 2021; Marsili, 2023)	Translates the threat context into facility-level requirements (D-DIL survivability, EMS-denial scenarios)
Tactical edge / cloudlets	(Lewis et al., 2014; Satyanarayanan, 2017; Shi et al., 2016)	Couples edge execution with facility power, thermal and redundancy engineering and quantifies the latency benefit
Dependability and resilience	(Avizienis et al., 2004; Kott & Linkov, 2019; Ross et al., 2021)	Scenario-based availability and MTTR model with explicit failure classes and uncertainty
Zero Trust	(Chandramouli & Butcher, 2023; Defense Information Systems Agency & National Security Agency, 2022; Rose et al., 2020)	Places PEP/PDP within the facility's availability and latency budgets
DSR evaluation	(Hevner et al., 2004; Larsen et al., 2025; Peffers et al., 2007; Venable et al., 2016)	Simulation-based artificial evaluation with stated propositions and released code

Source: Author's own elaboration based on the works cited in the table.

3. Research Design

The study follows the six DSR activities of Peffers et al.: problem identification, definition of objectives, design and development, demonstration, evaluation and communication (Peffers et al., 2007). Following the FEDS framework, the evaluation strategy is "technical risk and efficacy": an artificial, ex ante evaluation by simulation that precedes, and does not replace, a naturalistic evaluation (Venable et al., 2016). Three evaluation instruments are used:

- **E1 Latency analysis.** The decomposition model of Section 5.1, propagated by Monte Carlo sampling (20,000 draws) and analysed with variance-based Sobol' indices (Saltelli et al., 2008; Sobol', 2001).
- **E2 Availability and recovery simulation.** An alternating-renewal Monte Carlo model of the reliability block structure (Section 5.2). It is run as (a) two fixed parameter cases (planning and requirement envelope), each over 3000 simulated years per configuration and scenario, and (b) 300 Latin-hypercube parameter sets (McKay et al., 1979) $\times$ 40 simulated years each, for uncertainty and Spearman rank sensitivity.
- **E3 Model-based fault injection.** Eleven fault classes injected 20,000 times per configuration, recording the probability of mission interruption and the time to restore (Hsueh et al., 1997).

The unit of evaluation is the architecture, compared with two comparators defined in Section 4.6. Evidence from E1–E3 can support or refute the propositions of Table 1 *conditional*

on the model. Structured expert evaluation and a lab-scale prototype are planned for the next DSR iteration (Section 10). No expert-panel or hardware result is reported here.

4. Architecture Specification

4.1 System boundary and layers

The system boundary is the forward EWNOC-integrated facility: its power and thermal plant, internal transport, compute, storage, integration, fusion and Zero Trust layers, and the human-authorisation gate. Sensors, reach-back and federated C2 networks and operators are external actors (Fig. 1). Layer 1 provides physically diverse outside-plant and inside-plant transport designed to BICSI TDM-15 and OSPDRM (BICSI, 2024; BICSI, n.d.), together with 2N power, N+1 cooling and an NVMe priority storage tier. Layer 2 hosts tactical-edge inference and storage so that mission functions continue when reach-back is lost. Layer 3 is a brokerless, vendor-neutral publish/subscribe integration bus without a single message broker. Layer 4 is the EWNOC fusion engine, which correlates facility telemetry, subsystem data and digital-twin predictive alerts into one operational picture. Layer 5 contains the Zero Trust PEP and PDP. A human-authorisation gate sits below the fusion layer, and a governed-automation envelope limits autonomous action to low-risk diagnostics and pre-approved remediation.




Fig. 1. Five-layer EWNOC-integrated architecture, its system boundary and external actors. Dashed arrows denote intermittent, best-effort reach-back.

Source: Author's own elaboration.

4.2 Data flows and interfaces

Table 3 defines the nine principal data flows (F1–F9) shown in Fig. 2. For each flow it gives the source, the destination, the content, the interface class and the trust boundary crossed.

Interface classes are stated generically so the design remains vendor-neutral, consistent with modular open-systems practice.




Fig. 2. Data flows F1–F9 and trust zones TZ-0 to TZ-3. Every zone crossing passes a policy enforcement point and is logged.

Source: Author’s own elaboration.

Table 3. Data flows, interfaces and trust-boundary crossings.

Flow	Source → destination	Content	Interface class	Boundary / enforcement
F1	Sensors/subsystems → PEP	Telemetry, detections, status	Authenticated publish (mTLS), schema-validated	TZ-0 → TZ-1; device identity, schema check, rate limit
F2	PEP → Layer 3 bus	Authorised, normalised events	Brokerless pub/sub, topic ACLs	TZ-1 → TZ-2; per-topic authorisation
F3	Reach-back/CJADC2 mesh ↔ PEP	Federated data, models, tasking	Store-and-forward, intermittent	TZ-0 ↔ TZ-1; mutual authentication, data tagging, cached policy
F4	PDP → PEP	Access decisions	Policy decision API	Inside TZ-1; signed decisions with bounded validity
F5	BMS/DCIM → Layer 4	Power, thermal, environmental telemetry	One-way OT gateway (read-only)	TZ-3 → TZ-2; unidirectional, no write-back
F6	Digital twin → Layer 4	Predictive alerts, health indices	Pub/sub on dedicated topic	TZ-3 → TZ-2; signed alerts
F7	Layer 4 → authorisation gate	Fused picture, proposed actions	Internal API	Inside TZ-2; action classification (routine / consequential)
F8	Gate → operator console	Approved actions and rationale	Signed messages	TZ-2 → TZ-0 via PEP; operator identity and role
F9	All layers → audit log	Decisions, access events, alerts	Append-only log	To TZ-3; integrity-protected, retained

Source: Author’s own elaboration.

4.3 Trust boundaries and Zero Trust enforcement

Four trust zones are defined. TZ-0 covers everything outside the facility. TZ-1 is the enforcement tier. TZ-2 is the operations enclave (Layers 2–4 and the gate). TZ-3 is facility operational technology (OT) and management. Enforcement follows SP 800-207 and SP 800-207A (Chandramouli & Butcher, 2023; Rose et al., 2020) and targets the DoD target-level outcome set (Defense Information Systems Agency & National Security Agency, 2022; U.S. Department of Defense, 2025). Table 4 lists the mechanisms. Two design decisions bear directly on availability and latency. First, the PEP and PDP are each deployed as redundant active-active pairs, so they appear as 1-of-2 groups in the availability model rather than as single points of failure. Second, under D-DIL the PDP evaluates against locally cached, signed policy with a bounded validity period. It fails closed for consequential actions and degraded-open only for read-only telemetry, so that loss of an enterprise identity provider does not halt local operations.

Table 4. Zero Trust enforcement mechanisms by pillar.

Pillar	Mechanism in the architecture	Placement
User	Role- and attribute-based authorisation; step-up authentication for consequential actions	PDP; authorisation gate
Device	Device identity certificates; posture attestation before admission	PEP (F1, F3)
Network / environment	Micro-segmentation by trust zone; unidirectional OT gateway	TZ boundaries; F5
Application / workload	Workload identity; mutual TLS between services; per-topic ACLs on the bus	Layer 3; F2
Data	Data tagging and labelling; schema validation; signed alerts	F1, F3, F6
Visibility and analytics	Append-only audit of every decision; anomaly detection in Layer 4	F9; Layer 4
Automation and orchestration	Policy-as-code; cached signed policy with bounded validity under D-DIL	PDP

Source: Author’s own elaboration based on the Zero Trust pillars defined by Defense Information Systems Agency and National Security Agency (2022) and U.S. Department of Defense (2022a).

4.4 Governed autonomy

System behaviour is partitioned into observation, prediction, authorisation and execution. Consequential operational actions always pass the human-authorisation gate. Low-risk, deterministic diagnostics and pre-approved failover procedures run automatically inside an auditable policy envelope. Automation therefore affects only facility recovery (the failover coverage and failover time parameters of Section 5.2), never consequential decisions.

4.5 Energy resilience and thermal management at Layers 1 and 2

Containerised forward facilities face tight power and cooling trade-offs. The design uses 2N power paths, each consisting of a generator set, an uninterruptible power supply (UPS) and distribution, with physically separated A/B routes to keep the common-cause fraction low. Cooling is N+1 (2-of-3 modules). Under fuel or generator constraints, IT load is shed in priority tiers. Accelerators are power-capped first, then non-mission analytics, while fusion, Zero Trust and the integration bus are protected. This approach is consistent with the concept of energy-proportional computing advocated by Barroso and Hölzle (2007). Equipment is assumed qualified to MIL-STD-810H for environmental exposure (U.S. Department of Defense, 2019),

and IT inlet conditions follow the ASHRAE TC 9.9 envelopes (ASHRAE Technical Committee 9.9, 2021).

A cooling outage becomes a mission outage only after the facility's thermal ride-through time is exhausted. With a lumped thermal capacitance C_{th} , an allowable temperature rise ΔT and IT heat load P_{IT} :

$$t_{ride} = C_{th} \cdot \Delta T / P_{IT} \quad (1)$$

Illustratively, assume $C_{th} \approx 1.07 \text{ MJ/K}$ (about 2,000 kg of IT and rack mass at $\approx 0.5 \text{ kJ kg}^{-1} \text{ K}^{-1}$, plus enclosure air), $\Delta T = 10 \text{ K}$ and $P_{IT} = 40 \text{ kW}$. Then $t_{ride} \approx 268 \text{ s}$ ($\approx 4.5 \text{ min}$), and at 20 kW it is $\approx 9 \text{ min}$. These are assumed values, so the ride-through parameter is varied between 3 and 20 min in the evaluation and halved in the environmental-stress scenario. Low thermal mass is a real weakness of containerised designs, and Section 7.3 quantifies its effect.

4.6 Comparator architectures

Two comparators are defined so that the claimed advantages can be tested (Table 5). The centralized baseline is a thin forward node that depends entirely on a rear data centre over two diverse wide-area paths. The rear facility is represented by a single equivalent element whose planning availability ($\approx 99.982\%$) corresponds to the commonly cited Tier III class (Uptime Institute, n.d.). The edge + diverse fabric comparator has local compute and three external paths, but its local layers are single-string and it can bridge reach-back loss only for a cache hold-over period h_c . The local components of all configurations share the same reliability parameters, and the baseline omits cooling, bus, fusion and storage groups. Both choices are conservative: they favour the comparators.

Table 5. Configurations evaluated (redundancy as n units, k required).

Group	Centralized baseline	Edge + diverse fabric	Proposed five-layer EWNOC
Power path	1 (single)	2N (1-of-2)	2N (1-of-2)
Cooling module	— (thin node)	1 (single)	N+1 (2-of-3)
Compute node	1 (thin client)	1 (single edge node)	3-of-4 cluster
Storage	—	1	1-of-2 replicated
Internal fabric path	1	1-of-2	1-of-2
Integration bus	—	1	1-of-2 brokerless
Fusion engine	—	1	1-of-2 active/standby
Zero Trust PEP/PDP	1	1	1-of-2 active-active
External (reach-back) paths	1-of-2, required	1-of-3, required after hold-over h_c	1-of-3, not required for mission capability
Rear data centre	required ($\approx$ Tier III class)	—	—

Source: Author's own elaboration.

4.7 Standards traceability

Table 6. Layer-to-metric-to-standard traceability.

Layer	Monitored metric	Primary failure mode	Engineered control	Standard basis
1	Path availability; jitter; power and thermal margin	Single-path collapse; common-cause power loss; thermal runaway	Diverse OSP/ISP routes; 2N power with route separation; N+1 cooling	(ASHRAE Technical Committee 9.9, 2021; BICSI, 2024; BICSI, n.d.; Telecommunications Industry Association, n.d.)
2	Edge inference throughput; reach-back dependency	D-DIL isolation; node loss	3-of-4 cluster; local models and data	(Lewis et al., 2014; Satyanarayanan, 2017)
3	Integration overhead; message loss	Broker single point of failure; schema heterogeneity	Brokerless redundant bus; schema validation	Open-systems practice
4	Data-convergence ratio ρ_{conv} ; alert precision/recall	Dashboard fragmentation; fusion-instance failure	Active/standby fusion; digital-twin inputs	(Fuller et al., 2020)
5	Policy coverage; decision latency	Identity or endpoint compromise; PDP loss	Active-active PEP/PDP; cached signed policy	(Chandramouli & Butcher, 2023; Defense Information Systems Agency & National Security Agency, 2022; Rose et al., 2020)
Facility	EMS and environmental exposure	Interference; HEMP; climatic extremes	Shielding and filtering; qualification	(U.S. Department of Defense, 2015; U.S. Department of Defense, 2019; U.S. Department of Defense, n.d.)

Source: Author's own elaboration based on the standards and works cited in the table.

5. Formal Models and Failure Model

5.1 Decision-cycle latency

End-to-end machine-side latency is the sum of five contributors:

$$L_{total} = L_{phys} + L_{store} + L_{int} + L_{reach} + L_{op} \quad (2)$$

Here L_{phys} is Layer 1 transport delay, L_{store} storage access and commit delay, L_{int} integration-bus normalisation and routing delay, L_{reach} expected reach-back delay, and L_{op} decision-support and authorisation-gate processing delay. In the original manuscript L_{op} was described as "operator cognitive" delay. That was inconsistent with its millisecond scale. It is redefined here as machine-side processing up to presentation at the console and arrival at the authorisation gate. Human deliberation time lies outside the model boundary. For normalised workload $w \in [0,1]$, each contributor is modelled as:

$$L_i(w) = a_i + b_i \cdot w^{p_i}, \quad i \in \{phys, store, int, reach, op\} \quad (3)$$

Here a_i (ms) is idle-load delay, b_i (ms) is congestion growth at peak and $p_i \geq 1$ is a convexity exponent. The reach-back term is derived from physical quantities rather than asserted:

$$a_{reach} = f_{nl} \cdot RTT, \quad RTT = 2D / v_f, \quad b_{reach} = \kappa \cdot f_{nl} \cdot RTT \quad (4)$$

Here f_{nl} is the fraction of decision transactions that need non-local data or compute, D the fibre route length to the reach-back node, $v_f \approx 2 \times 10^8$ m/s the propagation speed in silica fibre, and κ a queueing-growth coefficient (assumed $\kappa = 5$). For the proposed design $f_{nl} = 0.10$ and $D = 300$ km give $RTT = 3.0$ ms and $a_{reach} = 0.30$ ms. The edge comparator uses $f_{nl} = 0.30$. The baseline uses $f_{nl} = 1.0$, $D = 1,000$ km ($RTT = 10.0$ ms) and 2 ms of WAN equipment delay.

EMS interference raises the packet-error rate (PER) on radio-dependent segments. With automatic repeat request, the expected number of transmissions is $1/(1 - \text{PER})$, so:

$$L_{total}^{EMS} = L_{total} + (L_{phys} + L_{reach}) \cdot [1/(1 - \text{PER}) - 1] \quad (5)$$

Link margin under interference is governed by the signal-to-interference-plus-noise ratio. A jammer term J (dBm) is added to the noise floor N_0 , and $P_r = P_t + G_t + G_r - L_p - L_a$ is the received power:

$$\text{SINR}_{dB} = P_r - 10 \cdot \log_{10}(10^{N_0/10} + 10^{J/10}) \quad (6)$$

In the fusion engine, sub-threshold SINR flags the affected stream as degraded-confidence and down-weights it rather than discarding it, so that partial situational awareness is preserved under partial denial.

5.2 Availability and recovery

Each functional group of Table 5 is a k-of-n redundant group of identical units. Each unit alternates between up times, which are exponential with mean MTBF, and repair times, which are log-normal with mean MTTR and $\sigma = 0.6$. Steady-state unit availability and constant-failure-rate reliability are:

$$A_u = \text{MTBF} / (\text{MTBF} + \text{MTTR}), \quad R_u(t) = e^{-\lambda t}, \quad \lambda = 1/\text{MTBF} \quad (7)$$

For independent units the group availability follows the standard k-of-n expression (International Electrotechnical Commission, n.d.):

$$A_{k/n} = \sum_{j=k}^n C(n, j) \cdot A_u^j \cdot (1 - A_u)^{n-j} \quad (8)$$

Three mechanisms violate the assumption of independent component failures and are therefore explicitly represented in the simulation. (i) Common-cause failures occur at rate $\beta \cdot \lambda_u$ per group and take down all units (beta-factor method; International Electrotechnical Commission, 2010). (ii) Imperfect failover: each unit failure in a redundant group is covered with probability c , causing a transient interruption of duration τ , and is otherwise uncovered, causing an outage until repair. To first order the per-group contribution to unavailability is:

$$U_{fo} \approx n \cdot \lambda_u \cdot [c \cdot \tau + (1 - c) \cdot \text{MTTR}] \quad (9)$$

(iii) **Ride-through and hold-over:** cooling-group outages shorter than t_{ride} (Eq. 1) are absorbed, and only their excess counts. For the edge comparator, reach-back outages shorter than h_c are bridged. Mission capability is the logical AND of all required groups. For the proposed configuration reach-back is *not* required (DO4). Mission availability A is the fraction of simulated time that the system is mission-capable. System-level MTTR is total mission downtime divided by the number of distinct mission-outage events, where overlapping group outages are merged into one event.

5.3 Failure model and operating scenarios

Table 7 lists the fault classes and how each is represented. Table 8 defines the four scenarios used to answer RQ3 and to address the effect of EMS and environmental degradation on the pre-prototype metrics. The facility is assumed to be shielded and filtered to MIL-STD-461G and MIL-STD-188-125 (U.S. Department of Defense, 2015; U.S. Department of Defense, n.d.). EMS effects are therefore modelled on external radio and satellite paths, not on internal equipment. Failure of that hardening is outside the model and is noted as a limitation.

Table 7. Failure model.

Fault class	Layer	Representation in the model
Random hardware/software unit failure	1–5	Exponential up time; log-normal repair
Uncovered failover	1–5	Bernoulli($1 - c$) per unit failure $\rightarrow$ outage until repair
Failover transient	1–5	Covered failures interrupt service for τ
Common-cause failure	Redundant groups	Poisson events at $\beta \cdot \lambda_u$; all units down
Cooling loss	1	Excess over thermal ride-through t_{ride} counts as outage
EMS denial of external links	External	Poisson episodes; each path jammed with probability q per episode
Physical attrition	1, 2, external	Poisson hits on a random unit (power, fabric, compute or external path); mean repair 48 h
Environmental stress	1	Power and cooling failure rates $\times 2$, repair $\times 1.5$, ride-through $\times 0.5$

Source: Author's own elaboration.

Table 8. Operating scenarios.

Scenario	EMS episodes per year (mean duration; q)	Environmental stress factor	Physical hits per year
S0 Benign	0	none	0
S1 EMS-contested	24 (3 h; 0.8)	none	0
S2 Environmental stress	0	$\times 2$ failure rate, $\times 1.5$ repair, $\times 0.5$ ride-through	0
S3 Combined contested	24 (3 h; 0.8)	$\times 2$ failure rate, $\times 1.5$ repair, $\times 0.5$ ride-through	6

Source: Author's own elaboration.

6. Parameterisation, Simulation Design and Reproducibility

No field or laboratory data on military C5I facilities are available to the author for publication. Every parameter is therefore an **assumed engineering planning value**. Tables 9 and 10 state its unit, its uncertainty range and its basis. Where a value is derived from physics (propagation delay) or tied to a recognised class (rear facility $\approx$ Tier III), this is stated. The latency allocation is a *design budget* calibrated so that the nominal total equals the original 4.60 ms target at $w = 1$. The Monte Carlo analysis then tests how robust that budget is. The availability parameters were fixed before the simulations were run and were not tuned to the targets.

Table 9. Latency-model parameters (proposed configuration, nominal values; ms unless stated).

Contributor	a_i (ms)	b_i (ms)	p_i	$L_i(1)$ (ms)	Uncertainty	Basis
L_{phys}	0.08	0.22	1.0	0.30	a, b: triangular $\pm 30\%$; p: uniform ± 0.25 ($p \geq 1$)	Design budget allocation (assumed)
L_{store}	0.12	0.34	1.5	0.46	a, b: triangular $\pm 30\%$; p: uniform ± 0.25 ($p \geq 1$)	Design budget allocation (assumed)
L_{int}	0.18	0.52	1.5	0.70	a, b: triangular $\pm 30\%$; p: uniform ± 0.25 ($p \geq 1$)	Design budget allocation (assumed)
L_{reach}	0.30	1.50	2.0	1.80	a, b: triangular $\pm 30\%$; p: uniform ± 0.25 ($p \geq 1$)	Derived (Eq. 4): $f_{n1} = 0.10$, $D = 300$ km, $\kappa = 5$
L_{op}	0.28	1.06	2.0	1.34	a, b: triangular $\pm 30\%$; p: uniform ± 0.25 ($p \geq 1$)	Design budget allocation (assumed)
Total	0.96	3.64	—	4.60	Propagated	Sum (Eq. 2)

Source: Author's own elaboration.

Table 10. Reliability parameters (per unit, nominal planning values) and uncertainty ranges.

Group	MTBF (h)	MTTR (h)	Uncertainty range	Basis
Power path (genset + UPS + distribution)	2,000	4	MTBF $\times [0.5, 2]$; MTTR $\times [0.5, 2]$ (log-uniform)	Assumed planning value for field-deployed generation
Cooling module	4,000	8	as above	Assumed
Compute node	6,000	6	as above	Assumed (ruggedised, accelerator-equipped)
Storage replica	8,000	4	as above	Assumed
Internal fabric path	4,000	3	as above	Assumed
Integration-bus instance	6,000	2	as above	Assumed (software restart dominated)
Fusion-engine instance	3,000	1	as above	Assumed (software-intensive)
Zero Trust PEP/PDP	5,000	2	as above	Assumed
External (reach-back) path	700	10	as above	Assumed (exposed tactical link)
Rear data centre (equivalent)	22,000	4	MTTR $\times [0.5, 2]$	Chosen to give $\approx 99.982\%$ (Tier III class) (Uptime Institute, n.d.)
Failover coverage c	—	—	$U[0.95, 0.999]$; planning 0.975	Assumed
Common-cause fraction β	—	—	$U[0.01, 0.10]$; planning 0.05	Typical beta-factor range (International Electrotechnical Commission, 2010)
Failover transient τ	—	—	$U[0.25, 2]$ min; planning 1 min	Assumed
Thermal ride-through t_{ride}	—	—	$U[3, 20]$ min; planning 10 min	Eq. 1 illustration
Edge hold-over h_c	—	—	$U[0.25, 2]$ h; planning 1 h	Assumed

Source: Author's own elaboration.

Two fixed parameter cases are reported. The **planning case** uses the nominal values (all multipliers 1, $c = 0.975$, $\beta = 0.05$, $\tau = 1$ min). The **requirement-envelope case** states what the design would need to achieve: unit MTBF $\times 2$ (through qualification and preventive maintenance), MTTR $\times 0.5$ (through forward spares and hot-swap), $c = 0.995$, $\beta = 0.02$ (through

physical route separation) and $\tau = 0.25$ min. Each case is simulated for 3,000 years per configuration and scenario, in 30 independent batches of 100 years, and 95% intervals are taken across batches. The uncertainty study draws 300 Latin-hypercube samples across all ranges of Table 10 and simulates 40 years per sample. Sensitivity is measured with Spearman rank correlation for availability and with first-order and total Sobol' indices (Saltelli/Jansen estimators, $N = 20,000$) for latency [27, 28]. The code (Python 3, NumPy, SciPy; fixed seed 20261012) regenerates every number, table and figure in Sections 6–7 and is provided as Supplementary Material S1.

7. Results

7.1 Latency (RQ1)

Fig. 3a shows the contributor curves for the proposed design with the 5–95% uncertainty band. Fig. 3b compares configurations. The nominal total rises from 0.96 ms at idle to 4.60 ms at peak. Reach-back and decision-support processing together account for 68% of the peak total, which supports the original paper's qualitative claim. The edge and centralized configurations reach 8.20 ms and 64.8 ms at peak, so the proposed design's peak latency is 44% lower than the edge comparator and 93% lower than the centralized baseline. The probability that the proposed total stays within 5 ms at peak is 0.946.




Fig. 3. (a) Modelled latency contributors $L_i(w)$ and total for the proposed architecture with the 5–95% Monte Carlo interval; (b) total modelled latency of the three configurations (log scale).

Source: Author's own elaboration based on simulation results.

Table 11. Modelled decision-cycle latency (ms): Monte Carlo median and 5–95% interval.

Configuration	w = 0	w = 0.5	w = 1.0	Reach-back RTT
Centralized baseline	12.65 (10.19–15.12)	25.85 (21.96–30.05)	64.78 (54.27–75.33)	10.0
Edge + diverse fabric	1.56 (1.36–1.76)	3.36 (3.01–3.74)	8.21 (7.24–9.16)	3.0
Proposed five-layer EWNOC	0.96 (0.87–1.05)	2.01 (1.85–2.18)	4.60 (4.20–5.01)	3.0

Source: Author's own elaboration based on simulation results.

Sobol' analysis (Fig. 4a) shows that peak-load uncertainty is driven by the congestion coefficients of reach-back ($S_T = 0.56$) and decision-support processing ($S_T = 0.28$). Integration follows ($S_T = 0.07$). At $w = 0.7$ the reach-back exponent p_{reach} also matters ($S_T = 0.07$).

Engineering effort on reach-back avoidance and on decision-support pipeline efficiency will therefore reduce latency more than faster physical media. Under EMS degradation (Fig. 4b), a PER of 30% raises the peak total to 5.50 ms for the proposed design, against 10.64 ms (edge) and 91.5 ms (baseline). The proposed design is less exposed because a smaller share of its latency travels over radio-dependent paths.




Fig. 4. (a) Total-order Sobol' indices of L_{total} for the eight most influential parameters at $w = 1.0$ and $w = 0.7$; (b) peak latency against EMS-induced packet-error rate.

Source: Author's own elaboration based on simulation results.

7.2 Availability and recovery (RQ2, RQ3)

Table 12. Modelled mission availability (%) / system-level MTTR (min). Means over 3,000 simulated years; 95% batch intervals for availability are within ± 0.04 percentage points for the proposed configuration.

Case	Configuration	S0 Benign	S1 EMS	S2 Environmental	S3 Combined
Planning	Baseline	99.404 / 88.5	98.877 / 116.8	99.009 / 131.9	96.119 / 354.9
Planning	Edge	99.355 / 132.3	99.082 / 132.2	98.911 / 155.1	97.818 / 237.0
Planning	Proposed	99.907 / 14.2	99.910 / 13.6	99.808 / 23.3	99.795 / 24.5
Requirement	Baseline	99.866 / 38.7	99.341 / 103.6	99.766 / 59.6	96.813 / 426.9
Requirement	Edge	99.878 / 52.2	99.589 / 93.1	99.771 / 67.5	98.669 / 234.8
Requirement	Proposed	99.993 / 2.2	99.993 / 2.2	99.986 / 3.5	99.979 / 5.0

Source: Author's own elaboration based on simulation results.

Table 12 and Fig. 5 give the main result. Under **planning** parameters the proposed design achieves 99.907% availability (≈ 491 min of downtime per year) with an MTTR of 14.2 min. That is 6–7 times less downtime than either comparator (99.40% and 99.35%), but it does **not** meet DO2 (99.985%) or DO3 (12 min). Proposition P4 is supported. Within the **requirement envelope** the proposed design reaches 99.993% and 2.2 min in S0, 99.993% in S1 and 99.986% in S2, which meets DO2 and DO3 in S0–S2 (P3 supported). Under combined stress (S3) it reaches 99.979%, just short of DO2, while MTTR stays at 5.0 min.

Reach-back denial (S1) leaves the proposed configuration unchanged (99.907% $\rightarrow$ 99.910% in the planning case). The baseline falls from 99.40% to 98.88% and the edge comparator from 99.35% to 99.08%, both reductions greater than 0.25 points (P5 supported). The cost of this independence is that the proposed facility runs in D-DIL mode, without reach-

back, for 0.60% of the time in S1. Its local picture is complete during those periods, but federated data are stale. One result runs against intuition. In the benign planning case the edge comparator is slightly *less* available than the centralized baseline. Adding non-redundant local layers adds series failure points, so edge placement only pays off in availability terms when it is combined with redundancy, or when reach-back is contested.




Fig. 5. Modelled mission downtime per year by configuration and scenario for (a) planning parameters and (b) the requirement envelope. The dotted line marks the 99.985% target (≈ 79 min/year).

Source: Author's own elaboration based on simulation results.

7.3 Uncertainty and sensitivity

Across the 300 parameter sets (Fig. 6b), the median availability of the proposed design in S0 is 99.902% (5–95%: 99.756–99.966%). Only 0.3% of sets meet 99.985%, while 12% reach 99.95% and 40% meet the 12-min MTTR. The comparators do not reach 99.95% in any set. Fig. 6a ranks the parameters. Mission unavailability is most strongly associated with the repair-time multiplier ($\rho = 0.69$), failover coverage ($\rho = -0.42$), the common-cause fraction ($\rho = 0.34$) and power-path MTBF ($\rho = -0.22$). Under combined stress the power and cooling terms grow in importance. This ranking is the evidence base for the requirement envelope: the targets depend more on logistics (forward spares and hot-swap), automated failover quality and physical separation of redundant paths than on the intrinsic reliability of any single component.




Fig. 6. (a) Spearman rank correlation between parameters and the proposed design’s mission unavailability (S0 and S3); (b) distribution of mission unavailability across parameter sets (box: interquartile range; whiskers: 5–95%).

Source: Author’s own elaboration based on simulation results.

7.4 Model-based fault injection

Table 13. Fault-injection results: probability of mission interruption / expected mission outage per injected fault (min). "n/a" = group not present.

Injected fault	Centralized baseline	Edge + diverse fabric	Proposed
F1 single power-path failure	100.0% / 239.1	100.0% / 7.0	100.0% / 6.7
F2 single cooling-module failure	n/a	100.0% / 470.8	2.6% / 12.5
F3 single compute-node failure	100.0% / 362.3	100.0% / 359.7	100.0% / 10.2
F4 fusion-engine instance failure	n/a	100.0% / 60.5	100.0% / 2.4
F5 Zero Trust gateway (PEP/PDP) failure	100.0% / 119.4	100.0% / 119.1	100.0% / 3.9
F6 integration-bus instance failure	n/a	100.0% / 120.4	100.0% / 4.2
F7 loss of one external (reach-back) path	masked (0)	masked (0)	masked (0)
F8 double power-path failure	n/a	100.0% / 160.6	100.0% / 161.3
F9 double cooling-module failure	n/a	n/a	100.0% / 311.0
F10 common-cause power failure	n/a	100.0% / 239.3	100.0% / 240.4
F11 EMS full-denial of reach-back (mean 3 h)	100.0% / 178.8	89.2% / 122.4	masked (0)

Source: Author’s own elaboration based on simulation results.

Table 13 and Fig. 7 report the campaign at planning parameters. In the proposed design, single faults in power, compute, fusion, Zero Trust and bus groups always cause a brief failover transient (100% interruption probability, about 1 min). Expected outage per fault is 2.4–10.2 min because uncovered failovers occasionally require repair. A single cooling-module failure interrupts the mission in only 2.6% of trials, because the remaining 2-of-3 modules and the ride-through absorb it. In the non-redundant comparators, the same faults produce 60–470 min of expected outage. That is a ratio of at least 25 in every case, so P6 is supported. Loss of a single external path is masked in all configurations. A full EMS denial of reach-back (F11) causes no mission outage in the proposed design, against a mean of about 3 h for the baseline. Faults that exhaust redundancy (F8, F9) or strike common causes (F10) cost the proposed design as much

as the comparators (160–310 min). This is why β and route separation appear in the requirement envelope.




Fig. 7. Expected mission outage per injected fault class (planning parameters, log scale).

Source: Author’s own elaboration based on simulation results.

7.5 Assessment of propositions

Table 14. Verdict on design propositions (conditional on the model).

Proposition	Evidence	Verdict
P1 latency ≤ 5 ms with $p \geq 0.90$	$P = 0.946$	Supported
P2 $\geq 40\%$ / $\geq 90\%$ lower latency	44% vs edge; 93% vs baseline	Supported
P3 targets met in S0–S2 (requirement envelope)	99.993 / 99.993 / 99.986%; MTTR ≤ 3.5 min	Supported (S3: availability narrowly missed)
P4 targets not met at planning values	99.907%, 14.2 min	Supported
P5 D-DIL independence	Δ proposed 0.004 pp; baseline -0.53 pp; edge -0.27 pp	Supported
P6 $\geq 10\times$ lower single-fault outage	Ratios ≥ 25 across F1–F6	Supported

Source: Author’s own elaboration based on simulation results.

8. Discussion

The central thesis of the original study—that decision-cycle latency is determined by multiple system-level factors rather than bandwidth alone—is supported by the explicit modelling results. Reach-back and decision-support processing dominate both the nominal budget and its uncertainty. Local execution addresses the first directly, and fusion addresses the second by removing manual cross-dashboard correlation. The latency advantage over the centralized comparator is mainly a matter of physics, since a 1,000 km fibre path costs 10 ms round trip before any queueing. The advantage over the edge comparator comes from the lower non-local fraction f_{nl} .

The availability results change the paper’s claims materially. The original 99.985% availability and 12-min MTTR are not intrinsic properties of the five-layer topology. They are achievable only if failover coverage approaches 0.995, the common-cause fraction is held near 0.02 through physical separation, and repair is roughly halved through forward spares and hot-

swap. The design contribution is therefore the topology together with this requirement envelope, which can be written into procurement and acceptance-test specifications. For context, the planning-case result ($\approx 99.91\%$) lies between the commonly cited Tier II and Tier III availability classes, and the requirement-envelope result ($\approx 99.99\%$) lies between Tier III and Tier IV (Uptime Institute, n.d.). These are comparisons of order of magnitude only, because tier classes describe fixed civilian facilities with grid supply and are not like-for-like comparators.

On the reviewers' question of EMS and environmental degradation, EMS denial matters for the proposed design mainly through the share of time spent in D-DIL mode and through EMS-related latency penalties, not through loss of mission capability. Environmental stress, by contrast, acts directly on the power and cooling plant. It roughly doubles downtime in the planning case, and in the requirement envelope it moves the design from 99.993% to about 99.986%. Thermal ride-through is a containerised facility's weak point. Raising thermal mass, using chilled-water storage or derating IT load during cooling faults are therefore effective design levers, and the model can test each one.

Finally, the architecture places Zero Trust enforcement inside the same availability and latency budgets as the mission functions. A single PDP would have been a series element of planning availability $\approx 99.96\%$, enough on its own to break DO2. Redundant PEP/PDP pairs and cached signed policy are therefore availability requirements, not only security features.

9. Threats to Validity and Limitations

- **Construct validity.** Mission capability is binary, and degraded modes such as stale federated data are reported but not scored. Human decision time is excluded from latency by definition.
- **Internal validity.** Parameters are assumed planning values, not field data, and the latency allocation is a design budget. Results are conditional on them. The uncertainty and sensitivity analyses bound, but do not remove, this dependence. Units within a group are assumed identical and failures independent apart from the modelled common-cause and scenario effects. Repair crews are not constrained.
- **External validity.** Comparators are stylised, and real centralized or edge deployments may have different redundancy. The EMS model affects only external paths and assumes hardening is effective. Effects of directed energy or HEMP on unhardened equipment are excluded.
- **Evaluation scope.** The evaluation is artificial and *ex ante*. It establishes internal consistency and conditional performance, not empirical validity (Larsen et al., 2025). No expert panel or prototype data are reported. Thresholds for P2 and P6 were set by the author in this revision and are stated so that others can re-test them.

Programmatic context is cited only from primary government sources. Market forecasts, vendor briefs and news reports used in the first version have been removed.

10. Conclusion and Future Work

This paper has presented a design-stage, DSR-framed architecture for an EWNOC-integrated forward data centre. It specifies system boundaries, data flows, trust boundaries and Zero Trust enforcement, and evaluates the design with reproducible latency, availability and fault-injection models against two comparators. The proposed design shows a large modelled latency advantage (4.60 ms versus 8.20 ms and 64.8 ms at peak), substantially lower downtime in every scenario, and insensitivity of mission capability to reach-back denial. Its availability and recovery targets are achievable only within a stated requirement envelope. The next DSR iteration will (i) conduct a structured expert evaluation of the architecture and the requirement envelope with facility, network-operations and EW practitioners; (ii) build a lab-scale prototype and run physical fault-injection and D-DIL emulation campaigns that report p50/p95/p99 latency, availability, MTTR, predictive-alert precision and recall, and the data-convergence ratio $\rho_{\text{conv}} = N_{\text{corr}}/N_{\text{total}}$; and (iii) replace the planning parameters with measured failure and repair distributions.

References

- Avizienis, A., Laprie, J. C., Randell, B., & Landwehr, C. (2004). Basic concepts and taxonomy of dependable and secure computing. *IEEE Transactions on Dependable and Secure Computing*, 1(1), 11–33.
- Barroso, L. A., & Hölzle, U. (2007). The case for energy-proportional computing. *Computer*, 40(12), 33–37.
- Fuller, A., Fan, Z., Day, C., & Barlow, C. (2020). Digital twin: Enabling technologies, challenges and open research. *IEEE Access*, 8, 108952–108971.
- Gregor, S., & Hevner, A. R. (2013). Positioning and presenting design science research for maximum impact. *MIS Quarterly*, 37(2), 337–355.
- Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75–105.
- Hsueh, M. C., Tsai, T. K., & Iyer, R. K. (1997). Fault injection techniques and tools. *Computer*, 30(4), 75–82.
- Kott, A., & Linkov, I. (2019). Fundamental concepts of cyber resilience: Introduction and overview. In A. Kott & I. Linkov (Eds.), *Cyber resilience of systems and networks*. Springer.
- Kott, A., Swami, A., & West, B. J. (2016). The Internet of Battle Things. *Computer*, 49(12), 70–75.
- Larsen, K. R., Lukyanenko, R., Mueller, R. M., Storey, V. C., Parsons, J., VanderMeer, D., & Hovorka, D. S. (2025). Validity in design science. *MIS Quarterly*, 49(4), 1267–1294. <https://doi.org/10.25300/MISQ/2024/18064>
- Lewis, G., Echeverría, S., Simanta, S., Bradshaw, B., & Root, J. (2014). Tactical cloudlets: Moving cloud computing to the edge. In *Proceedings of the IEEE Military Communications Conference (MILCOM 2014)* (pp. 1440–1446). IEEE. <https://doi.org/10.1109/MILCOM.2014.238>.
- Marsili, M. (2021). The Russian influence strategy in its contested neighbourhood. In H. Mölder, V. Sazonov, A. Chochia, & T. Kerikmäe (Eds.), *The Russian Federation in global knowledge warfare: Influence operations in Europe and its neighbourhood*. Springer. https://doi.org/10.1007/978-3-030-73955-3_8.

- Marsili, M. (2023). Guerre à la carte: Cyber, information, cognitive warfare and the metaverse. *Applied Cybersecurity & Internet Governance*, 2(1), 105–120. <https://doi.org/10.60097/ACIG/162861>
- McKay, M. D., Beckman, R. J., & Conover, W. J. (1979). A comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2), 239–245.
- Peffer, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45–77.
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., & Tarantola, S. (2008). *Global sensitivity analysis: The primer*. Wiley.
- Satyanarayanan, M. (2017). The emergence of edge computing. *Computer*, 50(1), 30–39.
- Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
- Sobol', I. M. (2001). Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates. *Mathematics and Computers in Simulation*, 55(1–3), 271–280.
- Venable, J., Pries-Heje, J., & Baskerville, R. (2016). FEDS: A framework for evaluation in design science research. *European Journal of Information Systems*, 25(1), 77–89.

Other Sources

- ASHRAE Technical Committee 9.9. (2021). *Thermal guidelines for data processing environments* (5th ed.). ASHRAE.
- BICSI. (2024). *Telecommunications distribution methods manual (TDMM)* (15th ed.).
- BICSI. (n.d.). *Outside plant design reference manual (OSPDRM)*.
- Chandramouli, R., & Butcher, Z. (2023). *A zero trust architecture model for access control in cloud-native applications in multi-location environments* (NIST Special Publication 800-207A). National Institute of Standards and Technology.
- Congressional Research Service. (2022, January 21). *Joint All-Domain Command and Control (JADC2)* (In Focus IF11493).
- Defense Information Systems Agency, & National Security Agency. (2022, July). *Department of Defense zero trust reference architecture* (Version 2.0).
- International Electrotechnical Commission. (2010). *IEC 61508-6: Functional safety of electrical/electronic/programmable electronic safety-related systems—Part 6: Guidelines on the application of IEC 61508-2 and IEC 61508-3*.
- International Electrotechnical Commission. (n.d.). *IEC 61078: Reliability block diagrams*.
- Rose, S., Borchert, O., Mitchell, S., & Connelly, S. (2020). *Zero trust architecture* (NIST Special Publication 800-207). National Institute of Standards and Technology.
- Ross, R., Pillitteri, V., Graubart, R., Bodeau, D., & McQuaid, R. (2021). *Developing cyber-resilient systems: A systems security engineering approach* (NIST Special Publication 800-160, Vol. 2, Rev. 1). National Institute of Standards and Technology. <https://doi.org/10.6028/NIST.SP.800-160v2r1>
- Telecommunications Industry Association. (n.d.). *ANSI/TIA-942-C: Telecommunications infrastructure standard for data centers*.
- U.S. Department of Defense. (2015). *MIL-STD-461G: Requirements for the control of electromagnetic interference characteristics of subsystems and equipment*.

- U.S. Department of Defense. (2019). *MIL-STD-810H: Environmental engineering considerations and laboratory tests*.
- U.S. Department of Defense. (2022a, November). *DoD zero trust strategy*.
- U.S. Department of Defense. (2022b, March). *Summary of the Joint All-Domain Command and Control (JADC2) strategy*.
- U.S. Department of Defense. (2025, July 17). *Directive-type memorandum 25-003: Implementing the DoD zero trust strategy*.
- U.S. Department of Defense. (n.d.). *MIL-STD-188-125-1 and -2: High-altitude electromagnetic pulse (HEMP) protection for ground-based C4I facilities performing critical, time-urgent missions*.
- Uptime Institute. (n.d.). *Tier standard: Topology*.
- Yasin, M. R. (2026). *Companion design studies on tactical-edge C5I computing, integrated defence data-centre infrastructure, digital-twin predictive situational awareness, and autonomous infrastructure management* [Unpublished manuscripts].