


THE DISTRIBUTED SHIELD

Multi-Sensor Fusion, Resilient Tactical Networks and Air Defence in the Drone-Swarm Era

From point-to-point battlefield communications to thousands of cooperative defence nodes



Written by Muhammad Sohail Ashraf

An educational, well-researched open-source analysis

THE DISTRIBUTED SHIELD

*Multi-Sensor Fusion, Resilient Tactical Networks and Air Defence in the Drone-Swarm Era
From Point-to-Point Battlefield Communications to Thousands of Cooperative Defence
Nodes*

Written by

Muhammad Sohail Ashraf

Based on the author's 2023 MSc research at the University of Hertfordshire

PUBLICATION STATEMENT

An educational, well-researched analysis based on the author's research, open-source information, official publications, academic literature and claims reported worldwide. It contains no classified information. Simulations are illustrative and are not operational predictions.

Copyright © 2026 Muhammad Sohail Ashraf. All rights reserved.

Preface

This book is a major expansion of my original university research into military communication and surveillance using point-to-point wireless devices. The original work was built around a practical question: can a local network of directional links, access points, cameras, sensors and command centres provide more bandwidth, responsiveness and control than traditional communications alone?

The events and technological developments of the following years made the wider implication clearer. Modern air defence is no longer only a contest between a radar and a missile. It is a contest between networks. Small drones and loitering munitions attack the lower layer, while electronic warfare and cyber operations attack the information layer. Expensive strategic systems can remain technically powerful yet become operationally stressed if they depend on too few nodes or too few interceptors.

This study therefore develops the original network into a distributed shield: a system in which large strategic sensors and interceptors are supported by many smaller nodes, local fusion, low-cost effectors and resilient command. The aim is not to argue that cheap devices replace high-end technology. It is to show how distributed sensing and communications can protect high-end systems, preserve situational awareness and maintain defence through repeated attacks.

The book also revises parts of the original research where stronger evidence is now available. It distinguishes simulation from field validation, rejects unsupported claims of automatic quantum resistance, and places artificial intelligence under human command authority. These revisions strengthen rather than weaken the research because credible defence analysis must identify its limits.

Muhammad Sohail Ashraf

August 2026

Contents

Chapter	Title	Page
1	Research Origin and the Central Thesis	6
2	The Communications Problem: From HF and Satellite Links to Resilient Local Networks	9
3	The Original Point-to-Point Architecture	13
4	From Tactical Communications to Distributed Air Defence	17
5	Building the Lower-Layer Air Picture	20
6	Multi-Sensor Fusion and Track Confidence	23
7	AI-Assisted Decision Support and Human Authority	26
8	Network-Centric Warfare, Command and Data Links	28
9	Few High-End Nodes versus Thousands of Small Nodes	31
10	Legacy Anti-Aircraft Guns and Low-Cost Effectors	34
11	High-Tier Systems: S-400, THAAD, Aegis/Standard Missile and Iron Dome	37
12	Drone Swarms, Saturation and the Cost-Exchange Problem	40
13	Resilience against Jamming, Spoofing, Cyber Attack and Physical Attrition	43
14	Interoperability, Gateways and Edge Computing	46
15	Electronic Warfare, Emission	49

Chapter	Title	Page
	Control and Passive Sensing	
16	Human Factors, Cognitive Load and Command Design	51
17	Defending Bases, Cities and Critical Infrastructure	53
18	External Validation: Policy and Experimentation in 2025-2026	55
19	Illustrative War-Gaming Method	57
20	Simulation Results and Operational Interpretation	60
21	Deployment and Implementation Roadmap	63
22	Economics, Logistics, Training and Governance	65
23	Findings, Limitations and Future Research	67
A	Simulation Assumptions and Formulae	70
B	Reference Architecture and Message Types	71
C	Glossary	72
D	Bibliography and Image Credits	73

Research Origin and the Central Thesis

How a 2023 point-to-point communication project develops into a theory of distributed air defence

AUTHOR'S ORIGINAL CONTRIBUTION

The 2023 project joined local communications, cameras, IoT sensors, command centres and remote control through directional point-to-point links. This book extends that architecture into a lower-layer air-defence and situational-awareness network.

This book grows from Muhammad Sohail Ashraf's 2023 MSc project, Military Communication and Surveillance System Using Point-to-Point Wi-Fi Devices, completed at the University of Hertfordshire. The original project was not an air-defence weapon design. It was a communications and surveillance proof of concept: central and field command centres were linked to access points, directional antennas, cameras, Internet of Things sensors and field devices so that information could move quickly and securely across a defined operational area. Its most important intellectual contribution was the decision to treat communications, surveillance and remote control as one integrated system rather than as separate pieces of equipment.

The project began with a practical observation. Conventional military radio and satellite communications are indispensable, but each has limitations. High-frequency radio can provide long reach, yet it offers limited bandwidth and can be affected by propagation, interference and operator workload. Satellite communications can provide wide-area coverage, but they introduce dependence on scarce infrastructure, can suffer latency, and may be vulnerable to disruption or denial. The proposed local point-to-point network used directional links, unicast traffic and distributed access points to create a lower-latency local communications layer. The aim was not to abolish radio or satellite systems, but to provide a complementary architecture for the parts of the battlefield where local bandwidth, surveillance video and rapid device control matter most.

The central thesis of this expanded study is that the same architecture can be developed into a distributed lower-layer air-defence network. A point-to-point communications grid can carry not only voice and video, but also radar plots, electro-optical observations, passive radio-frequency detections, acoustic bearings, friendly-force locations, track confidence scores and defensive tasking. When hundreds or thousands of low-cost nodes contribute to a common picture, the result can be more resilient than an architecture built around only a few expensive sensors and command centres.




Figure 1.1. The original point-to-point communication concept expanded into a distributed sensing and response architecture.

This does not mean that small nodes replace strategic radars, airborne early-warning aircraft, S-400-class long-range systems, THAAD, Aegis or the Standard Missile family. High-tier systems provide reach, discrimination, altitude, interceptor performance and battle-management functions that low-cost nodes cannot reproduce. The argument is architectural: high-end systems become more survivable and more economical when they are supported by a dense local layer that detects low and slow threats, filters clutter, preserves local situational awareness and assigns the cheapest effective response.

The distinction is especially important in the drone-swarm era. A modern defensive system can be defeated without every interceptor or radar being technically inferior. It can be stressed by too many simultaneous tracks, insufficient low-altitude coverage, limited ready magazines, slow reloads, expensive interceptors, electronic attack, cyber intrusion and the destruction of a small number of critical nodes. The failure is therefore often a failure of system design, capacity and economics rather than a failure of one missile.

The research proposition is that a dense network changes this equation. Numerous local sensors provide overlapping observations. Edge fusion cells correlate them before forwarding recognised tracks. Alternate data paths reduce dependence on a single hub. Local effectors can engage suitable threats without consuming premium missiles. Damaged nodes can be isolated and replaced. The network may lose accuracy or capacity under attack, but it should degrade gradually rather than collapse suddenly.

The book uses the phrase distributed shield for this combination of high-tier and low-tier defence. It includes strategic radars, airborne sensors and premium interceptors, but it also includes micro-radars, EO/IR cameras, passive RF receivers, acoustic arrays, field observers, mobile relays, local fusion software, legacy anti-aircraft guns, short-range interceptors, electronic warfare and secure command networks. The value comes from how these elements cooperate.




Figure 1.2. An original topology image from the 2023 project, showing the research emphasis on connecting command centres, access points and field devices.

The analysis is educational and based on open-source information, the author's original project, official policy documents, public technical descriptions and academic research. It contains no classified material. The simulation results are illustrative, not predictions of actual national systems. Their purpose is to test the logic of sparse and dense architectures under declared assumptions.

A second purpose is to improve the original research by identifying claims that require refinement. Local wireless links do not automatically provide resistance to quantum computing; security depends on cryptography, authentication, key management, hardware trust and post-quantum migration. Likewise, facial emotion analysis should not be treated as a reliable method of determining hostile intent. The stronger and more defensible use of artificial intelligence lies in track association, anomaly detection, sensor management, confidence scoring and decision support under human authority.

The book's final judgement is therefore conditional. A distributed network is not inherently secure, accurate or effective merely because it has many nodes. A badly governed dense network can create false alarms, cyber exposure, maintenance burden and command confusion. The architecture becomes valuable only when it is built around data quality, interoperability, redundancy, zero-trust security, disciplined command and realistic training.

Selected sources: Ashraf (2023); NATO (2025); NATO ACT (2026a).

CHAPTER 02

The Communications Problem: From HF and Satellite Links to Resilient Local Networks

Why modern warfare requires multiple communications layers rather than one universal network

CRITICAL REASSESSMENT

Local point-to-point links reduce some exposure, but they are not automatically quantum-resistant. Security against future cryptanalytic threats depends on post-quantum algorithms, disciplined key management, authenticated devices and upgradeable cryptographic modules.

Military communication has always been shaped by a trade-off among range, bandwidth, mobility, survivability and control. No single part of the electromagnetic spectrum provides every desired property. High-frequency radio can reach beyond line of sight through ionospheric propagation, but the channel is variable, narrow and operationally demanding. Very-high-frequency and ultra-high-frequency systems provide clearer tactical communications, yet terrain, buildings and curvature limit reach. Satellite systems extend coverage, but they depend on space and ground infrastructure and can introduce latency, congestion and strategic vulnerability.

The original research correctly treated the communications problem as a system problem. A field unit needs more than voice. It may need maps, video, imagery, sensor alerts, biometric information, friendly-force tracking, remote device control and rapid access to commanders. These demands increase bandwidth and reduce tolerance for delay. A network that performs well for voice messages may not perform well for multiple high-definition video feeds or time-critical sensor cueing.

Point-to-point wireless links are attractive because directional antennas concentrate energy between known locations. This can improve link budget, reduce unnecessary radiation outside the intended path and support higher data rates over local distances. Directionality also helps frequency reuse: two links operating in the same general band may interfere less when their beams are separated. The project added dynamic antenna alignment and remote control so that field users could adjust links as positions changed.




Figure 2.1. The communications hierarchy considered in the original research.

However, line-of-sight systems introduce their own constraints. Terrain masking, foliage, buildings, weather, antenna movement and mast damage can break a link. A network based on a single chain of directional links may be efficient in normal conditions but fragile under manoeuvre or attack. The book therefore modifies the original design by adding mesh relays, alternate frequencies, wired or fibre backhaul where available, store-and-forward modes, local processing and automatic route recovery.

The most credible architecture is hybrid. High-frequency radio remains useful for long-distance fallback. Tactical radios remain essential for mobile units. Satellite communications provide strategic reach. Fibre provides high capacity and electromagnetic silence along protected routes. Point-to-point wireless links create high-bandwidth local bridges. Mesh links allow data to move around damaged nodes. The question is not which medium wins, but which combination maintains the mission under changing conditions.

Resilience also requires separation of services. Voice command, radar tracks, video, logistics traffic and software updates should not compete without priority rules. Time-critical air tracks require low latency and high integrity. Video can be compressed or delayed. Administrative traffic can wait. Quality-of-service policy must therefore be part of the architecture, not an afterthought.

Security is similarly layered. Encryption protects confidentiality, but it does not by itself prove that a device is trustworthy. Mutual authentication is needed so that nodes can verify one another. Signed software and configuration updates reduce supply-chain and maintenance risk. Segmentation prevents a compromised camera or access point from reaching every command service. Continuous monitoring detects unusual traffic, failed authentication and inconsistent sensor reports.



Figure 2.2. A directional wireless antenna similar to the type examined in the original point-to-point concept.

A distributed network has an apparent paradox. It is harder to destroy completely because there are many nodes and paths, yet it creates more devices that can fail or be compromised. The answer is not to trust every node equally. The network should assign confidence according to device identity, calibration, location, health and consistency with independent sensors. A single low-cost sensor may be wrong; ten independent sensors observing the same event can produce a useful track.

The transition from communications to air defence becomes possible when the network handles time and location precisely. Sensor reports must share a common clock, coordinate system and data model. Without time alignment, two reports of the same aircraft may appear to be different objects. Without position accuracy, a camera may be pointed to the wrong sector. The communications layer therefore becomes part of the sensing system.

The operational lesson is that bandwidth is not the same as combat value. A network creates value when it delivers the right information to the right user with known confidence and sufficient time for action. A slow but trustworthy report may be more useful than a fast false

alarm; a local track may be more important than a distant raw video stream. Architecture must be designed around decisions, not around maximum data volume.

Communications Layer Comparison

Layer	Strength	Primary limitation	Best role
HF radio	Beyond-line-of-sight reach	Variable propagation and low bandwidth	Long-range fallback and dispersed command
Tactical VHF/UHF	Mobile voice and data	Terrain and line-of-sight limits	Manoeuvre units and local coordination
Satellite	Wide-area connectivity	Latency, congestion and infrastructure dependence	Strategic reach and remote operations
Fibre/wired	High capacity and low emissions	Physical route vulnerability	Protected command and fixed sites
Point-to-point wireless	High local bandwidth and directionality	Alignment and terrain dependence	Local sensor, video and command links
Mesh/relay	Alternate paths and graceful degradation	More complex security and routing	Resilient local distribution

Selected sources: Ashraf (2023); Monzon Baeza et al. (2025).

CHAPTER 03

The Original Point-to-Point Architecture

Command centres, access points, directional antennas, cameras, sensors and field devices

RESEARCH STATUS

The original implementation demonstrated network and control logic in simulation. It did not establish operational air-defence performance. This book treats it as the foundation for a new testable architecture.

The original project defined three broad functional groups: command-and-control centres, access points and field communication devices. The command centre collected data, displayed maps and video, managed communications and supported decisions. Access points acted as bridges between the point-to-point backbone and cameras or IoT devices. Field devices allowed soldiers to communicate, receive updates, view surveillance feeds and control selected remote equipment.

This division remains useful because it separates strategic processing from local connectivity and user interaction. The central command centre does not need to communicate directly with every camera. Field access points aggregate local traffic, apply policy and forward information. The architecture can therefore expand by adding more access points without redesigning the entire command centre.

The original design also used separate central and field network segments connected through routed links. Firewalls, switches, servers and access-control rules were included. This was a valuable design choice because military networks should not be flat. A flat network allows a compromised endpoint to move laterally toward critical services. Segmentation creates boundaries and allows commanders to isolate a damaged field sector while preserving the wider mission.




Figure 3.1. Original field command-and-control network design from the project.

The use of cameras in the project was more than a surveillance feature. Cameras were treated as controllable sensors. Pan-tilt-zoom systems could be redirected toward an area of interest, while directional antennas could be aligned to preserve communications. This introduces the idea of active sensor management: the network does not simply receive data; it can task sensors to collect better data.

The project also explored remote weapon control. In a modern defensive architecture, the safest interpretation is remote cueing and supervised control rather than unrestricted autonomous firing. A command system may provide a track, predicted bearing and engagement window to a weapon crew. The crew verifies identity, rules of engagement and danger areas before action. This retains human responsibility while reducing the time needed to understand where the threat is located.

The original simulation used Cisco Packet Tracer, Google Earth, radio emulation and vendor configuration tools. These tools were appropriate for demonstrating network topology, routing, addressing and conceptual deployment. They could not validate radar physics, electronic warfare, real-time operating systems or live weapon integration. The project was therefore a proof of concept, not a field qualification.

A publishable research book must make this distinction explicit. Network connectivity can be demonstrated in software, but military effectiveness requires hardware-in-the-loop testing, realistic latency, packet loss, radio interference, clock error, sensor noise, operator workload and adversarial behaviour. The revised architecture therefore adds a test hierarchy: digital simulation, emulated networks, laboratory sensors, live field trials and adversarial exercises.




Figure 3.2. Refined central/field network design from the project.

The project remains important because it anticipated a trend now visible in counter-UAS programmes: sensors, command systems and effectors are being connected through common standards so that one platform can detect a target and another can engage it. In 2026 the U.S. Army described demonstrations in which a sensor on one vehicle classified a hostile drone and transmitted data to a weapon system on another platform. NATO's Layered Counter-UAS Initiative similarly focuses on integrating cells, sensors, command systems and response options.

The architecture therefore has contemporary relevance. Its strongest element is not the use of commercial-style Wi-Fi itself, because military deployments may require different radios and waveforms. The enduring idea is the functional pattern: distributed access, local aggregation, secure routing, remote sensing and central-local coordination.

The next step is to convert the original topology into a reference architecture with clear interfaces. Sensors publish standardised observations. Fusion cells produce recognised tracks. Command services assign priority and authority. Effectors receive only the minimum data needed for a safe response. Health monitoring reports node status. This interface-based approach allows old and new equipment to cooperate without requiring every device to use the same internal technology.



Figure 3.3. Google Earth-based link simulation used in the original research.

Prototype Components and Their Expanded Roles

Original component	Original role	Expanded air-defence role
Central C2	maps, communications, surveillance display	regional track fusion, tasking, policy and audit
Field C2	local coordination	edge fusion, local autonomy and degraded-mode control
Access point	connect IoT and field devices	sensor gateway, time synchronisation and health monitoring
PTZ camera	visual surveillance	EO identification and track confirmation
Directional antenna	point-to-point communications	low-latency data bridge and alternate relay
Field device	voice, video and control	local air picture, alerts and supervised effector cueing

Selected sources: Ashraf (2023); U.S. Army (2026a); NATO ACT (2026b).

CHAPTER 04

From Tactical Communications to Distributed Air Defence

The communications network becomes a sensor, command and response system

CENTRAL ARGUMENT

The network does not become valuable because it contains many devices. It becomes valuable when the devices create a recognised air picture, preserve command intent and deliver the right response before the threat reaches its objective.

Air defence is often represented as a radar circle around a missile battery. This image is useful for public understanding but incomplete. A functioning air-defence system is a chain: sensors detect an object; networks move observations; software and operators build a track; identification processes estimate whether it is friendly, neutral or hostile; command authorities prioritise threats; a suitable effector is assigned; the engagement is assessed; and the network adapts.

The point-to-point research architecture supplies the connective tissue for this chain. Cameras and IoT devices become sensor nodes. Access points become gateways. Field command centres become local fusion cells. The central command centre becomes a theatre or regional air-defence operations node. Directional links and relays connect them. The transformation is primarily one of data models, timing, resilience and command policy.

A distributed air-defence network should distinguish raw observations from recognised tracks. A radar report may contain range, bearing and Doppler information. A camera may provide an image and line of sight. A passive RF sensor may detect an emitter. An acoustic sensor may estimate a bearing. A field observer may report a visual contact. None of these observations is automatically a target. The fusion process determines whether multiple reports correspond to the same object.




Figure 4.1. The lower-layer air picture combines strategic reach with dense local sensing.

The network should also distinguish detection from engagement quality. A low-cost sensor may indicate that something is present without providing sufficient accuracy for a weapon. The system can then task a better sensor to confirm the track. This prevents the common mistake of expecting every node to perform every function. Cheap nodes provide persistence and coverage; premium nodes provide precision and discrimination.

The distributed model is particularly useful for the lower air layer. Strategic radars are optimised for wide-area surveillance and higher-value threats. Small drones may fly slowly, close to terrain, among buildings or in clutter. Their radar cross section may be small, and they may emit little or no radio energy. Local cameras, acoustic sensors and short-range radars can fill these gaps when their reports are fused.

The command design should be hierarchical but not rigidly centralised. National or theatre command defines policy, identification rules and priorities. Regional nodes coordinate sectors and allocate scarce resources. Local fusion cells maintain the immediate picture and can continue limited operations if disconnected. This is mission command applied to air defence: central intent with delegated local action inside defined limits.

The idea of thousands of nodes should not be interpreted as thousands of independent weapon systems. Most nodes may be passive or dual-use sensors, communication relays, observation posts or inexpensive cameras. Their collective role is to maintain awareness and prevent the destruction of a small number of strategic sensors from causing blindness.




Figure 4.2. The continuous sense-share-fuse-decide-engage-assess cycle.

A dense network can also protect civilian infrastructure. Ports, airports, power stations, rail hubs and urban areas already contain cameras, weather sensors and communications equipment. Selected feeds might contribute to warning under emergency law and strict governance. This creates opportunities but also serious privacy and cybersecurity risks. Military use of civilian sensors must therefore be controlled, temporary, auditable and legally authorised.

The distributed shield is best understood as a mosaic. Each cell contains a small group of sensors, a local processing node, communications and one or more response options. Cells can be added, moved, isolated or replaced. Regional command sees the mosaic as a coherent picture, while local cells retain enough information to function during disruption.

The operational objective is continuity. A sparse architecture may deliver excellent performance until a key radar, command centre or data link is lost. A dense architecture accepts that individual nodes will fail and is designed so that the mission continues. This is the difference between protecting equipment and protecting function.

Selected sources: NATO (2025); MDA (2026a); NATO ACT (2026a).

CHAPTER 05

Building the Lower-Layer Air Picture

How micro-radars, EO/IR, passive RF, acoustic sensors and field observers complement strategic surveillance

DESIGN PRINCIPLE

Every sensor report must include time, location, uncertainty, identity and health. A fusion engine cannot produce trustworthy tracks from unqualified data.

The lower-layer air picture is a continuously updated representation of activity close to the ground, sea surface, buildings and terrain. It is the part of the battlespace where small drones, loitering munitions, cruise missiles and helicopters may exploit clutter and radar geometry. It is also the layer in which friendly aircraft, civilian aviation and ground weapons create the greatest identification challenge.

No single sensor is sufficient. Radar provides range and velocity but may struggle with small radar cross sections, slow targets and clutter. EO/IR systems can confirm shape and behaviour, but their performance depends on weather, visibility, background and line of sight. Passive RF sensors can detect control links or emissions, but autonomous drones may remain silent. Acoustic sensors are inexpensive and useful at short range, yet wind, vehicles and urban noise produce false alarms. Human observers add context but vary in training and accuracy.

Fusion begins by understanding these limitations. The goal is not to force all sensors to agree. It is to combine partially independent evidence. A radar track and an EO/IR image may provide high confidence. A passive RF bearing without a radar plot may still justify re-tasking a camera. An acoustic detection near a protected site may trigger local observation even when strategic radar reports nothing.




Figure 5.1. Strategic and local sensors cooperate to build a lower-layer air picture.

Distributed sensing also improves geometry. A small drone hidden behind terrain from one sensor may be visible to another. Multiple bearings can be combined to estimate location. A moving target can be handed from one cell to the next. Overlapping coverage reduces the chance that the destruction or obstruction of one sensor creates a permanent gap.

Time synchronisation is critical. A drone travelling at modest speed can move tens or hundreds of metres while data is delayed. Each observation should therefore include a reliable timestamp, sensor position, uncertainty and health status. The fusion engine should account for latency instead of assuming that every report describes the present moment.

The recognised picture should display uncertainty. A track can have a confidence score, position error, classification probability and source history. Operators should be able to see whether a track is supported by radar only, camera confirmation, passive RF or human observation. This reduces automation bias and allows commanders to decide when more evidence is required.

The lower-layer picture also needs friendly-force information. Counter-drone operations create fratricide risk because friendly drones may resemble hostile systems. Blue-force tracking, mission plans, electronic identification and temporary airspace control measures should be fused with sensor data. A system that detects every drone but cannot distinguish friendly from hostile activity is not operationally mature.

Edge processing reduces bandwidth and latency. Rather than transmitting every video frame to central command, a local node can detect movement, extract features and forward a short report or selected clip. This also improves resilience because the local cell can continue to track threats when wide-area communications are degraded.

The architecture should support sensor management. A fusion cell can command a PTZ camera to observe a radar plot, increase the update rate of a micro-radar, request another

cell's bearing, or launch a small observation drone. This active process converts a collection of sensors into a coordinated sensing system.

The lower-layer picture should not be confused with complete control of the airspace. It is a decision-support product with known gaps and confidence. Its value lies in providing earlier warning, reducing false tracks and preserving local awareness under pressure.

Sensor Contribution Matrix

Sensor	Primary contribution	Typical weakness	Best fusion partner
Long-range radar	wide-area detection and cueing	terrain and small-target clutter	local radar and EO/IR
Micro-radar	local range, bearing and velocity	limited coverage and power	EO/IR and passive RF
EO/IR camera	visual identification and behaviour	weather and line of sight	radar cueing
Passive RF	emitter detection and direction	silent/autonomous targets	radar and acoustic
Acoustic array	short-range bearing and signature	noise and weather	camera and micro-radar
Field observer	context and unusual-event reporting	human error and delay	all technical sensors
Friendly-force feed	identification and deconfliction	data freshness and compliance	fusion engine and command plan

Selected sources: Farlik et al. (2019); Park et al. (2021); Rojhani and Shaker (2024).

CHAPTER 06

Multi-Sensor Fusion and Track Confidence

Turning observations into recognised tracks without hiding uncertainty

FUSION RULE

A low-cost node may contribute to a high-confidence track, but it should not create one by itself. Confidence comes from independent, time-consistent evidence and transparent uncertainty.

Multi-sensor fusion is often described as combining data, but the difficult part is deciding which observations belong together. Two sensors may report one object at slightly different times and positions. One sensor may report clutter. Several drones may cross. A spoofed node may inject a false target. The fusion engine must associate reports while preserving uncertainty and source provenance.

A practical workflow contains six stages. Detection produces raw observations. Normalisation converts them into common units, coordinates and timestamps. Correlation estimates which reports refer to the same object. Classification estimates object type and behaviour. Scoring calculates confidence and threat priority. Cueing sends the recognised track to another sensor, a commander or an effector.

Track confidence should not be a single unexplained percentage. It should reflect the number and independence of sources, sensor quality, geometry, time consistency, track continuity, classification evidence and known interference. A radar and a camera mounted on the same vehicle do not provide the same independence as two sensors separated by distance and technology.





Figure 6.1. A six-stage fusion workflow from raw detection to safe cueing.

The system also needs negative evidence. If a radar reports a target but several cameras with clear line of sight see nothing, confidence should fall. If a passive RF sensor detects a control link but no physical sensor sees an object, the system may classify the event as an emitter of interest rather than an air track. Fusion must be able to disagree with its inputs.

Quorum or confidence voting can make the network more resilient to false data. A single low-trust node cannot create a high-confidence engagement track. Several independent nodes can raise confidence. A previously reliable sensor that suddenly reports impossible movement can be quarantined for calibration or cyber investigation.

Track management also requires rules for initiation, confirmation, maintenance and deletion. Immediate confirmation may be impossible during saturation, but keeping every uncertain report can overload operators. The system should use different thresholds for surveillance, warning and engagement. A low-confidence report may justify observation; it should not automatically justify lethal action.

Artificial intelligence can assist with association and classification, especially when the network produces thousands of observations. Machine-learning models can detect visual objects, recognise radio signatures, estimate trajectories and identify anomalies. Their outputs should remain advisory and should include confidence, input sources and model status.



Visual Plate 6A. Multi-sensor fusion for lower-layer air defence.

Classical methods remain important. Kalman filters, Bayesian inference, probabilistic data association and rule-based checks are transparent and mathematically grounded. Machine learning can complement them but should not replace basic physics and consistency tests. A

hybrid system is more defensible because it combines learned patterns with known constraints.

The data architecture should record provenance. Operators and investigators should be able to reconstruct which sensors contributed to a track, how the confidence changed, which software version produced the recommendation and who approved the response. This supports accountability, training and improvement.

The most important fusion result is not perfect classification. It is a sufficiently reliable, timely and explainable picture that allows commanders to allocate attention and defensive resources. A good system knows what it knows, what it does not know and what sensor should be tasked next.

Track Confidence Factors

Factor	Effect on confidence
Independent corroboration	raises confidence when different sensor types agree
Fresh timestamp and low latency	raises confidence in current position
Good geometry	reduces position error and ambiguity
Continuous history	raises confidence in identity and trajectory
Sensor health warning	reduces confidence or removes source
Inconsistent classification	reduces confidence and triggers re-tasking
Known jamming or spoofing	increases uncertainty and requires corroboration
Friendly mission-plan match	supports deconfliction and reduces hostile classification

Selected sources: Park et al. (2021); Farlik et al. (2019); Smailov et al. (2025).

CHAPTER 07

AI-Assisted Decision Support and Human Authority

Using automation to manage scale while preserving accountability and safety

RESEARCH CORRECTION

The book retains AI for track fusion, anomaly detection, sensor management and decision support. It rejects emotion analysis as a dependable basis for assessing hostile intent.

The strongest case for artificial intelligence in air defence is scale. A dense network may produce more observations than human operators can evaluate individually. AI can filter duplicate reports, associate observations, detect anomalous behaviour, estimate short-horizon trajectories and recommend sensor tasking. It can reduce cognitive load and allow operators to focus on ambiguous or high-consequence decisions.

AI is also useful for resource management. A system can estimate which sensor has the best geometry, which local cell is overloaded, which effector has sufficient range and which response preserves scarce interceptors. These recommendations can shorten the decision cycle, especially during simultaneous attacks.

The most dangerous error is to confuse recommendation with authority. Lethal action requires identity, legal authority, rules of engagement, proportionality and awareness of friendly or civilian activity. An AI model cannot assume responsibility for these judgements. Human command should remain the authority for lethal engagement, with automation limited to clearly defined defensive functions and fail-safe conditions.



Figure 7.1. AI-assisted air defence with human command authority and explicit safeguards.

The original project proposed facial recognition, behaviour analysis and emotion detection as future features. Facial recognition can support identity checks in controlled contexts, but behaviour and emotion inference are unreliable indicators of hostile intent. Expressions vary across people and cultures, surveillance conditions are imperfect, and models can reproduce bias. The revised architecture therefore removes emotion detection from the core military decision process.

AI systems are themselves attack surfaces. Adversaries may spoof sensors, poison training data, exploit software vulnerabilities or deliberately create patterns that confuse classifiers. The system should therefore use diverse models, independent rule-based checks, signed software, controlled update processes and continuous monitoring for drift.

Explainability is operational, not merely ethical. An operator needs to know why a track was prioritised. The interface should display contributing sensors, confidence, predicted path, friendly-force conflicts and alternative explanations. A recommendation that cannot be explained should receive less authority, not more.

Human-machine teaming requires calibrated trust. Operators should understand where automation is strong and where it is weak. Training should include false alarms, degraded sensors and deliberately wrong recommendations. The objective is neither blind trust nor automatic rejection, but informed use.

The command structure should define automation levels. At the lowest level, AI filters and displays. At a higher level, it re-tasks sensors. It may allocate non-kinetic responses such as warning or electronic protection within policy. It may recommend an engagement. Only under narrow, pre-authorised conditions should automated defensive action occur, and even then the system requires audit, abort mechanisms and positive safety constraints.

AI can also support network defence. It can identify unusual traffic, inconsistent device behaviour, calibration drift and improbable sensor patterns. This turns the fusion system into a guardian of its own inputs. A sensor that behaves abnormally can be down-weighted or isolated before it corrupts the picture.

The final principle is that AI should increase the quality of human command. It should create more time, better evidence and clearer options. Any design that accelerates action while hiding uncertainty has made the command problem worse.

Selected sources: Monzon Baeza et al. (2025); Smailov et al. (2025); U.S. Army (2026b).

CHAPTER 08

Network-Centric Warfare, Command and Data Links

Why the decisive capability is the networked process rather than the individual platform

VALIDATION IN 2026

NATO and U.S. Army counter-UAS programmes increasingly describe connected sensors, command systems and effectors, distributed cells, common data standards and formation-level layered protection.

Network-centric warfare seeks to convert distributed information into coordinated action. The concept is sometimes reduced to connecting every platform, but connectivity is only the first step. A useful network must create shared awareness, distribute command intent, support local initiative and adapt when links fail.

The distributed shield follows a sense-share-fuse-decide-engage-assess cycle. Sensors detect and track. Networks share observations. Fusion nodes create a common picture. Commanders prioritise threats. Effectors act. Assessment feeds back into the system. The cycle is continuous because the environment changes during every engagement.

Data links must therefore carry more than coordinates. They must carry identity, confidence, time, source, tasking and status. They also need priority and congestion control. A track update may be urgent; a software update is not. During saturation, the network should preserve command and track data even if video quality falls.





Figure 8.1. The network-centric air-defence process.

Interoperability is a central challenge. Legacy equipment may use proprietary interfaces, analogue signals or incompatible message formats. A gateway can translate these into a common data model, but translation must not create false precision. If an old radar reports a broad bearing, the network should not display it as an exact coordinate.

The architecture should use local autonomy because communications will not always be available. A local cell should maintain tracks, apply identification rules and defend a protected site within delegated authority. When links return, it synchronises with the wider picture. This prevents an adversary from paralysing the system merely by isolating a sector.

At the same time, local autonomy must not become uncoordinated firing. Airspace control, friendly-force plans and engagement zones need common rules. Command intent defines what local units may do, against which threats and under what confidence. The network supports decentralised execution inside a central policy framework.

The system is harder to compromise completely because information can travel through alternate paths and no single node holds the entire mission. A destroyed relay can be bypassed. A compromised sensor can be isolated. A local cell can continue with reduced information. The network loses performance gradually rather than becoming blind at once.



Visual Plate 8A. Faults, compromise and recovery in a distributed defence network.

This resilience depends on architecture. If every low-cost node still depends on one central authentication server or one data-fusion hub, the network remains centralised in practice. Critical services must be replicated, and local cells must be able to operate with cached policy and verified keys.

NATO's 2025 integrated air and missile defence policy explicitly emphasises resilient and redundant air command-and-control, distributed nodes, resilient networks and backup solutions. NATO's 2026 Layered Counter-UAS Initiative moves from individual counter-UAS cells toward connected mosaics. These developments independently validate the central logic of the author's research extension.

Network-centric warfare is therefore a discipline of organisation. It requires common standards, shared training, trusted identities, resilient communications and commanders who understand the quality of information. Technology can shorten a kill chain, but only doctrine prevents it from becoming a chain of fast mistakes.

Selected sources: NATO (2025); NATO ACT (2026a); NATO ACT (2026b); U.S. Army (2026a); U.S. Army (2026c).

CHAPTER 09

Few High-End Nodes versus Thousands of Small Nodes

Coverage, resilience, cost, battle-damage tolerance and the importance of hybrid design

HYBRID JUDGEMENT

The dense layer should not imitate the premium layer. Its purpose is persistence, local awareness, redundancy and low-cost response. The premium layer provides reach, discrimination and engagement against the most demanding threats.

A sparse premium network concentrates capability in a small number of powerful systems. A long-range radar can search a large volume. An airborne early-warning aircraft can extend the horizon. A strategic missile battery can engage demanding threats. A central command post can maintain high-quality control. Such systems are indispensable, but their concentration creates risk.


The first risk is geometry. A wide-area sensor may still have terrain, clutter or horizon limitations. The second risk is capacity. A radar or command centre can process many tracks, but every system has limits. The third risk is battle damage. Losing one of twelve critical nodes is more significant than losing one of a thousand small nodes. The fourth risk is economics: premium interceptors and specialised maintenance are difficult to replace rapidly.

A dense distributed network trades individual quality for quantity and overlap. A micro-radar may cover only a local sector. A camera may be weather-dependent. An acoustic sensor may be noisy. Yet many nodes of different types can cover one another's gaps. Local fusion can produce a useful picture even when individual observations are imperfect.


Sparse Premium Architecture vs Dense Distributed Architecture

The strongest design combines both: premium reach with mass local resilience

SPARSE PREMIUM NETWORK



DENSE DISTRIBUTED NETWORK



Coverage	wide-area, high-quality	dense local coverage and overlap
Single-node loss	large local gap	graceful degradation
Cost	high unit and sustainment cost	lower unit cost, larger quantity
Scalability	slow procurement cycle	incremental and modular
Distributed nodes are not a replacement for high-end systems; they are the survivable foundation that protects and extends them.		
Low-altitude picture	limited by geometry and clutter	improved through local sensing
Cyber burden	smaller attack surface	larger surface, but compartmentalised


Figure 9.1. Sparse premium and dense distributed architectures compared.

The dense architecture also supports reconstitution. A damaged node can be replaced by a commercial or locally produced unit. A new relay can be added without redesigning the national system. Software can redirect data around failed areas. This makes the network better suited to prolonged conflict and repeated attacks.

However, density creates its own burdens. Thousands of devices require power, maintenance, calibration, identity management and software updates. They increase the cyber attack surface. They can generate false alarms and network congestion. A distributed system without disciplined management may become an expensive collection of unreliable data.

The correct comparison is therefore not premium versus cheap. It is an architecture dominated by a few premium nodes versus a hybrid architecture in which premium nodes are protected and extended by a mass layer. High-end sensors provide reach and discrimination. Low-cost sensors provide persistence, local detail and redundancy. High-tier interceptors defeat the threats that require them. Low-cost effectors absorb the mass.

The concept of thousands of nodes should also be scalable. A national network might contain many regional cells. A brigade or city may operate hundreds. A protected base may operate dozens. The same data model and security principles allow these levels to cooperate without requiring every sensor to report directly to national command.



Visual Plate 9A. Few high-end nodes versus mass distributed nodes.

The book's simulation therefore compares two idealised architectures. The sparse model contains twelve major sensors and three command nodes. The dense model contains 1,024 low-cost sensors and sixty-four local fusion cells, while retaining a smaller set of high-tier

systems. The exact numbers are hypothetical; the purpose is to observe how architecture behaves under node loss, jamming and saturation.

The dense network performs better in continuity because its functions are distributed. The sparse network performs better in some individual measurements because its sensors are more capable. This is why a hybrid system is superior: the dense layer keeps the premium layer informed, preserves local awareness and gives commanders more options.

The strategic insight is that survivability comes from replacing single points of excellence with networks of complementary excellence. A nation does not need every node to be world-class. It needs the system to continue producing useful decisions after the first attack.

Architecture Comparison

Dimension	Sparse premium network	Dense distributed layer	Hybrid judgement
Wide-area reach	excellent	limited per node	premium sensors remain necessary
Low-altitude coverage	variable	strong through local overlap	dense layer fills gaps
Battle-damage tolerance	concentrated risk	graceful degradation	distribute critical functions
Cost per node	very high	low to moderate	use quality where required
Maintenance burden	complex specialist support	large quantity and calibration burden	standardise and modularise
Cyber exposure	fewer devices	many endpoints	zero trust and segmentation are essential
Reconstitution	slow replacement	rapid incremental replacement	stock common nodes and relays

Selected sources: NATO (2025); NATO ACT (2026a); Harrell (2026).

CHAPTER 10

Legacy Anti-Aircraft Guns and Low-Cost Effectors

How existing weapons can become useful nodes in a modern lower-layer defence

SAFETY RULE

A network cue is not an engagement order. Legacy weapons require positive identification, weapon-limit checks, danger-area control and trained human authority.

Many armed forces retain anti-aircraft guns, heavy machine guns, autocannons and manually controlled mounts that were designed before modern drone warfare. These systems may appear obsolete when compared with long-range missiles, yet they can become useful against suitable low, slow and close threats if they receive accurate cueing and operate under disciplined safety rules.

Digital integration does not mean that an old gun becomes a precision autonomous weapon. It means that the network can provide bearing, elevation, estimated range, target direction and engagement timing to a trained crew or a modernised fire-control unit. The crew confirms the target and acts within rules of engagement. The network then re-acquires the target and assesses the result.

The value is economic. Guns and programmable ammunition can offer a lower cost per engagement than premium missiles. They can preserve scarce interceptors for faster or more dangerous threats. They can also provide a final defence layer around bases, command posts and critical infrastructure.




Figure 10.1. Digital cueing of legacy anti-aircraft guns with human confirmation.

The limitations are serious. Guns have range, accuracy and ballistic constraints. Firing into urban airspace creates danger to civilians and friendly forces. Ammunition may fall beyond the target. Identification must be positive. A false track or inaccurate cue can produce catastrophic consequences. The integration gateway must therefore include safety zones, weapon limits and human confirmation.

Other low-cost effectors include electronic warfare, interceptor drones, nets, short-range missiles, directed energy and small-calibre airburst systems. Each is suitable for different conditions. Electronic warfare can be effective against controlled drones but less effective against autonomous navigation. Lasers offer low cost per shot but depend on weather, line of sight, dwell time and power. Interceptor drones can be economical but require their own sensors and control.

The correct response is selected according to threat, range, confidence, collateral risk and magazine state. A suspected drone may first be observed. A confirmed control link may be jammed. A small interceptor may be launched. A gun may be cued for a close target. A missile is used when speed, altitude or consequence requires it.

This decision is a cost and risk optimisation problem. The cheapest option is not always correct, because a failed interception may be more expensive than the interceptor. The network therefore recommends the lowest-cost response that has sufficient probability of success and acceptable risk.




Figure 10.2. Layered effectors and the cost-exchange logic.

Legacy integration also creates a training benefit. Existing crews become part of the air-defence picture rather than isolated local units. They receive alerts, understand the wider threat and report outcomes. Their observations can also improve the picture, especially when electronic sensors are degraded.

The architecture should use standard gateways rather than custom integration for every weapon. A gateway translates recognised tracks into safe cueing messages and reports weapon status. This allows the network to add or remove effectors without rebuilding the fusion system.

The final principle is proportionality. A distributed network should reduce the pressure to use a premium missile against every object. By offering more response levels, it improves endurance and keeps high-end systems available for their designed missions.

Effector Selection Logic

Effector	Strength	Weakness	Appropriate use
Electronic warfare	low cost and reusable	limited against autonomous or hardened systems	control-link disruption and area protection
Legacy gun / autocannon	low cost per round and existing inventory	range, safety and accuracy limits	close defence against suitable confirmed targets
Interceptor drone	flexible and potentially low cost	guidance, weather and recovery burden	local pursuit of small UAS
Short-range missile	high probability against manoeuvring threats	finite magazine and higher cost	dangerous UAS and cruise threats
Directed energy	low marginal shot cost	weather, dwell and power limits	repeated close engagements in favourable conditions
High-tier interceptor	reach and performance	scarce and expensive	ballistic, high-speed or high-consequence threats

Selected sources: U.S. Army (2026d); U.S. Army (2026e); Harrell (2026).

CHAPTER 11

High-Tier Systems: S-400, THAAD, Aegis/Standard Missile and Iron Dome

Their real missions, their strengths and why saturation is an architectural challenge

PRECISE FINDING

The predicted vulnerability is not that every high-end interceptor fails. It is that a defence centred on a few expensive nodes can be saturated, isolated or economically exhausted unless it is supported by dense sensing, local effectors and resilient command.

Public debate often places every air-defence system in one category, but S-400, THAAD, Aegis with the Standard Missile family and Iron Dome were designed for different threat sets. A serious analysis must begin with role rather than advertised range.

S-400 is a long-range integrated air-defence system intended to engage aircraft, cruise missiles and selected missile threats through a family of sensors, command components and interceptors. Its strength lies in wide-area coverage and multiple missile options. It is not a universal solution for large numbers of small drones operating close to terrain. Such threats require local sensors, point defence and low-cost effectors around the strategic battery.

THAAD is a terminal ballistic-missile defence system. The U.S. Missile Defense Agency describes it as a rapidly deployable hit-to-kill capability against ballistic missiles inside or outside the atmosphere during their terminal phase. Judging THAAD by its ability to defeat small quadcopters would misunderstand its mission. The correct question is whether the wider defence protects THAAD's radar, launchers and command components from lower-tier threats.




Figure 11.1. High-tier systems perform different missions and face different saturation pressures.

Aegis is a naval combat system that integrates radar, command and weapons. The Standard Missile family supports fleet air defence, over-the-horizon engagements and ballistic-missile defence depending on variant. Aegis demonstrates the value of a system designed from detection to engagement, but ships still face magazine limits and cost-exchange pressure during mass raids. Cooperative engagement and off-board sensing extend reach, yet local point defence remains necessary.

Iron Dome was designed for shorter-range rockets, artillery and mortars and has also developed broader air-defence applications. It is more relevant to mass short-range threats than THAAD, but no interceptor-based system has infinite capacity. Prolonged saturation creates questions of ready rounds, reload time, production and cost. Israel's development of laser and additional layered systems reflects the need for different engagement options.

The author's prediction should therefore be expressed precisely. These systems may fail to protect a defended objective when they are used as isolated high-end nodes against a dense, mixed and prolonged attack. The failure mechanism may be sensor overload, track ambiguity, insufficient low-altitude coverage, command saturation, interceptor exhaustion, reload delay, electronic attack or an unfavourable economic exchange. This is not the same as claiming that the systems are technically ineffective.

The distributed shield reduces these risks in four ways. First, local sensors provide earlier warning of small threats. Second, fusion filters false tracks and prevents the high-tier system from wasting attention. Third, low-cost effectors defeat suitable targets. Fourth, alternate command and communications allow the high-tier battery to continue operating if one local link or sensor is lost.

The high-tier system also contributes to the distributed layer. Its radar can cue local cells. Its command system can provide policy and identification. Its interceptors can defeat threats that escape local defences. The relationship is therefore reciprocal.

Magazine management should be treated as an operational function. The network must know which interceptors are available, their reload status, threat priority and the consequences of conserving or firing. During a mass raid, a commander may accept risk against low-value objects in order to preserve interceptors for ballistic or high-consequence threats.

The overall conclusion is that the future of air defence is layered and connected. High-tier systems remain essential, but their survivability and strategic value depend increasingly on dense lower layers, resilient command and a cost-effective family of responses.

System Role Comparison

System	Primary public role	What it contributes	Why it still needs a distributed lower layer
S-400	long-range integrated air defence	wide-area sensing and multiple interceptor classes	small UAS, terrain masking, local defence and battery protection
THAAD	terminal ballistic-missile defence	high-altitude hit-to-kill interception	protection from drones, cruise threats and local attacks
Aegis / Standard Missile	naval air and missile defence	integrated radar, C2 and missile family	magazine depth, local point defence and distributed sensing
Iron Dome	short-range rocket, artillery and mortar defence	automated threat evaluation and short-range interception	prolonged saturation, reloads and cheaper engagement options

Selected sources: MDA (2026b); U.S. Navy (2021a); U.S. Navy (2021b); RTX (2026); Rosoboronexport (2024).

CHAPTER 12

Drone Swarms, Saturation and the Cost-Exchange Problem

How mass low-cost threats attack the entire defensive process

MISSION-ORIENTED DEFENCE

Success is not the destruction of every drone. Success is the continued operation of the protected mission at an acceptable cost and level of risk.

A drone swarm is not merely a large number of aircraft. It is a method of creating simultaneous demands across detection, classification, command, engagement and resupply. Some drones may carry sensors, some may jam, some may act as decoys and some may attack. The defender must determine which objects are real threats before they reach protected assets.

Saturation can begin at the sensor layer. Numerous small tracks appear in clutter. Decoys and civilian objects complicate identification. Electronic attack reduces radar or communications performance. The system may be forced to devote high-quality sensors to low-value objects while a more dangerous threat approaches from another direction.

The command layer can also be saturated. Operators must evaluate more tracks, assign more sensors and select more effectors. Automation helps, but poor automation can produce a larger number of fast mistakes. The system needs prioritisation rules based on intent, proximity, payload, protected asset and confidence.




Figure 12.1. Drone swarms stress detection, classification, command, magazines and economics simultaneously.

The engagement layer faces capacity and economics. A launcher has a finite number of ready rounds. Reloading takes time and may expose crews. Production cannot instantly replace expenditure. A high-cost interceptor may defeat a low-cost drone, but repeated use can be strategically unsustainable even when tactically successful.

Swarms also exploit recovery time. A first wave may reveal radar locations, consume ammunition and damage communications. A second wave attacks during reload or repair. A third wave targets the remaining command nodes. The defender's ability to reconstitute becomes as important as its first engagement probability.

A distributed network improves endurance because it creates many local decision points and engagement options. Small sensors continue to report after a strategic radar is suppressed. Local effectors defeat low-end threats. Mobile relays restore connectivity. Damaged sectors can be isolated while neighbouring cells cover the gap.

The network should also use deception. Decoy sensors and emitters complicate enemy targeting. Emission control reduces the time that valuable radars reveal their location. Passive sensors allow the system to observe without transmitting. Mobile nodes change the defensive geometry between waves.




Figure 12.2. Layered effectors preserve high-tier magazines.

The cost-exchange problem is not solved merely by buying cheap interceptors. The defender must also reduce the cost of sensing, command and maintenance. A low-cost drone defence that requires a highly specialised crew at every site may still be expensive. Standardised nodes, common software and simple replacement are therefore part of the design.

The swarm threat also creates a data problem. The network must distinguish a coordinated attack from unrelated activity. Group behaviour, common timing, shared origin and converging paths can indicate coordination. AI can identify these patterns, but commanders should understand the evidence and uncertainty.

The strategic objective is not to intercept every object. It is to preserve the defended function. The network may allow harmless or low-confidence objects to continue under observation while concentrating resources on threats to command, airfields, radars, logistics and population centres. This is mission-oriented defence rather than object-oriented defence.

Selected sources: Park et al. (2021); Khawaja et al. (2022); Harrell (2026); NATO ACT (2026a).

CHAPTER 13

Resilience against Jamming, Spoofing, Cyber Attack and Physical Attrition

How a distributed network detects defects, isolates compromise and continues the mission

RESILIENCE PRINCIPLE

It is unrealistic to prevent every node from failing. The defensible objective is to detect failure early, contain its effect, route around it and preserve mission function.

A distributed network is designed to survive physical loss, but physical destruction is only one failure mode. Jamming can break links or blind sensors. Spoofing can create false identities and tracks. Cyber intrusion can alter software or configuration. Power loss can remove entire clusters. Calibration drift can produce incorrect data without any hostile action.

The first requirement is observability. Every node should report health, time quality, software version, authentication status, calibration state and link performance. The network cannot manage defects that it cannot see. Health information should be separated from mission data so that operators can distinguish a real threat from a failing sensor.

Cross-sensor verification is the primary defence against false reports. A compromised node may inject a target, but independent sensors should confirm its movement. Confidence voting and physics-based checks can reject impossible acceleration, inconsistent location or an emitter without a corresponding object.




Figure 13.1. Failure modes, compromised nodes and resilient recovery paths.


Zero-trust segmentation limits cyber movement. Nodes authenticate every connection, receive only the privileges required for their function and are isolated when behaviour changes. A

camera gateway should not have administrative access to a regional command server. A local fusion cell should not be able to rewrite national policy.

Software supply-chain security is essential because thousands of nodes require updates. Updates should be signed, tested and distributed through controlled stages. The system needs an inventory of hardware and software components so that a vulnerability can be identified across the network.

Resilient routing provides alternate paths. If a point-to-point link is jammed, data may move through another relay, fibre, tactical radio or store-and-forward process. Not every path has equal capacity, so the network should preserve high-priority track and command data before video or administrative traffic.

Local autonomy reduces dependence on central services. A cell can continue with cached identification rules, trusted keys and recent tracks. It can defend a protected site within delegated authority. When communications return, the cell reconciles its local picture with the regional network.



Visual Plate 13A. Distributed defence network compromise and recovery.

Quarantine is more valuable than immediate shutdown when uncertainty exists. A suspicious sensor can be down-weighted, restricted to observation-only status or isolated for inspection. This preserves some information while preventing it from controlling effectors.

Physical resilience includes mobility, camouflage, decoys, spare power, protected cabling and rapid replacement. A cheap sensor node is useful only if it can be installed, aligned and authenticated quickly. Standard connectors and pre-configured identities reduce reconstitution time.

The central insight is graceful degradation. A resilient network should lose accuracy, coverage or speed gradually as damage increases. It should not move directly from full capability to blindness. This property is measurable and becomes a central metric in the war-gaming chapter.

Threat-to-Mitigation Matrix

Failure or attack	Immediate effect	Architectural mitigation
Jamming	lost or degraded links	alternate paths, frequency diversity, local autonomy
Spoofed sensor report	false track or wrong classification	cross-sensor checks, provenance and confidence voting
Compromised endpoint	malicious data or lateral movement	zero trust, least privilege and quarantine
Power loss	node or cluster offline	battery, generator, energy monitoring and mobile replacement
Physical destruction	coverage gap	overlap, mobility and hot-swap nodes
Calibration drift	systematic track error	health monitoring, comparison and scheduled calibration
Central-service loss	authentication or C2 disruption	replicated services and cached local policy

Selected sources: NATO (2025); U.S. Army (2026f); Monzon Baeza et al. (2025).

CHAPTER 14

Interoperability, Gateways and Edge Computing

How legacy and modern devices contribute to one picture without losing their limitations

INTEROPERABILITY PRINCIPLE

The network should make different systems cooperate without hiding their differences. A gateway translates and validates; it does not manufacture accuracy.

The distributed shield depends on interoperability because its value comes from diversity. A modern micro-radar, an older surveillance radar, a commercial-style camera, a passive RF receiver and a field report do not produce the same type or quality of data. The architecture must allow them to cooperate without pretending that their measurements are equally precise.

A gateway is the boundary between a device and the wider network. It translates proprietary or analogue information into a common observation message, attaches a timestamp and location, records uncertainty, signs the message and reports node health. A gateway should never improve a measurement merely by changing its format. A broad bearing remains a broad bearing after translation.

The common data model should separate observations from tracks. Sensors publish what they measured. Fusion cells decide how observations relate to objects. Command nodes decide priority and authority. Effector gateways receive cueing only after the track has reached the required confidence and safety level. This separation reduces the risk that a faulty or compromised sensor directly controls a weapon.




Figure 14.1. Gateways and edge computing translate, validate and preserve uncertainty across diverse systems.

Edge computing is important because thousands of nodes cannot stream all raw data to one centre. A camera can detect movement and forward a selected clip. A passive RF node can report a signature and bearing rather than the entire sampled spectrum. A local radar can maintain a track and transmit updates instead of every return. The edge therefore reduces bandwidth and preserves local function when long-range links fail.

Edge processing must be controlled. Different nodes may run different software versions or models, producing inconsistent results. The network needs signed software, configuration baselines, health reporting and the ability to compare edge conclusions with raw evidence. Important raw data should be retained for a limited period when bandwidth and storage allow.

Interoperability also applies to command. A regional node may receive tracks from military sensors, civil aviation systems and coalition partners. Security labels, release rules and confidence need to travel with the data. The system should share the minimum information required for the mission while protecting sensitive sources.

Common standards reduce vendor dependence. A new sensor can join by implementing observation and health interfaces. A new effector can join through a supervised cueing gateway. The fusion engine does not need to understand the internal electronics of every device. This makes incremental procurement and local industrial participation more practical.

Legacy systems benefit most from gateways. An older radar may lack digital networking, but a local processor can convert its plot into a time-stamped observation. An anti-aircraft gun may lack modern fire control, but a gateway can provide supervised bearing and elevation cues. The network recovers operational value without claiming that the old equipment has become equivalent to a new system.

The main technical risk is semantic mismatch. Two systems may use the same word, such as track or confidence, with different meanings. Interface standards must therefore define units, coordinate frames, update rates, uncertainty and authority. Interoperability is not achieved when packets move; it is achieved when every participant interprets them correctly.

A successful distributed network will treat gateways and data models as strategic infrastructure. They are the means by which thousands of varied nodes become one system rather than a collection of disconnected products.

Core Interoperability Rules

Rule	Reason
Preserve source uncertainty	prevents false precision
Separate observation from track	prevents raw sensors from commanding effectors
Attach time and provenance	supports correlation and accountability
Use signed interfaces	protects integrity and software trust
Operate locally when disconnected	preserves mission continuity
Translate semantics, not only packets	ensures systems interpret data consistently

Selected sources: Ashraf (2023); U.S. Army (2026b); NATO ACT (2026b).

CHAPTER 15

Electronic Warfare, Emission Control and Passive Sensing

Surviving when the electromagnetic spectrum is contested

SPECTRUM PRINCIPLE

A network is more survivable when it can change how it senses and communicates. Continuous high-power transmission should be an option, not a dependency.

A networked air-defence system depends on the electromagnetic spectrum, and an adversary will attempt to deny, deceive or exploit it. Jamming reduces sensor or communication quality. Direction finding locates transmitters. Spoofing creates false signals or identities. Cyber and electronic attack may occur together, making it difficult to determine whether a failure is physical, electromagnetic or digital.

High-power active sensors provide valuable information but reveal their presence. A survivable architecture should not require every radar and relay to transmit continuously. Emission schedules, remote sensors, passive observation and short data bursts can reduce exposure. Strategic radars may operate selectively while distributed local sensors maintain awareness.

Passive RF sensors detect emissions without transmitting. They may identify control links, navigation signals, radars or communication activity. Several passive bearings can estimate location. Their weakness is that autonomous or silent systems may provide little RF evidence, so they must be fused with radar, EO/IR and acoustic sensing.




Figure 15.1. Electronic warfare pressures and survivable responses.

EO/IR and acoustic sensors provide another passive layer. A camera can confirm a target without revealing itself through radio transmission. Acoustic arrays can indicate a bearing at short range. These systems are affected by weather, noise and line of sight, but their diversity complicates an attacker's attempt to blind the entire network with one method.

Directional point-to-point links can reduce unnecessary radiation outside the intended path, but they are not invisible. Side lobes, reflections, compromised endpoints and physical observation can reveal them. Frequency agility, alternate routes and careful power management should complement directionality.

The network should distinguish degradation from compromise. A jammed link may lose packets but remain trustworthy. A spoofed node may deliver plausible but false data. Health monitoring and cross-sensor checks are needed to decide whether to route around a node, reduce its confidence or isolate it entirely.

Electronic protection also includes timing and navigation. Sensors require accurate clocks and positions for fusion. Jamming or spoofing satellite navigation can create track errors even when communications remain available. Local time sources, inertial references, surveyed positions and cross-checks reduce this dependency.

Decoys and mobility protect valuable nodes. A decoy emitter can attract attack. Mobile radars and relays can change position between waves. Passive nodes can remain in place while active nodes move. The distributed layer creates more possible targets and makes it harder for the adversary to know which node is essential.

Electronic warfare should be included in the simulation and training programme. Operators must practise reduced update rates, false tracks, navigation error and complete loss of a regional link. The system should continue with a smaller but trustworthy picture rather than display precise-looking false information.

The aim of emission control is not silence at all times. It is disciplined use of the spectrum: transmit when the information gain justifies the exposure, remain passive when possible, and preserve alternate methods when the primary spectrum is contested.

Selected sources: NATO (2025); Monzon Baeza et al. (2025); MDA (2026a).

CHAPTER 16

Human Factors, Cognitive Load and Command Design

Why more information can make defence worse unless the system supports judgement

COMMAND PRINCIPLE

The objective is not to place a human in every technical loop. It is to keep human authority, judgement and accountability at every consequential decision point.

A dense sensor network can produce an impressive amount of data, but operators do not need more data; they need a clearer understanding of what requires attention. Poorly designed systems create alarm fatigue, automation bias and delayed decisions. Human factors are therefore part of air-defence performance.

The interface should present recognised tracks rather than raw feeds by default. Operators can open source details when needed. Colour, sound and priority should reflect threat and confidence, not merely detection. A low-confidence distant object should not compete visually with a confirmed threat approaching a protected site.

Workload should be distributed across levels. Local operators manage immediate sectors. Regional command reconciles tracks and allocates scarce resources. National command sets policy and priorities. Sending every decision upward creates delay; allowing every local unit to act independently creates fratricide and duplication.

The concept of mission command provides a balance. Higher command defines intent, protected priorities, identification rules and delegated authorities. Local cells act within those boundaries when time or communications prevent central approval. Decisions and evidence are recorded for later review.

Automation should be adjustable. In normal conditions, the system may request human confirmation for sensor tasking. During saturation, it may automatically re-task cameras or allocate non-lethal responses within policy. Lethal engagement remains subject to the required human authority and safety checks.

Training should include wrong recommendations. If operators see only correct AI outputs, they will learn blind trust. Exercises should inject false confidence, delayed tracks, friendly drones and compromised sensors. Operators must practise questioning the system without becoming unable to use it.

Crew coordination matters. Sensor operators, air-defence commanders, cyber defenders and weapon crews may interpret the same track differently. Common terminology and shared displays reduce misunderstanding. Voice procedures and fallback methods remain necessary when digital systems fail.

The system should measure cognitive performance. Decision time, missed alerts, false engagements, manual corrections and trust calibration can be recorded in simulation. These metrics may reveal that a technically accurate system still overwhelms its users.

Command design should also protect against political or organisational pressure. A system that rewards high interception counts may encourage unnecessary engagements. Performance should be measured against protected mission continuity, civilian safety and resource preservation.

The final human-factors principle is that the network should create decision advantage, not decision speed alone. A fast decision based on misunderstood data is not an advantage.

Human-Factors Design Requirements

Requirement	Purpose
Prioritised display	focus attention on high-consequence tracks
Visible confidence and sources	reduce automation bias
Delegated authority	preserve action during disconnection
Audit trail	support accountability and learning
Degraded-mode interface	prevent collapse when data is incomplete
Adversarial training	calibrate trust and expose failure modes
Workload metrics	measure whether operators can use the system

Selected sources: NATO (2025); U.S. Army (2026e).

CHAPTER 17

Defending Bases, Cities and Critical Infrastructure

A mosaic of local cells connected to regional command

SITE-DEFENCE PRINCIPLE

Air defence protects a mission, not only an airspace volume. Interception, passive protection, rapid repair and continuity planning must be designed together.

The distributed shield is most practical when organised around protected sites and corridors. Airfields, ports, power stations, command centres, rail hubs and urban areas have different geometry and risk. Each requires a local cell adapted to its terrain, civilian activity and available effectors.

An airfield cell may emphasise runway approaches, fuel storage, aircraft shelters and radar sites. A port cell may integrate coastal radar and maritime observation. A power-station cell may use fixed cameras and acoustic sensors around a smaller perimeter. An urban cell must manage dense civilian drones, buildings and restricted fields of fire.

Local cells should overlap. A drone leaving one sector is handed to the next. Regional command reconciles duplicate tracks and allocates high-tier support. The cells do not need identical equipment; they need common interfaces, identification rules and communications.




Figure 17.1. Protected sites organised as overlapping local cells connected to regional command.

Protected-site design should include passive defence. Camouflage, shelters, dispersion, fire suppression, redundant power and rapid repair may preserve the mission when interception

fails. Air defence cannot guarantee that no object reaches the site. Resilience of the protected function is part of the defensive architecture.

Urban defence requires strict governance. Civilian cameras and communications may offer useful observations, but access must be authorised, limited and audited. The system should avoid permanent collection unrelated to the emergency. False identification in populated areas has unacceptable consequences.

The local cell should maintain a site-specific library of friendly activity, authorised drone routes, terrain masks and danger areas. This context improves classification and prevents weapons from being cued into unsafe sectors.

Mobile nodes can reinforce threatened sites. A regional reserve may contain micro-radars, relays, interceptor drones and power modules. The network moves them according to warning and damage. This creates a defensive manoeuvre capability rather than fixed coverage alone.

The mosaic also supports deception. Some sites may transmit decoy activity, while real command nodes operate elsewhere. Sensors can be distributed beyond the protected perimeter so that attackers cannot assume the centre of the network is the centre of the site.

Regional command should avoid micromanaging every cell. Its most important functions are track reconciliation, scarce-resource allocation, airspace coordination and support. Local cells should handle immediate sensing and point defence.

The concept scales from one base to a national network. Standardised cells can be added where risk increases. The network grows by connecting local capability rather than by building one enormous central system.

Selected sources: NATO (2025); NATO ACT (2026a).

CHAPTER 18

External Validation: Policy and Experimentation in 2025-2026

How recent NATO and U.S. Army programmes support the direction of the research

VALIDATION JUDGEMENT

Recent official programmes support the direction of distributed sensing, common data, layered effectors and resilient C2. They validate the concept family, not every technical assumption of this book.

The original project was completed in 2023. Since then, official defence programmes have increasingly adopted the same architectural direction: many sensors, common data standards, distributed command cells, layered effectors and experimentation under realistic conditions.

NATO's 2025 integrated air and missile defence policy calls for resilient and redundant air command-and-control, including distributed nodes, resilient networks and backup solutions. It also describes layered defence across ranges and altitudes against aircraft, missiles and unmanned systems. This policy supports the book's argument that command resilience and layering are as important as individual interceptors.

NATO Allied Command Transformation launched the Layered Counter-UAS Initiative-X as a 2026 beacon project. Its stated objective is to connect sensors, command systems and response options into an integrated approach. The programme moves from individual cells toward a mosaic that can be deployed, connected and scaled.




Figure 18.1. Open-source policy and experimentation supporting the architectural direction.

The LCI-X Crucible events in Romania and Finland tested interoperability and connected counter-UAS teams across locations. NATO described the challenge as linking teams, sensors, command systems and response options, which closely resembles the distributed-cell architecture developed in this book.

The U.S. Army's Project Flytrap experiments focus on moving passive and active sensor data securely across coalition networks and testing defeat options. The U.S.-UK agreement on counter-drone technology standards likewise emphasises common data sharing and sensor fusion.

Exercise Golden Shield demonstrated an autonomous sensor on one platform detecting and classifying a hostile drone and transmitting data to a weapon system on another platform. This is a practical example of separating sensor and shooter while reducing the sensor-to-shooter timeline.

The Army's Warden event included multi-sensor tracking, terrain and signal-propagation analysis and simulated swarm attacks. Other Army work connects vehicle sensors to remotely operated weapons and explores formation-based layered protection.

These programmes do not prove that every element of the author's concept is correct. They do not validate point-to-point Wi-Fi as the preferred military waveform, the exact simulation assumptions or the proposed scale of 1,024 nodes. They do, however, independently validate the central architectural direction.

The external evidence also reinforces the importance of standards and experimentation. Modern counter-UAS capability is developing through repeated integration events rather than through one perfect system. The research roadmap in this book follows the same logic.

The author's contribution is therefore best understood as an early systems concept that connected communications, local sensors, command and remote control. The current book places that concept inside a broader, more rigorous framework of air-defence fusion, resilience, safety and war-gaming.

Selected sources: NATO (2025); NATO ACT (2026a); NATO ACT (2026b); U.S. Army (2026a); U.S. Army (2026b); U.S. Army (2026c); U.S. Army (2026f).

CHAPTER 19

Illustrative War-Gaming Method

A transparent simulation for comparing architectural behaviour rather than national systems

SIMULATION BOUNDARY

The model is conceptual and non-operational. It demonstrates relative architecture behaviour and is not a prediction of any real conflict or named weapon system.

The war-gaming model in this book is designed to test a research question: how do sparse premium and dense distributed architectures behave when attacks increase, communications degrade and nodes are lost? It is not intended to estimate the performance of any named national system. No classified specifications, geographic targets or operational plans are used.

The sparse model contains twelve major sensor nodes, three command-and-control hubs, a limited group of local effectors and a high-tier interceptor inventory. Its sensors have high individual detection probability, but its coverage and mission continuity depend heavily on the survival of key nodes.

The dense model contains 1,024 low-cost sensor nodes, sixty-four local fusion cells, multiple relays and a larger inventory of low-cost effectors. Individual sensors are less capable, but overlapping observations and local processing improve continuity. The dense model retains the same notional high-tier interceptor inventory because the purpose is to compare support architectures rather than eliminate premium systems.




Figure 14.1. The three-phase illustrative war-gaming design.

Three phases are simulated. The first represents reconnaissance and probing by fifty mixed low-signature threats. The second represents a saturation attack of 160 drones and decoys under jamming. The third represents severe attrition and a follow-on wave of 300 threats. Node loss and jamming increase across the phases.

The model measures detection probability, classification quality, fraction of threats engaged, fraction defeated and mission continuity. Mission continuity combines surviving connectivity, command resilience, local autonomy and repair. The model runs thousands of random trials so that results are expressed as distributions rather than single outcomes.

Several assumptions favour the sparse architecture. Its individual sensors begin with much higher detection probability and lower false-alarm rates. Several assumptions favour the dense architecture. It has overlapping local sensing, more local effectors, more command cells and faster reconstitution. The comparison therefore tests the trade between quality concentration and distributed resilience.

The model deliberately simplifies physics. It does not calculate radar cross section, detailed missile kinematics, terrain propagation, weather, human fatigue or exact weapon probability of kill. Those factors would be needed for a national-level operational model. The simplified model is appropriate for testing architectural behaviour.

The results should be interpreted comparatively. A value of 0.75 does not mean that a real system will defeat seventy-five percent of threats. It means that, under the declared assumptions, one architecture preserves more of a selected function than another.

Sensitivity analysis varies node loss from zero to fifty percent and jamming severity from zero to sixty percent. This reveals whether an architecture degrades gradually or crosses a sudden collapse threshold. A resilient design should maintain a useful mission even after significant local loss.

The war-game also identifies where the dense network can fail. If fusion is poor, false alarms can overwhelm command. If identities are weak, many nodes create cyber risk. If local effectors are absent, sensing does not solve magazine pressure. If command authority is unclear, decentralisation creates confusion. These failure conditions are as important as the positive results.

Notional Simulation Parameters

Parameter	Sparse premium model	Dense distributed model
Major sensor nodes	12	1,024 low-cost nodes
Command/fusion nodes	3	64 local fusion cells
Individual sensor quality	high	low to moderate
Overlap	limited	high
Local effector capacity	limited	large
High-tier inventory	held constant	held constant
Reconstitution speed	slow	rapid and modular
Primary risk	single-point loss and capacity	false alarms, cyber burden and maintenance

Selected sources: Author simulation (2026); NATO ACT (2026a); U.S. Army (2026c).

CHAPTER 20

Simulation Results and Operational Interpretation

What the model suggests about awareness, saturation tolerance and graceful degradation

KEY RESULT

The dense network does not make the defender invulnerable. It changes the failure mode from sudden blindness toward gradual degradation and faster recovery.

The simulation produces a consistent pattern. Under limited pressure, both architectures maintain useful detection, but the dense model benefits from overlapping local observations. Under heavy saturation, the sparse architecture loses engagement capacity rapidly because a small number of nodes and effectors carry most of the burden. The dense architecture preserves more local tracks and distributes engagement across more cells.

Detection probability remains relatively strong in the dense model because multiple low-quality sensors contribute to the same track. This demonstrates a fundamental property of fusion: several independent imperfect observations can produce a useful result. The benefit disappears if the observations are correlated, poorly timed or compromised, which is why data quality remains essential.

The largest difference appears in mission continuity. As node loss and jamming increase, the sparse model approaches a threshold at which key command or sensing functions become unavailable. The dense model degrades more gradually because alternate paths, local fusion and local autonomy remain.




Figure 15.1. Illustrative Monte Carlo results across limited, heavy and severe conditions.

The dense model does not achieve perfect defence. In the severe phase, total threat volume exceeds engagement capacity. Some threats penetrate. This is an important finding because distributed sensing cannot create unlimited weapons. The architecture improves endurance and decision quality, but effectors, ammunition and protected-site design remain necessary.

The sparse architecture retains advantages in long-range track quality and management simplicity. Its lower number of nodes reduces maintenance and cyber complexity. A hybrid design should therefore preserve premium sensors and command nodes while distributing local functions and alternate pathways.

Sensitivity analysis shows that jamming and node loss interact. A sparse network can tolerate one form of disruption but may collapse when both occur together. A dense network retains more continuity, although its performance declines when local fusion cells and relays are simultaneously affected.

The results support the author's prediction about high-end systems in swarm warfare, but with an important qualification. A high-end interceptor may continue to perform correctly while the defended system fails because it lacks tracks, command capacity, reloads or economical responses. The system failure is architectural.




Figure 15.2. Mission continuity under increasing node loss and jamming.

The model also demonstrates that thousands of nodes are not automatically superior. When false-alarm rates are increased or fusion quality is reduced, operator workload rises and engagement decisions slow. The architecture therefore requires strong data standards, confidence scoring and automated filtering.

A further result is the importance of reconstitution. The dense model recovers because failed nodes can be replaced and routes can be rebuilt. This suggests that logistics should stock sensors, relays, power modules and pre-configured identities, not only interceptors.

The practical conclusion is that air-defence effectiveness should be measured by defended mission continuity over time. Interception percentage is useful but incomplete. A force that defeats most of the first wave and becomes blind may be less effective than a force that accepts limited penetration but continues defending through repeated attacks.

Interpretation of Results

Observed pattern	Operational meaning
Dense detection remains higher under attrition	overlapping local sensing preserves warning
Sparse engagement capacity falls quickly	few effectors and command nodes become bottlenecks
Dense continuity degrades gradually	alternate routes and local autonomy preserve function
Both models suffer under extreme volume	sensing cannot replace magazine depth
Dense false-alarm burden is higher	fusion quality and confidence scoring are decisive
Reconstitution changes later waves	spares, power and identity management are combat capabilities

Selected sources: Author simulation (2026).

Deployment and Implementation Roadmap

How to move from a university proof of concept to safe, repeatable operational experimentation

NEXT RESEARCH STEP

Build a small multi-sensor cell with declared interfaces and measured uncertainty, then attack it with jamming, false data, node loss and operator overload before attempting scale.

The research should advance through controlled phases. The first phase establishes a clean digital baseline: network segmentation, device identity, time synchronisation, logging and a common data model. Without this foundation, adding sensors will create more confusion rather than more awareness.

The second phase integrates a small number of sensor types in a laboratory environment. A micro-radar, EO/IR camera, passive RF receiver and acoustic sensor should report through a common interface. The objective is not to defeat a threat but to measure latency, time error, false alarms and correlation quality.

The third phase introduces local fusion and edge processing. Raw video should be reduced to alerts or selected clips. Track confidence should be displayed. Operators should be able to inspect source provenance. The system should continue operating after the wide-area link is disconnected.




Figure 16.1. A phased implementation and test roadmap.

The fourth phase integrates an effector gateway in a non-lethal environment. The gateway may cue a camera, spotlight, training simulator or inert weapon mount. Safety zones, authority and abort functions are tested. This phase validates the separation between track information and engagement authority.

The fifth phase is adversarial. Red teams jam links, inject false reports, compromise test nodes, remove power and destroy simulated relays. The system is evaluated on detection, isolation, graceful degradation and recovery. A network that works only in benign conditions has not demonstrated resilience.

The sixth phase scales through a digital twin. Hundreds or thousands of virtual nodes are combined with real hardware. Different terrain, weather, traffic and swarm conditions are tested. The simulation model in this book becomes part of a larger test environment rather than a final answer.

Measures of performance should be declared before testing. These include track latency, position error, false-alarm rate, identification confidence, network availability, time to isolate compromise, operator workload, engagement decision time, reconstitution speed and cost per defended hour.

Open interfaces are essential. The project should define sensor, track, health, tasking and effector messages. Vendors can implement different internal technologies while remaining interoperable. This prevents dependence on a single supplier and allows legacy equipment to participate.

Training should grow with the technology. Operators must understand sensor limitations, confidence, cyber warnings and degraded modes. Commanders must practise delegated authority and communication loss. Maintainers must replace nodes and restore identities quickly.

The roadmap should culminate in a defensible field experiment, not a marketing demonstration. Assumptions, failures and negative results should be documented. The research becomes credible when it can explain not only how the system succeeds, but also the conditions under which it fails.

Core Test Metrics

Category	Metric examples
Sensing	detection probability, false alarms, track initiation time
Fusion	association accuracy, confidence calibration, track continuity
Network	latency, packet loss, route recovery and priority traffic delivery
Security	time to detect, isolate and recover from compromise
Command	operator workload, decision time and authority compliance
Effectors	correct cueing, safety checks and magazine management
Logistics	node replacement time, power endurance and calibration burden
Economics	cost per defended hour and cost per successful response

Selected sources: NATO ACT (2026a); U.S. Army (2026c); U.S. Army (2026f).

CHAPTER 22

Economics, Logistics, Training and Governance

Why distributed technology succeeds only when the institution can sustain it

ECONOMIC FINDING

The distributed architecture changes the cost curve only when nodes are standardised, replaceable and governed. Quantity without support creates a maintenance crisis rather than resilience.

A distributed network appears inexpensive because individual sensors and relays can be cheap. The total system cost includes installation, power, communications, software, calibration, spares, cybersecurity, training and replacement. Economic analysis must therefore consider lifecycle cost rather than unit price.

Standardisation is the main source of savings. Common mounts, power connectors, network interfaces and software reduce training and spares. A small number of node families is easier to support than hundreds of unique devices. Modularity allows damaged sensors to be replaced without replacing the entire access point or communications package.

Power is a strategic constraint. Remote nodes need batteries, generators, solar charging or connection to protected infrastructure. Power status should appear in the network health picture. A sensor that is technically operational but has thirty minutes of power remaining is a planning problem.

Calibration is another hidden burden. Radar alignment, camera orientation, acoustic calibration and clock accuracy affect fusion. The system should support automated self-test and comparison with neighbouring nodes. Maintenance schedules should be risk-based rather than identical for every device.

Personnel remain decisive. Operators need to understand uncertainty and automation. Maintainers need networking, sensor and cybersecurity skills. Commanders need to practise partial information and delegated authority. Training must include false alarms, friendly drones, civilian activity and communications loss.

Governance is especially important when civilian-style cameras or commercial components are used. Data access, retention, privacy, legal authority and audit must be defined. Emergency integration of civilian sensors should not become permanent uncontrolled surveillance.

Procurement should reward interoperability and upgradeability. A cheap proprietary sensor can become expensive if it requires a unique command system. Contracts should require documented interfaces, signed software, vulnerability disclosure, lifecycle support and the ability to operate in disconnected modes.

Supply-chain diversity improves resilience but complicates assurance. Components from multiple sources reduce dependence on one supplier, yet each source introduces security and quality risks. The system needs component inventory, testing and trusted configuration baselines.

The economics of defence also involve opportunity cost. Every premium interceptor used against a low-cost drone may be unavailable for a higher-consequence threat. Every operator assigned to a manual camera is unavailable elsewhere. Fusion and layered effectors aim to allocate scarce resources according to threat value.

The institutional objective is sustainable readiness. A system that performs impressively during one demonstration but cannot be maintained, updated or trained at scale is not a defence capability. The distributed shield must be designed as an organisation, supply chain and doctrine as much as a network.

Lifecycle Cost Categories

Category	Examples
Acquisition	sensors, relays, servers, gateways and effectors
Installation	masts, shelters, fibre, power and site preparation
Software	fusion, command applications, licences and updates
Security	identity, monitoring, red teaming and incident response
Personnel	operators, maintainers, analysts and trainers
Sustainment	spares, calibration, batteries and replacement
Training	simulation, field exercises and certification
Governance	data protection, audit, safety and legal compliance

Selected sources: Harrell (2026); NATO ACT (2026a); U.S. Army (2026c).

Findings, Limitations and Future Research

The final assessment of the distributed shield concept

FINAL RESEARCH STATEMENT

A few high-end systems can provide extraordinary performance, but a defence that depends on them alone can be saturated or paralysed. A dense, secure and intelligently fused lower layer converts expensive systems from isolated points of failure into protected members of a resilient network.

The first finding is that the original point-to-point communication research contains a valid architectural insight. Communications, surveillance, sensors and control should be designed as one system. Directional links and local access points can provide a useful high-bandwidth layer inside a wider hybrid network.

The second finding is that the architecture can be extended into a distributed lower-layer air-defence network. Cameras, IoT sensors, micro-radars, passive RF receivers, acoustic sensors and field observers can contribute to a recognised picture when their reports are time-aligned, correlated and confidence-scored.

The third finding is that high-end air and missile defence systems remain essential but should not operate as isolated islands. S-400, THAAD, Aegis/Standard Missile and Iron Dome represent different roles. Their common vulnerability in swarm-era warfare is architectural saturation: too many tracks, insufficient local coverage, finite magazines, expensive interceptors and concentrated command dependencies.




Figure 18.1. The distributed shield as a combined communications, sensing, command and response architecture.

The fourth finding is that thousands of low-cost nodes can improve persistence, redundancy and local awareness, but they create cyber, maintenance and false-alarm burdens. Their value depends on zero-trust identity, segmentation, edge fusion, open standards and rapid replacement.

The fifth finding is that legacy anti-aircraft guns and other low-cost effectors can regain relevance through digital cueing. They should remain under human control and strict safety constraints. Their role is to absorb appropriate low-tier threats and preserve premium interceptors.

The sixth finding is that artificial intelligence is most credible as decision support. Track association, anomaly detection, sensor tasking and confidence scoring are valuable. Emotion detection and opaque lethal autonomy are not defensible foundations for military decision-making.

The seventh finding is that the central performance metric should be mission continuity over time. A system that intercepts a high percentage in the first wave but becomes blind or empty is not necessarily effective. Reconstitution, reloads, power, spares and local autonomy are combat functions.

The simulation supports the architecture but does not validate real system performance. It lacks detailed terrain, radar physics, electronic warfare models, human behaviour and national inventories. Future work should build a digital twin with real sensor data and hardware-in-the-loop testing.

Future research should also examine post-quantum cryptography, time synchronisation under attack, Byzantine-resilient fusion, explainable AI, coalition data standards, urban sensor governance and the economics of directed energy and interceptor drones.

The concept's strongest external validation comes from recent official developments. NATO's 2025 IAMD policy emphasises resilient distributed command-and-control. NATO's 2026 LCI-X programme seeks a mosaic of connected counter-UAS cells. U.S. Army experiments connect sensors on one platform to effectors on another and test multi-sensor fusion against simulated swarms. These programmes do not prove the author's complete design, but they confirm the direction of the central argument.

The final judgement is that the future air-defence network will be neither purely centralised nor purely distributed. It will combine high-tier reach with low-tier mass, national policy with local initiative, automated fusion with human authority, and military systems with carefully governed civil infrastructure. The decisive advantage will belong to the force that can keep sensing, sharing, deciding and acting after its first plan and first nodes have failed.

The distributed shield is therefore not a single product. It is a design philosophy: many sensors, shared standards, trusted data, layered effectors, resilient command and continuous adaptation.

Selected sources: Ashraf (2023); NATO (2025); NATO ACT (2026a); U.S. Army (2026a); U.S. Army (2026c).

APPENDIX A


Simulation Assumptions and Formulae

The illustrative simulation compares architectural behaviour, not real national systems. It uses a Monte Carlo process in which sensor survival, jamming effect, classification quality, engagement capacity and mission continuity vary across repeated trials.

Sparse model: twelve major sensors, three command nodes, limited local effectors and a shared high-tier inventory. Dense model: 1,024 low-cost nodes, sixty-four local fusion cells, many local effectors and the same notional high-tier inventory.

Detection probability is approximated from individual sensor probability and effective overlap. Mission continuity combines surviving nodes, command resilience, jamming degradation, local autonomy and repair. These formulae are conceptual and intentionally avoid real system parameters.

The model should be extended through terrain-aware propagation, real sensor data, hardware-in-the-loop testing, operator workload measurement and adversarial cyber/electronic warfare.



Appendix Figure A.1. War-gaming design and declared boundaries.

Simulation Variables

Variable	Meaning
Sensor survival	fraction of sensor nodes remaining available
Jamming severity	reduction in communications and sensing effectiveness
Effective overlap	number of independent observations contributing to a track
Classification quality	probability that a detected object is correctly categorised
Engagement capacity	number of threats that can be assigned an available response
Mission continuity	combined measure of sensing, C2, autonomy and repair

APPENDIX B

Reference Architecture and Message Types

The reference architecture separates observations, recognised tracks, command decisions, effector tasking and health data. This prevents a raw sensor from directly controlling a weapon and allows old and new systems to interoperate through gateways.

Every message should include identity, timestamp, position or reference frame, confidence, source, security label and integrity protection. Nodes should receive only the message types required for their function.


Local fusion cells may operate independently with cached policy during disconnection. Regional nodes reconcile tracks and resources. National command defines authority, identification policy and priorities.

Core Message Types

Message	Required fields	Purpose
Observation	sensor ID, time, location, uncertainty and measurement	raw sensor report
Recognised track	track ID, position, velocity, confidence and source history	shared air picture
Sensor tasking	target area, priority, duration and authority	re-task radar, camera or observer
Effector cue	track, engagement window, limits and command authority	support supervised response
Node health	power, link, time quality, software and calibration	resilience and maintenance
Security event	identity, anomaly, severity and containment action	cyber defence and audit

From the 2023 Research Prototype to a Distributed Defence Architecture

The original point-to-point network becomes the communications spine for sensing, command and local response



Appendix Figure B.1. Reference architecture derived from the original research.

APPENDIX C

Glossary

AEW&C: Airborne Early Warning and Control.

AI: Artificial intelligence; in this book, primarily decision support and data processing.

C2: Command and control.

C-UAS: Counter-uncrewed aircraft systems.

EO/IR: Electro-optical and infrared sensing.

Fusion cell: A local or regional node that correlates sensor reports into tracks.

IAMD: Integrated air and missile defence.

IoT: Internet of Things; networked sensors and devices.

Legacy effector: An existing weapon integrated through a digital cueing gateway.

Lower-layer air picture: The local low-altitude situational picture around terrain and protected sites.

Mission continuity: The ability to preserve the defended function through attack and degradation.

Node: A sensor, relay, command, processing or effector component connected to the network.

Point-to-point link: A directional communications link between known endpoints.

Recognised track: A correlated representation of one object based on one or more observations.

ROE: Rules of engagement.

Saturation: An attack designed to exceed sensing, command, engagement or reload capacity.

Zero trust: A security model requiring continuous verification and least privilege.

APPENDIX D

Bibliography and Image Credits

Bibliography

- Ashraf, Muhammad Sohail. Military Communication and Surveillance System Using Point-to-Point Wi-Fi Devices. MSc Advanced Computer Science Project, University of Hertfordshire, 2023.
- Farlik, Jan, Miroslav Kratky, Josef Casar, and Vadim Stary. 'Multispectral Detection of Commercial Unmanned Aerial Vehicles.' *Sensors* 19, no. 7 (2019): 1517. <https://doi.org/10.3390/s19071517>.
- Harrell, John. 'David's Sling, Goliath's Bill: The Economics of Drone Defense - A Layered Defense Framework.' *Cyber Defense Review* 11, no. 1 (2026).
- Khawaja, Wahab, et al. 'Threats from and Countermeasures for Unmanned Aerial and Underwater Vehicles.' *Sensors* 22, no. 10 (2022): 3896. <https://doi.org/10.3390/s22103896>.
- Missile Defense Agency. 'Terminal High Altitude Area Defense (THAAD).' U.S. Department of Defense, accessed August 2026. <https://www.mda.mil/system/thaad.html>.
- Missile Defense Agency. 'Sensors: AN/TPY-2.' U.S. Department of Defense, accessed August 2026. <https://www.mda.mil/system/sensors.html>.
- Monzon Baeza, Victor, Raul Parada, Laura Concha Salor, and Carlos Monzo. 'AI-Driven Tactical Communications and Networking for Defense: A Survey and Emerging Trends.' *arXiv:2504.05071* (2025).
- NATO. 'NATO Integrated Air and Missile Defence Policy.' 13 February 2025. <https://www.nato.int/en/about-us/official-texts-and-resources/official-texts/2025/02/13/nato-integrated-air-and-missile-defence-policy>.
- NATO Allied Command Transformation. 'Layered Counter-UAS Initiative.' 2026. <https://www.act.nato.int/activities/lci-x/>.
- NATO Allied Command Transformation. 'Layered Counter-UAS Initiative Is Building NATO's Approach to a Fast-Moving Threat.' 18 May 2026.
- NATO Allied Command Transformation. 'NATO Advances Counter-UAS Integration Through LCI-X Crucible 2-26 in Finland.' 10 June 2026.
- Park, Seongjoon, Hyeong Tae Kim, Sangmin Lee, Hyeontae Joo, and Hwangnam Kim. 'Survey on Anti-Drone Systems: Components, Designs, and Challenges.' *IEEE Access* 9 (2021). <https://doi.org/10.1109/ACCESS.2021.3065926>.
- Rojhani, Neda, and George Shaker. 'Comprehensive Review: Effectiveness of MIMO and Beamforming Technologies in Detecting Low RCS UAVs.' *Remote Sensing* 16, no. 6 (2024): 1016. <https://doi.org/10.3390/rs16061016>.
- Rosoboronexport. S-400 Triumf Air Defence Missile System. Official product material, 2024. https://roe.ru/pdfs/pdf_6221.pdf.
- RTX. 'Iron Dome System and SkyHunter Missile.' Accessed August 2026. <https://www.rtx.com/raytheon/what-we-do/integrated-air-and-missile-defense/irondome>.
- Smailov, Nurzhigit, et al. 'Review of AI-Augmented Multisensor Architectures for Detecting and Neutralizing UAV Threats.' *International Journal of Innovative Research and Scientific Studies* 8, no. 5 (2025).
- U.S. Army. 'Army Enhances C-UAS Data Flow and Interoperability during Project Flytrap 4.5.' 9 January 2026.
- U.S. Army. 'US, UK Set Joint Counter-Drone Technology Standards.' 12 March 2026.
- U.S. Army. '1st Cavalry Division Tests Golden Shield Counter-Drone System.' April 2026.
- U.S. Army. 'Army Armaments Center Develops New Counter-UAS Capability.' June 2026.
- U.S. Army. 'Joint Interagency Task Force 401, Paratroopers Test New Counter-UAS.' April 2026.
- U.S. Army. 'Army-Industry Collaboration Advances Counter-UAS Technology.' 22 July 2026.
- U.S. Navy. 'AEGIS Weapon System.' Fact File, 20 September 2021. <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2166739/aegis-weapon-system/>.
- U.S. Navy. 'Standard Missile.' Fact File, 31 August 2021. <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2169011/standard-missile/standard-missile/>.

Image Credits

Original research images are reproduced from Muhammad Sohail Ashraf's 2023 University of Hertfordshire project.

All vector diagrams, simulation charts and the cover were created specifically for this book from the author's concepts and the stated open-source research base.

Visual Plate 6A, Visual Plate 9A and Visual Plate 13A are conceptual infographics prepared for this project. They are educational representations and do not depict classified architectures.

System descriptions are based on official public sources listed in the bibliography. No official logos imply endorsement.