Journal of Science Research and Reviews
PRINT ISSN: 1595-9074
E-ISSN: 1595-8329
DOI: https://doi.org/10.70882/josrar.2026.v3i4.272
Homepage: https://josrar.esrgngr.org
Original Research Article

Wireless Sensor Network-Based Border Security Using A Three-Tiered Hierarchical Surveillance Strategy: Architecture, Security and Deployment Considerations

[1]Ojoawo Akinwale Olusola, [1]Olorunisola Abiodun Emmanuel, [1]Idowu Mayowa Opakunle, *[2]Alabi Adewale Abayomi and [3]Adeaga I. I.

[1]Department of Computer Science, Adeseun Ogundoyin Polytechnic, Eruwa, Oyo State, Nigeria
[2]Department of Electrical and Electronic Engineering, Adeseun Ogundoyin Polytechnic, Eruwa, Oyo State, Nigeria
[3]Department of Computer Engineering, The Polytechnic, Ibadan, Oyo State, Nigeria
*Corresponding Author's email: adewalealabi01@gmail.com

KEYWORDS
Wireless Sensor Network,
Border Surveillance,
Hierarchical Architecture,
Edge Artificial Intelligence,
Raspberry Pi,
MICAz,
Intrusion Detection,
Cybersecurity.

ABSTRACT
Border regions in Nigeria and the Sub-Saharan African and Sahelian security environment are characterised by expansive terrain, sparse state presence, transnational criminal mobility, irregular telecommunications coverage and a need for timely intelligence. Patrol-centred surveillance remains indispensable, but its reach, endurance and responsiveness are limited across hostile terrain. This design-oriented integrative study proposes a wireless sensor network (WSN)-based border-security system organised through a three-tiered hierarchical surveillance strategy. A synthesis of peer-reviewed studies, routing research, IoT security standards and technical documentation was used to derive operational requirements, identify limitations and formulate a reference architecture. Tier 1 comprises low-power MICAz-class nodes and multimodal sensors for motion, acoustic, seismic, vibration and environmental events. Tier 2 uses Raspberry Pi-based motorised camera nodes for event-triggered image acquisition and lightweight object classification. Tier 3 integrates gateways, secure messaging, geospatial dashboards, incident records and human-authorised response. The design separates continuous low-energy detection from computationally intensive visual confirmation, thereby reducing camera activation, bandwidth use and false-alarm burden. It also incorporates energy-aware clustering, location-informed routing, message buffering, mutual authentication, encrypted transport, signed firmware, network segmentation and degraded-mode operation. Rather than claiming unverified field performance, the paper presents a multidimensional evaluation framework and scenario-based validation model covering detection quality, latency, packet delivery, energy consumption, cybersecurity, maintainability, operator workload and lifecycle cost. The synthesis indicates that sustainable border surveillance depends on coordinated sensing, networking, edge inference, platform services, governance and maintenance. The architecture provides a basis for phased experimentation along national borders, state boundaries, forest reserves and sparsely monitored security corridors.

CITATION
Ojoawo, A. O., Olorunisola, A. E., Idowu, M. O., Alabi, A. A., & Adeaga, I. I. (2026). Wireless Sensor Network-Based Border Security Using A Three-Tiered Hierarchical Surveillance Strategy: Architecture, Security and Deployment Considerations. *Journal of Science Research and Reviews*, *3*(3), 145-163. https://doi.org/10.70882/josrar.2026.v3i4.272

INTRODUCTION

The contemporary security landscape of West Africa and the Sahel is shaped by the interaction of violent extremism, armed banditry, kidnapping, illicit trafficking, irregular migration and organised criminal mobility. These threats do not respect administrative boundaries; they exploit porous frontiers, forest corridors, poorly illuminated access routes, seasonal tracks and locations where public authority is episodic. The resulting security problem is not merely the physical existence of long borders, but the absence of persistent, trustworthy and rapidly interpretable information concerning what is occurring along them. Recent regional assessments continue to associate Sahelian instability with the geographic expansion of militant organisations and transnational criminal markets, while studies of north-western Nigeria describe highly mobile armed groups that operate from forested and difficult-to-police environments (UNODC, 2024; UNIDIR, 2024; Africa Center for Strategic Studies, 2025).

Border security encompasses the coordinated protection of territorial boundaries against unauthorised movement of persons, weapons, contraband, vehicles and other threats, while preserving lawful trade, travel and community interaction. Border surveillance is the information-producing component of that broader mandate. Its function is to detect anomalous activities, classify their probable significance, estimate their location and movement, communicate evidence to authorised decision makers and support proportionate intervention. A surveillance system is therefore not equivalent to a camera, fence or radar. It is a socio-technical arrangement in which sensors, communications, computing, operating procedures, personnel and legal authority must function as a coherent whole.

In many African environments, surveillance still depends heavily on patrol teams, static observation posts, community reporting and periodic aerial reconnaissance. These mechanisms provide human judgement and contextual knowledge that cannot be completely automated. Nevertheless, they are expensive to sustain continuously across extensive terrain and may expose personnel to ambush, environmental hazards and fatigue. The operational difficulty becomes particularly acute where roads are poor, network coverage is intermittent, electricity is unreliable and the monitored corridor contains vegetation, uneven topography or seasonal flooding. Under such conditions, the central design question is how to extend the perceptual reach of security personnel without constructing an economically prohibitive or maintenance-intensive infrastructure.

Wireless sensor networks provide a promising response because they distribute small sensing and communication devices over an area of interest and permit them to cooperate in reporting physical phenomena. A typical node contains one or more sensors, a microcontroller, memory, a radio transceiver and an energy source. Nodes can be deployed in ad hoc or planned patterns, form multi-hop routes and send observations to a sink or gateway. Their historical association with military surveillance derives from precisely those characteristics: unattended operation, spatial redundancy, self-organisation and the ability to function where permanent wired infrastructure is impracticable (Akyildiz et al., 2002; Hammoudeh et al., 2017).

WSNs are, however, constrained systems. Radio communication consumes scarce energy; processors and memory are limited; wireless links are lossy; nodes can be captured or damaged; and environmental events can resemble genuine intrusions. A stand-alone motion sensor can therefore produce an abundance of ambiguous alerts, while a continuously active camera can exhaust power and bandwidth. Conversely, a high-capacity camera placed only at a command post cannot observe events beyond its line of sight. These limitations imply that no single sensor family, communication protocol or machine-learning model can provide reliable border awareness by itself.

The proposed solution addresses this problem through functional hierarchy. The first tier undertakes persistent, low-power detection in the remote monitored zone. The second tier provides event-triggered visual confirmation and edge inference in a semi-monitored zone. The third tier performs secure aggregation, geospatial presentation, human validation and response coordination in a protected command environment. The hierarchy is intended to conserve energy, distribute computational work according to device capability and progressively increase the evidential quality of an alert before it reaches operators.

The intellectual contribution of this study is fourfold. First, it reformulates border surveillance as a sequence of coordinated sensing, confirmation, fusion, decision and response responsibilities. Second, it develops a three-tier reference architecture that integrates resource-constrained WSN nodes with edge-computing cameras and conventional command infrastructure. Third, it incorporates security, resilience and maintenance as design requirements rather than post-deployment additions. Fourth, it proposes a transparent evaluation framework that avoids reducing system performance to detection accuracy alone.

Accordingly, the study seeks to answer the following questions: What technical and operational requirements should guide a WSN-enabled border-surveillance service? How can heterogeneous sensing and edge-computing resources be organised to reduce energy consumption and false alarms? Which communication, routing and security controls are appropriate across the three tiers? What measures should be used to evaluate the system under realistic African deployment conditions? The remainder of

the paper describes the design method, synthesises the evidence base, presents the architecture and algorithms, evaluates representative scenarios and identifies implementation priorities and research gaps.

Wireless Sensor Networks as a Surveillance Substrate

A WSN converts geographically dispersed physical events into networked digital observations. Its basic architecture contains sensing nodes, one or more sinks, communication links and an application that interprets the collected data. Depending on the application, the nodes may measure temperature, light, acoustic pressure, magnetic disturbance, vibration, acceleration, infrared radiation or chemical properties. In surveillance, the most useful capability is not the measurement of an isolated variable, but the temporal and spatial correlation of several variables that collectively indicate the passage of a person, animal or vehicle.

The network must simultaneously satisfy coverage and connectivity. Coverage concerns whether the sensing field is observed with sufficient probability, whereas connectivity concerns whether an observation can reach the sink. A deployment can have excellent geometric coverage and still fail operationally because nodes cannot establish reliable routes. Border-monitoring research therefore treats node density, belt geometry, radio range, sensing range and routing together. Hammoudeh et al. (2017), for example, formulated deployment and routing considerations for topologically linear border applications and demonstrated that protocols designed around the geometry of the monitored belt can outperform general-purpose ad hoc routing under relevant conditions.

Energy is the dominant constraint because replacing batteries in hostile or inaccessible terrain may be dangerous, expensive or impossible. The radio should remain asleep whenever transmission is unnecessary, sensing should be duty-cycled, and data should be aggregated before long-range communication. Hierarchical protocols such as LEACH distribute communication burden by rotating cluster-head roles, while geographic and real-time protocols such as GEAR and SPEED address location-informed forwarding and bounded delivery delay (Heinzelman et al., 2000; He et al., 2003). Yet protocol selection cannot be universal: a system that prioritises maximum lifetime may tolerate delays that are unacceptable in a time-critical intrusion event.

The practical implication is that the surveillance architecture must distinguish routine status traffic from urgent event traffic. Battery level, link quality and periodic health reports can tolerate delay and aggregation; a probable human intrusion requires immediate transmission, acknowledgement and escalation. This distinction is embedded in the proposed message classes and quality-of-service rules.

Existing border-surveillance proposals exhibit several recurring weaknesses. Some assume continuous Internet connectivity or mains power in locations where neither is dependable. Others concentrate on prototype detection without explaining maintenance, key management, gateway failure or the procedure for reconciling contradictory sensor and camera evidence. Several systems stream video continuously, thereby imposing bandwidth and energy requirements that are incompatible with remote deployment. In addition, security is frequently represented as generic encryption without attention to device identity, physical capture, credential rotation, signed updates, anti-replay protection or operator accountability.

False alarms constitute an equally important operational problem. Wind-blown vegetation, livestock, wildlife, rainfall, thermal variation and routine community movement may trigger single-modality sensors. Excessive alert volume produces alarm fatigue, weakens operator trust and can cause genuine incidents to be overlooked. Visual confirmation can reduce this burden, but computer vision itself is affected by darkness, occlusion, dust, camera movement and unfamiliar objects. Reliable classification therefore requires evidence fusion and explicit uncertainty rather than an unqualified binary assertion.

The proposed hierarchy is designed around these constraints. Low-power nodes remain the persistent first observers; more capable edge devices are activated selectively; the command tier receives a richer evidence package rather than raw sensor noise; and operators retain authority over consequential response. This arrangement is not presented as a replacement for border personnel. It is an intelligence-support system intended to expand situational awareness, reduce unnecessary patrol exposure and improve the speed with which limited resources can be directed to probable incidents.

MATERIALS AND METHODS

Research Design

The study adopted a design-science orientation supported by an integrative technical review. Design science was appropriate because the principal output is an artefact: a reference architecture, operating logic and evaluation model intended to address a defined surveillance problem. The integrative component permitted the combination of evidence from WSN deployment studies, border-surveillance prototypes, routing protocols, embedded computing, machine vision, IoT security standards and device documentation. A statistical meta-analysis was not attempted because the retained studies employed heterogeneous terrains, node counts, communication technologies, threat models and performance measures.

The method proceeded through five stages: evidence identification, eligibility appraisal, requirement synthesis, architectural design and scenario-based validation. Figure 1 illustrates the procedure. The objective was not to claim exhaustive retrieval, but to establish a traceable connection between recurring problems in the literature and the responsibilities assigned to each tier of the proposed system.

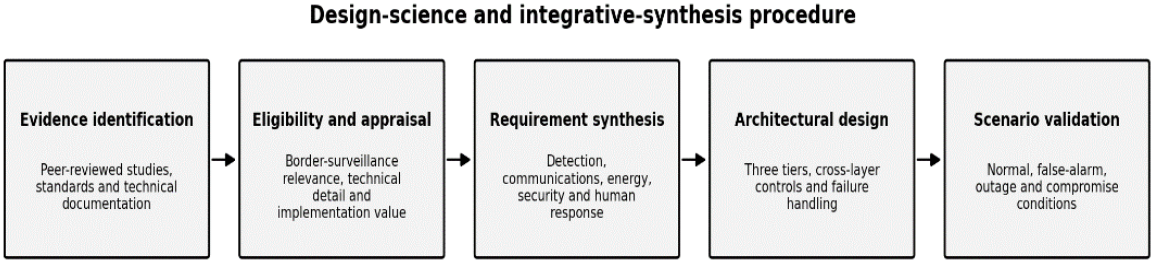

Figure 1: Design-science and integrative-synthesis procedure

Evidence Identification and Eligibility

Targeted searches were conducted across IEEE Xplore, SpringerLink, ScienceDirect, Google Scholar, institutional repositories and standards organisations. Search terms combined wireless sensor network, border surveillance, intrusion detection, perimeter monitoring, energy-efficient routing, LEACH, GEAR, SPEED, AODV, LoRaWAN, ZigBee, MQTT, edge computing, Raspberry Pi, Tiny YOLO, IoT cybersecurity and secure firmware. Reference lists of relevant publications were examined to identify foundational routing and architecture studies.

A source was retained when it addressed at least one of the following: sensor deployment for border or perimeter monitoring; routing and network lifetime in resource-constrained WSNs; heterogeneous sensor and camera integration; embedded or edge-based object detection; secure IoT communication; or operational evaluation of surveillance systems. Purely promotional materials, papers without sufficient technical detail, duplicated descriptions and claims unsupported by a discernible method were excluded. Manufacturer documentation and recognised standards were retained only where they informed hardware capabilities, messaging or cybersecurity controls.

The resulting evidence base contained thirty-six sources, including peer-reviewed journal and conference publications, standards and authoritative technical documentation. The sources span foundational WSN research and contemporary work on edge intelligence, energy management and IoT security. Because the article develops a design rather than reports a completed field trial, evidence was used to justify architectural decisions and evaluation criteria, not to manufacture empirical outcomes.

Data Extraction and Architectural Synthesis

For each source, information was extracted on surveillance context, sensing modality, node platform, communication technology, routing strategy, processing location, security provisions, reported metrics and limitations. The evidence was coded into eight themes: operational requirements; coverage and deployment; sensing and classification; communications and routing; energy management; edge and cloud processing; cybersecurity and resilience; and evaluation and maintainability.

Architectural synthesis separated stable responsibilities from replaceable products. For example, the requirement for event-triggered visual confirmation is stable, whereas the specific camera, single-board computer or neural network may change. Likewise, authenticated low-bandwidth messaging is a stable responsibility, while the deployment may use LoRaWAN, IEEE 802.15.4, private cellular or another appropriate transport. This separation prevents the system from becoming obsolete when a particular device reaches end of support.

Design Assumptions, Scope and Ethical Boundary

The design assumes that the monitored corridor can be partitioned into remote, semi-monitored and protected zones; that at least one gateway or relay path can connect each operational sector to a command facility; and that authorised personnel are available to interpret critical alerts. Terrain is not assumed to be uniformly flat. Rather, the deployment method requires a site survey that accounts for vegetation, elevation, water channels, settlements, animal routes and radio obstruction.

The system is limited to detection, classification, tracking support and communication. It does not include autonomous weapon activation or unreviewed coercive action. Consequential intervention remains subject to human authorisation, applicable law, rules of engagement and institutional accountability. This boundary is technically and ethically necessary because machine classification can be wrong, sensor identities can be spoofed and legitimate civilians may move within border communities.

The study also recognises the privacy implications of identifiable imagery and movement data. Data minimisation, restricted access, defined retention and

auditable use are therefore included within the security and governance layer. The system should collect only the evidence needed to support a legitimate security purpose and should not transform routine border management into indiscriminate population surveillance.

Design-Level Evaluation

Validation was performed at design level through requirement tracing and scenario analysis. Each architectural responsibility was mapped to one or more operational risks, and representative scenarios were used to test whether the proposed workflow could respond coherently to human movement, animal activity, vehicle movement, communication loss and suspected node compromise. The analysis does not substitute for a field experiment; it identifies what must be measured during subsequent simulation, prototype and operational trials.

RESULTS AND DISCUSSION

Characteristics of the Evidence Base

The literature reveals an evolution from isolated intrusion sensing towards heterogeneous, intelligence-assisted surveillance. Early work established the feasibility of using WSNs for persistent perimeter observation and concentrated on coverage, connectivity and energy-efficient forwarding. Later studies incorporated cameras, Raspberry Pi platforms, fuzzy logic, UAVs and machine-learning classifiers. Recent research increasingly emphasises secure routing, trust management, edge processing and resilience against malicious nodes.

Despite this progression, direct comparison remains difficult. Simulation studies often report network lifetime and packet delivery, whereas prototypes report detection demonstrations, camera streaming or notification functionality. Deployment scale ranges from a few laboratory nodes to mathematical models of extensive border belts. Security assumptions also differ substantially. Table 1 therefore summarises representative contributions and their implementation implications rather than presenting unlike metrics as if they were equivalent.

Table 1: Evidence Profile of Selected Border-Surveillance and Enabling Studies

Study	Context or approach	Principal contribution	Implementation implication
Hammoudeh et al. (2017)	Linear border-belt deployment and routing	Coverage metric and Levels Division Graph routing tailored to border geometry	Demonstrates the importance of cross-layer design; field robustness and adversarial behaviour require further validation.
Arfaoui et al. (2017)	Known or estimated crossing paths	Sensor deployment based on probable intruder trajectories	Can reduce node count, but performance depends on the stability and quality of crossing-path assumptions.
Bhadwal et al. (2019)	Computer-vision border surveillance	Central control and automated visual assistance	Human decision remains necessary; camera coverage and communications are major constraints.
Singh and Singh (2022)	Raspberry Pi, PIR, ultrasonic and camera prototype	Differentiation between animals and people using embedded vision	Useful feasibility evidence; limited scale and environmental diversity constrain generalisation.
Farghly et al. (2021)	Multisensor WSN with fuzzy decision logic	Fusion of PIR, sound, vibration and motion observations	Reduces reliance on a single sensor; membership functions require context-specific calibration.
Laouira et al. (2021)	Hybrid cameras, scalar sensors, radar and UAV architecture	Multilayer scheduling and load balancing for longer lifetime	Highlights heterogeneity; coordination complexity and lifecycle cost increase.
Jayachandran and Vimaladevi (2025)	Secure energy-efficient clustering	Trust-aware cluster-head selection and anomaly detection	Promising against malicious nodes; simulation assumptions need field testing.
Heinzelman et al. (2000)	Adaptive clustering for WSNs	Rotating cluster heads to distribute energy consumption	Foundational energy logic; classical LEACH assumes conditions that may not hold in long linear borders.
He et al. (2003)	Soft real-time routing	Stateless forwarding based on desired delivery speed	Relevant to urgent alerts; congestion and link asymmetry must be monitored.

NIST (2020); ETSI (2024)	IoT cybersecurity baselines	Device identity, configuration, data protection, updates and vulnerability management	Security requires lifecycle support, not encryption alone.

Operational and Technical Requirements

A border-surveillance service must provide persistent observation without assuming that all devices remain connected or fully powered. It should detect relevant physical changes, attach trustworthy time and location information, distinguish urgent events from routine status traffic, and preserve evidence until it is acknowledged by the next tier. The service must further support degraded operation during backhaul failure and prevent the compromise of one node from providing unrestricted access to the wider network.

Requirements are interdependent. Increasing camera coverage improves visual evidence but raises power, storage and privacy demands. Increasing sensor density improves redundancy but complicates channel access, routing and maintenance. Strong cryptography protects data but consumes processing and may extend transmission time. The design process must therefore optimise a portfolio of requirements rather than maximise one metric in isolation.

Table 2: Core Requirements and Their Architectural Responses

Requirement	Operational meaning	Architectural response
Persistent detection	Continuous or duty-cycled monitoring of relevant corridors	Multimodal Tier-1 sensing with sleep scheduling and health reports
Early localisation	Rapid estimate of the event sector and direction	Stable node identifiers, GPS or surveyed coordinates, timestamping and neighbour correlation
False-alarm control	Differentiate wildlife, weather and benign movement from probable intrusion	Temporal filtering, multisensor fusion and event-triggered edge vision
Timely communication	Deliver urgent alerts within an operationally meaningful delay	Priority queues, acknowledgement, store-and-forward buffering and redundant routes
Energy endurance	Minimise maintenance visits and battery replacement	Clustering, adaptive duty cycles, local aggregation, solar assistance where appropriate
Cybersecurity	Resist spoofing, eavesdropping, replay, malicious firmware and unauthorised access	Unique identities, encryption, mutual authentication, signed updates and segmentation
Resilience	Continue essential operation during node, gateway or backhaul failure	Redundancy, cached rules, local logging, degraded mode and controlled recovery
Human usability	Present concise, interpretable and prioritised evidence	Geospatial dashboard, confidence levels, alert grouping and auditable acknowledgement
Maintainability	Enable local diagnosis, replacement and support	Modular hardware, documented interfaces, spare strategy and technician training
Governance and privacy	Ensure lawful and proportionate collection and use	Role separation, retention limits, access logs and human-authorised response

Proposed Three-Tiered Hierarchical Surveillance Strategy

The central result of the synthesis is the architecture shown in Figure 2. The physical area is divided conceptually into three operational zones, although the boundaries may overlap according to terrain and radio coverage. The monitored or remote zone contains the largest number of low-power sensor nodes. The semi-monitored zone contains fewer but more capable edge-intelligence nodes. The protected zone contains gateways, command infrastructure and personnel. Data quality and computational capability increase as an event moves upward through the hierarchy.

The hierarchy is not simply a sequence of devices; it is a mechanism for allocating scarce resources. Tier 1 is optimised for endurance and spatial coverage. Tier 2 is optimised for selective visual confirmation and local intelligence. Tier 3 is optimised for integration, accountability and response. Security and maintenance responsibilities cut across all tiers.

Tier (Layer)	Name	Purpose / Main Function	Key Components	Sensing / Detection Capabilities	Communication Technology	Power Source	Coverage Range	Processing Location	Typical Response / Output
Tier 1 (First Layer)	Perimeter Detection Layer	Detect any intrusion at the outermost perimeter and trigger early alerts.	• MICAz sensor nodes • PIR motion sensor • acoustic sensor • vibration sensor • light sensor	• Motion detection • acoustic signature • ground vibration • light level change • tamper detection	IEEE 802.15.4 (ZigBee); multi-hop mesh network	Battery + solar energy harvesting	Up to 2–5 km between nodes	In-node processing (threshold filtering and event detection)	Intrusion event packet forwarded to Tier 2 gateway.
Tier 2 (Second Layer)	Intelligent Monitoring Layer	Verify, classify and track events using intelligent visual monitoring.	• Motorized AI camera (Raspberry Pi) • pan-tilt unit (PTZ) • YOLOv8 / Tiny YOLO • local storage (SD card) • edge computing module	• Human/vehicle detection • object classification • direction tracking • behavior analysis • false-alarm filtering	Wi-Fi / 4G/5G / LoRa / LTE backhaul link to base station	Solar + battery backup	Up to 5–15 km per camera (depending on terrain and optics)	Edge processing (Raspberry Pi) + event metadata to base station	Validated intrusion alert with image/ video clip sent to Tier 3.
Tier 3 (Third Layer)	Conventional Surveillance Layer	Provide wide-area situational awareness and strategic command control.	• PTZ long-range camera • radar (optional) • thermal/IR camera • surveillance towers • manned control posts	• Wide-area visual surveillance • thermal imaging • radar tracking • night operation • long-range identification	Fiber / microwave backhaul / satellite / LTE (VPN secure link)	Fiber power + generator + solar backup	15–30 km or more (depending on infrastructure)	Centralized control center (high-performance servers)	Operator decision, alert security forces, log and archive event.
Base Station (Command Center)	Central Processing and Decision Layer	Aggregate data, correlate, analyze and support decision making.	• Database server • data fusion engine • AI analytics module • operator dashboard • alert management system	• Multi-sensor data fusion • event correlation • threat-level assessment • historical analysis • report generation	High-speed backhaul (fiber / microwave / satellite)	Grid + UPS + generator + solar backup	Nationwide / regional coverage	Central command center (cloud / on-premise infrastructure)	Decision support, alerts, dispatch, reports and system updates.

Figure 2: Three-Tiered Hierarchical Border-Surveillance Architecture

Tier 1: Remote Multimodal Sensing

Tier 1 comprises clusters of MICAz-class or functionally equivalent low-power sensor nodes. A node may combine passive infrared sensing with acoustic, vibration, seismic, magnetic or light sensors, depending on the threat and environment. The node performs local filtering, attaches its identity and location, and transmits only a compact event record unless a diagnostic request requires additional data. Nodes remain in a low-power state for most of their duty cycle and wake on an interrupt, scheduled health interval or neighbour message.

Multimodality improves interpretive confidence. A single PIR transition may be caused by an animal or vegetation; a correlated acoustic and vibration pattern strengthens the probability of physical movement. For vehicles, magnetic disturbance, low-frequency vibration and characteristic acoustic energy may provide complementary evidence. Sensor fusion at this tier should remain computationally modest. Its purpose is to reject obvious noise and decide whether the event merits transmission, not to make the final operational judgement.

Cluster-heads aggregate observations from nearby nodes, remove duplicates and forward an event summary. Cluster-head responsibility should rotate or be assigned according to residual energy, link quality and topological importance. In long border belts, purely random cluster-head selection may produce inefficient routes; geographic partitions and gateway-aware clustering are therefore preferable.

Tier 2: Edge-Intelligence and Visual Confirmation

Tier 2 contains Raspberry Pi-based edge nodes connected to motorised day/night cameras and, where needed, ultrasonic or additional PIR sensors. These nodes are not required to process video continuously. They remain in a low-power or reduced-processing state until a Tier-1 event identifies an area of interest. The camera is then directed towards the reported bearing or sector, captures a short evidence sequence and executes a lightweight object-detection model.

YOLO-family models are appropriate candidates because they perform single-stage object detection and have lightweight variants that can be deployed on resource-constrained edge platforms. Nevertheless, model selection should be based on measured inference time, energy consumption and class-specific recall under the intended conditions. The system should not assume that a model trained on conventional urban images will recognise people, motorcycles, livestock or partially occluded objects in Sahelian dust, nocturnal infrared imagery or dense vegetation. Local data collection,

transfer learning and periodic drift assessment are necessary.

The output of Tier 2 is not merely an image. It is an evidence package containing the event identifier, camera identity, timestamps, probable classes, confidence values, bounding boxes or track summaries, a representative image or short clip, and device-health information. Where connectivity is weak, metadata and a thumbnail should be sent first, while the full evidence is buffered for later transmission. This prioritisation preserves operational timeliness without discarding evidential detail.

Tier 3: Protected Command, Analysis and Response

Tier 3 resides in a protected base station or regional command facility. A multi-radio gateway translates between the sensor network and IP-based services, validates message origin and forwards events to a secure broker. The platform maintains a geospatial representation of nodes, sectors, alerts, acknowledgements and incident status. It also monitors battery levels, link degradation, silent nodes and repeated false alarms, thereby treating maintenance data as part of operational awareness.

Operators receive prioritised alerts rather than unfiltered sensor messages. A critical alert should show the probable event type, confidence, location, time, evidence source, neighbouring observations and any communication uncertainty. The operator may acknowledge, request additional imagery, dispatch a patrol, coordinate an unmanned aerial vehicle or escalate to another authority. Every consequential action is logged. This human-in-the-loop arrangement constrains the effect of classification error and supports accountability.

Cloud services may be used for long-term storage, model training and cross-sector analysis, but the command tier should preserve essential local capability. If the external Internet is unavailable, a local broker, database and dashboard should continue receiving and displaying events from the border sector. Cloud dependence should therefore be an enhancement rather than a single point of operational failure.

End-to-End Operating Workflow

The operating workflow begins when a Tier-1 node detects a physical change. The node applies debounce rules, verifies that its sensing circuit is healthy and generates an event containing node identity, local time, measured features, residual energy and location. Neighbouring nodes may corroborate the event within a short temporal window. The cluster-head aggregates the evidence and forwards a trigger to the appropriate Tier-2 edge node and to the gateway.

The edge node orients its camera, captures visual evidence and executes classification. A fusion service combines the WSN evidence, image results, time of day, sector sensitivity and recent event history. The event is assigned one of several states: rejected noise, benign or animal activity, watch condition, probable intrusion or critical intrusion. A critical event is transmitted immediately and repeatedly until acknowledged; lower-priority evidence may be queued or grouped.

Failure handling is explicit. If the edge camera is unavailable, the WSN alert is not silently discarded; it is escalated with reduced confidence and a camera-failure flag. If the backhaul fails, the gateway stores events locally and activates an alternative path where available. If observations conflict, the dashboard communicates uncertainty rather than forcing an unjustified conclusion. Figure 3 summarises the workflow.

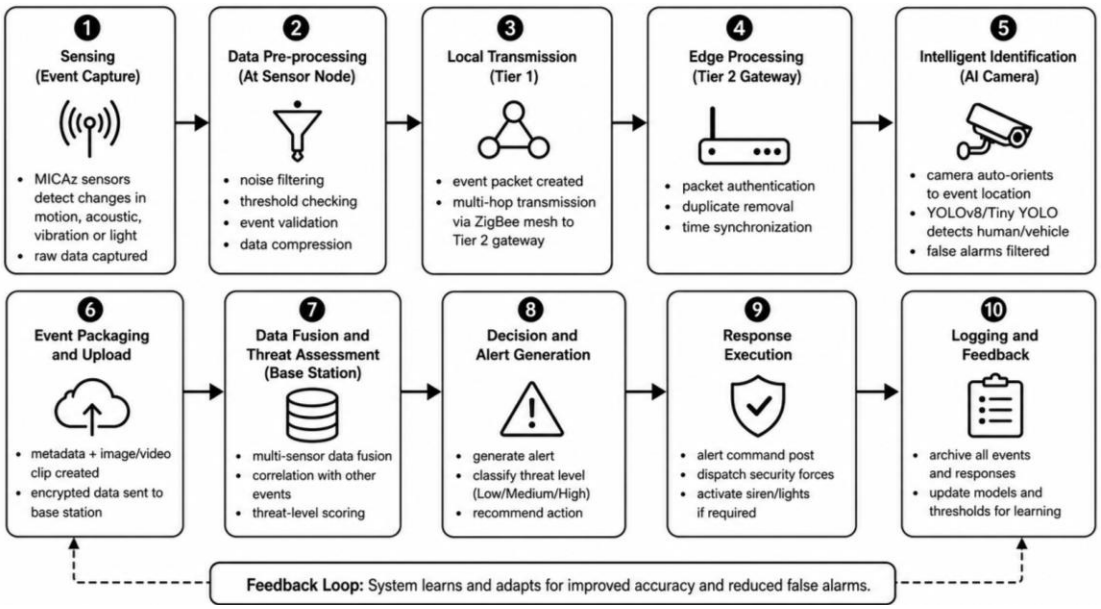

Figure 3: End-to-end intrusion-detection and decision workflow

Hardware and Software Components

MICAz-Class Sensor Nodes

The MICAz mote is a recognised WSN platform built around an Atmel ATmega128L microcontroller and a Chipcon CC2420 IEEE 802.15.4-compatible transceiver. Its limited memory, processing rate and battery supply make it unsuitable for image inference but appropriate for event-driven sensing, local thresholding and low-rate mesh communication. The platform illustrates the design principle that Tier-1 devices should execute only those functions necessary for dependable detection and communication.

In a new deployment, the exact MICAz product may be replaced by a contemporary low-power microcontroller platform with secure storage and sub-GHz communication. The architectural requirement is therefore expressed as a MICAz-class node: a small, battery-powered, programmable device capable of sensing, timestamping, authenticated messaging, sleep scheduling and basic self-diagnostics. This abstraction is important because legacy hardware may have limited commercial availability and weaker native security than current devices.

Sensor Modalities

Passive infrared sensors detect changes in infrared radiation across zones in their field of view. They consume little energy and are useful for motion triggering, but they do not measure identity or exact distance. Acoustic sensors capture sound energy and spectral patterns, while geophones or vibration sensors observe ground-transmitted energy. Ultrasonic sensors estimate distance by measuring echo time and are most useful at constrained paths, gates or narrow corridors rather than over broad open terrain. Magnetic sensors can strengthen vehicle detection.

Sensor selection should be informed by environmental testing. Rain, insects, heat, dust, livestock, vegetation and installation angle can affect measurements. A sensor that performs accurately in an indoor laboratory may generate unacceptable false alarms outdoors. Each node should therefore support site-specific calibration and preserve raw or minimally processed samples for a limited commissioning period so that thresholds can be refined.

Raspberry Pi Edge Nodes and Motorised Cameras

The Raspberry Pi 4 provides a quad-core ARM processor, multiple memory configurations, general-purpose input/output interfaces, camera connectivity, Ethernet, Wi-Fi and Bluetooth. These capabilities permit camera control, local inference, storage, message handling and device management on one platform. The device nevertheless requires more power and thermal management than a mote and should be installed in a protected enclosure with regulated power, surge protection and, where appropriate, solar charging and battery backup.

A motorised camera assembly uses pan-and-tilt actuation to inspect the event sector. The positioning mechanism must include mechanical limits, calibration and a known home position. Night operation may require infrared illumination or a thermal camera, but each option introduces distinct cost, visibility and privacy implications. Video should be acquired in short, event-centred windows unless continuous recording is justified by the site and supported by power and storage.

Gateways, Edge Routers and Platform Services

The gateway bridges low-power sensor protocols and the IP network. It performs protocol translation, authentication, message normalisation, buffering, network management and bidirectional configuration. Where feasible, gateway and edge-router functions can be combined in a hardened industrial device with Ethernet, cellular and LoRaWAN or IEEE 802.15.4 interfaces. Redundant power and backhaul are desirable because gateway failure can isolate an entire sector.

Platform services include an MQTT broker, device registry, time-series or event database, geospatial dashboard, alert engine, audit log and model-management repository. MQTT is suitable for compact asynchronous messaging because it’s publish-subscribe design decouples field devices from applications and provides configurable delivery semantics. The broker should not be exposed directly to the public Internet without strong authentication, encrypted transport, topic-level authorisation and rate control.

Table 3: Principal Hardware and Software Components

Component class	Illustrative technology	Primary responsibility
Tier 1 node	MICAz-class mote or modern low-power MCU	Sensing, local filtering, timestamping, sleep scheduling and authenticated event transmission
Motion sensing	PIR sensor	Low-energy movement trigger; requires environmental calibration
Ground/acoustic sensing	Geophone, vibration or microphone module	Corroborates footsteps, engines or mechanical movement
Distance sensing	Ultrasonic module	Useful at gates or constrained paths; susceptible to angle and weather effects
Tier 2 processor	Raspberry Pi 4/5 or equivalent SBC	Camera control, local inference, buffering, secure messaging and health monitoring
Imaging	Day/night or thermal motorised camera	Visual confirmation, object classification and track evidence
Actuation	Pan-and-tilt mechanism with driver	Points the camera towards a reported sector and returns to a calibrated home position
Gateway	Multi-radio industrial gateway	Protocol translation, device authentication, buffering and network management
Backhaul	Private cellular, microwave, satellite or secure fixed link	Connects remote sectors to the protected command environment
Platform	Local server and optional cloud services	Broker, database, dashboard, analytics, evidence archive and audit
Security hardware	Secure element or TPM where feasible	Protects long-term credentials and supports device attestation

Communication Architecture and Routing Strategy

Network Topology and Message Classes

The proposed network is heterogeneous. Tier-1 nodes form local clusters or multi-hop meshes. Tier-2 nodes communicate with nearby clusters and gateways using a higher-capacity link. Gateways forward normalised events over secure IP backhaul. This arrangement avoids requiring every battery-powered sensor to maintain a direct cellular or Internet connection.

Messages are divided into four classes. Class A contains critical intrusion alerts and requires immediate forwarding, acknowledgement and retransmission. Class B contains watch events and visual-confirmation results. Class C contains node-health and maintenance data. Class D contains configuration and software-management traffic. Separating classes permits different priorities, retention periods and security controls. A battery report should not delay a probable intrusion, while a firmware update should never be initiated through the same unaudited path as an urgent alert.

Energy-Aware and Location-Informed Routing

LEACH provides a useful conceptual basis for energy-aware clustering because it rotates cluster-head responsibility and aggregates local traffic. However, an unmodified randomised LEACH implementation may be inefficient along a narrow border belt where distance to the gateway is highly unequal. Cluster-head selection should therefore consider residual energy, geographic position, link quality, queue state and the availability of an alternative next hop.

GEAR is relevant because it uses geographic and energy information to route towards a target region, while SPEED is relevant because it seeks to maintain a desired packet-delivery speed for soft real-time communication. AODV or AOMDV can be used where routes must be discovered reactively and topology changes are significant; multipath variants provide resilience at the cost of additional state and signalling. The architecture does not mandate a single protocol for every link. It recommends a policy layer that chooses routing behaviour according to urgency, energy and link condition.

For local sensor clusters, IEEE 802.15.4 or a low-power sub-GHz mesh may be appropriate. LoRaWAN can provide long-range transmission of compact events where gateway coverage is available, but it is not suitable for continuous video. Camera evidence should travel through Ethernet, Wi-Fi, microwave or cellular links with adequate bandwidth, preferably after edge processing has reduced the data volume.

Table 4: Communication Options and Deployment Trade-Offs

Technology	Best-fit role	Principal advantage	Principal constraint
IEEE 802.15.4 / ZigBee	Short-range low-rate sensor clusters	Low power, mesh capability and mature mote support	Interference, limited payload and 2.4 GHz propagation in vegetation
LoRaWAN	Long-range compact event and health messages	Wide coverage and very low energy	Low throughput, duty-cycle constraints and gateway planning
Wi-Fi	Protected sites and short-range edge-node links	High throughput and widespread support	Higher energy demand and limited remote coverage
4G/5G cellular	Gateway backhaul and selected edge nodes	Managed wide-area IP connectivity	Recurring cost, coverage dependence and exposure to outages
Microwave / fixed wireless	Point-to-point sector backhaul	Predictable capacity where line of sight exists	Tower, alignment and weather considerations
Satellite	Fallback for highly remote locations	Coverage beyond terrestrial infrastructure	Cost, latency, terminal power and weather effects
MQTT over TLS	Gateway-to-platform messaging	Lightweight publish-subscribe model and delivery controls	Requires careful credential, topic and broker management

Data Model, Time and Delivery Semantics

Every event should contain a globally unique event identifier, device identifier, sector identifier, event time, receive time, location, sensor features, confidence, battery state, software version and integrity information. Event time and receive time must be distinguished because messages can be delayed or delivered out of order. The platform should apply idempotent processing so that retransmission does not create duplicate incidents.

Critical topics should use delivery semantics that provide acknowledgement, but quality of service does not remove the need for application-level confirmation. A broker may confirm message receipt even though an operator has not seen the alert. The incident workflow should therefore distinguish device transmission, broker receipt, platform processing, operator acknowledgement and response dispatch. Each transition should be timestamped for performance evaluation and audit.

Cybersecurity, Privacy and Resilience

A border-surveillance network is an attractive target because a successful attacker could conceal movement, generate diversionary alarms, obtain operational patterns or pivot into connected command systems. The threat model must include remote eavesdropping, message injection, replay, jamming, malicious routing, credential theft, physical node capture, firmware replacement, denial of service, insider misuse and unauthorised access to images.

Security should be implemented as defence in depth, as illustrated in Figure 4. At the device layer, every node requires a unique identity, protected credentials, restricted debug interfaces and verified software. At the communication layer, messages require authentication, integrity protection and anti-replay controls. At the gateway and platform layers, segmentation, role-based access, logging, backup and anomaly detection are required. Operational governance defines who can enrol devices, change thresholds, view identifiable imagery, initiate an update and close an incident.

NISTIR 8259A identifies core IoT capabilities such as device identification, configuration control, data protection, logical access restriction, secure software update and cybersecurity state awareness. ETSI EN 303 645 similarly emphasises the elimination of universal default passwords, vulnerability management, secure updates, protected communication and minimisation of exposed attack surfaces. These principles are directly applicable even though a border device is not a consumer product, because the underlying lifecycle risks are comparable.

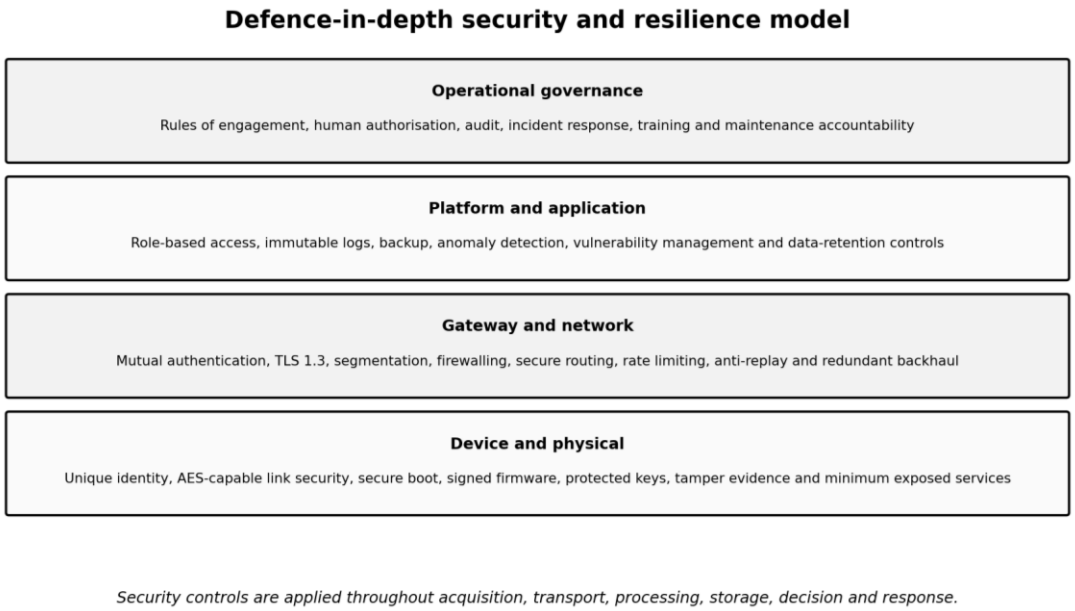

Figure 4: Defence-in-depth security and resilience model

Device Identity, Cryptography and Key Management

Sensor nodes should be provisioned with unique credentials before deployment. Shared network-wide keys create an unacceptable blast radius because the capture of one node compromises the entire sector. Where constrained hardware permits, link-layer AES protection can secure local frames, while gateway-to-platform communication should use TLS 1.3 with mutually authenticated certificates or carefully managed pre-shared keys. Cryptographic choices must be matched to device capability and supported through a practical renewal process.

Long-term keys should be stored in a secure element or protected hardware region where feasible. Commissioning records must link physical asset identifiers, cryptographic identities, installation location and responsible technician. Lost or captured nodes should be revocable without rekeying the whole network. Time or sequence counters are necessary to prevent replay of previously valid intrusion messages.

Secure Updates and Gateway Hardening

Remote software maintenance is necessary because vulnerabilities and model errors will emerge after deployment. Over-the-air updates should therefore be signed, version-controlled, recoverable and delivered through an authenticated management channel. The device should reject unsigned or downgraded software and retain a known-good image for recovery. Update campaigns should be staged by sector to prevent a defective release from disabling the entire surveillance service.

Raspberry Pi gateways and edge nodes require hardening beyond changing default passwords. Unnecessary services should be removed, root login disabled, host firewalls enabled, administrative access restricted to a management network, logs forwarded, storage encrypted where practical and brute-force attempts rate-limited. Containerisation can isolate application components but does not replace host security, patching or least-privilege execution.

Jamming, Physical Capture and Degraded Mode

Radio jamming cannot be solved by encryption because it attacks availability rather than confidentiality. Mitigation includes channel diversity, spread-spectrum mechanisms where supported, directional or redundant links, detection of abnormal noise floors and local buffering. The system should treat simultaneous loss of multiple neighbouring nodes as a potential security event rather than a routine maintenance condition.

Physical capture is plausible in remote terrain. Enclosures should be tamper-evident, installation should avoid conspicuous patterns, and the platform should detect unexpected movement or prolonged silence. A captured node must not reveal credentials that unlock other devices. During connectivity loss, nodes continue sensing, edge cameras continue confirmation where possible, and gateways preserve a local event queue. Degraded mode is successful when essential safety functions remain available and the system clearly indicates reduced confidence.

Table 5: Threat Model and Corresponding Controls

Threat	Operational consequence	Principal controls
Eavesdropping	Unauthorised observation of sensor or image traffic	Link protection, TLS 1.3, key rotation and minimised payloads
Spoofing and injection	Fabricated node or alert messages	Unique device identity, mutual authentication, message integrity and allow-listed enrolment
Replay	Retransmission of a previously valid alarm	Sequence numbers, timestamps, nonce handling and duplicate-event detection
Selective forwarding/sinkhole	Compromised router suppresses or attracts traffic	Multipath routing, neighbour monitoring, route diversity and trust assessment
Jamming	Radio channel made unavailable	Interference detection, alternative channels/backhaul, directional links and local buffering
Physical capture	Node removed, inspected or reprogrammed	Tamper evidence, secure key storage, rapid revocation and minimal stored data
Malicious firmware	Unauthorised software changes device behaviour	Secure boot, signed updates, anti-rollback and staged deployment
Gateway compromise	Attacker reaches multiple devices or platform services	Segmentation, hardened OS, least privilege, firewalling, monitoring and backups
Insider misuse	Authorised user accesses or alters data improperly	Role separation, audit logs, approval workflows and periodic access review
Privacy overcollection	Images or movement data retained without necessity	Data minimisation, local inference, retention limits and purpose-bound access

Energy Management and Network Lifetime

Energy management must be designed at node, link, routing and application levels. At node level, the microcontroller, sensors and radio remain asleep except during scheduled sampling or event response. At link level, payloads are compact and retransmission is bounded. At routing level, traffic is aggregated and high-burden roles are rotated. At application level, cameras are activated only when Tier-1 evidence justifies the energy expense.

A commonly used first-order radio model represents the energy required to transmit a k-bit message across distance d as the electronic energy plus a distance-dependent amplifier cost. For short distances, the amplifier cost is proportional to d^2; for longer multipath links, it can be proportional to d^4. Reception consumes electronic energy proportional to message length. The model can be written as:

$E_{TX}(k,d) = kE_{elec} + k\varepsilon_{fs} d^2, \; d < d_0$; $E_{TX}(k,d) = kE_{elec} + k\varepsilon_{mp} d^4, \; d \geq d_0$

$E_{RX}(k) = kE_{elec}$

The equations demonstrate why repeated long-range transmission from every node is inefficient. Local aggregation and appropriately placed gateways can reduce amplifier cost substantially. Nevertheless, analytical models must be calibrated against real hardware because sensing, wake-up, encryption, flash access, temperature and battery chemistry also affect consumption.

Network lifetime should be reported through several indicators: time to first node death, time to a defined percentage of node failures, coverage loss, number of unreachable sectors and maintenance visits. A network may technically remain alive after strategically positioned nodes have failed, yet no longer provide useful border coverage. Lifetime must therefore be interpreted spatially and operationally.

Energy harvesting can extend service but should not be treated as unlimited power. Solar panels require exposure, cleaning, theft protection and charge control; shaded forest corridors may produce insufficient energy; and high temperatures degrade batteries. The preferred strategy is to minimise demand first, then use harvesting and backup capacity to increase resilience.

Edge Artificial Intelligence and False-Alarm Suppression

The edge-classification pipeline should transform raw frames into operationally meaningful but uncertainty-aware evidence. A lightweight detector identifies candidate persons, vehicles, motorcycles or animals. A tracker can determine movement direction and persistence. Context rules compare the visual result with the WSN event, restricted-zone schedule and recent observations. The system should avoid treating a single low-confidence detection as definitive.

A simple fusion score may combine normalised evidence from PIR or motion sensing (P), acoustic or vibration

sensing (A), seismic or magnetic sensing (S) and visual classification (V):

$F = w_pP + w_aA + w_sS + w_vV, \quad \text{where } \Sigma w_i = 1$

The weights should be calibrated empirically and may vary by sector. A road sector may assign more weight to magnetic and visual evidence, while a forest footpath may emphasise PIR, acoustic and visual correlation. The fusion output should include both a class and a confidence interval or qualitative band. This prevents the dashboard from presenting uncertain machine output as unquestionable fact.

False-alarm reduction must be evaluated alongside missed detection. Increasing the classification threshold may suppress nuisance alerts while allowing partially occluded intruders to pass undetected. The appropriate operating point should be selected according to the security consequence of each error and the capacity of operators to investigate alerts. Precision, recall, class-specific confusion matrices and alert rate per sector are therefore more informative than overall accuracy alone.

Model maintenance is an operational responsibility. New clothing, vehicles, seasons, camera angles and adversarial behaviour can reduce performance. Representative edge cases should be retained for controlled retraining, and every model release should be versioned, tested and reversible. Images used for training or evaluation should be governed by access and retention rules.

System Algorithms

The algorithms below express the principal operating logic. They are implementation-neutral and should be adapted to the selected programming environment, real-time operating system and communication stack.

Algorithm 1: Overall three-tier surveillance process

```
Initialise device identities, clocks, sector maps, routing tables and security state.
Start Tier-1 sensing and periodic health reporting.
WHILE the mission is active:
    Wait for a sensor event, scheduled report or management command.
    IF a Tier-1 event is received:
        Validate integrity, freshness, node health and geographic plausibility.
        Correlate observations from neighbouring nodes within the event window.
        Send an event trigger to the relevant Tier-2 edge node and gateway.
        Capture visual evidence and execute edge classification.
        Fuse sensor, visual and contextual evidence.
        IF confidence is below the rejection threshold: log as noise/false alarm.
        ELSE IF confidence is within the watch band: publish a watch event.
        ELSE: publish a critical alert and retry until acknowledged.
        Record all state transitions, evidence versions and operator actions.
    IF a node, link or gateway failure is detected: enter the defined degraded mode.
    IF compromise is suspected: isolate, revoke and escalate the affected identity.
END WHILE.
```

Algorithm 2: Tier-1 low-power sensing and reporting

```
Wake on sensor interrupt or scheduled sampling interval.
Read enabled sensors and perform calibration/health checks.
Apply debounce, threshold and short-window feature extraction.
IF no relevant phenomenon is present: update health state and return to sleep.
Create event ID; attach timestamp, coordinates, features, battery and software version.
Authenticate and transmit the event to the cluster-head with priority based on severity.
Await acknowledgement for the configured interval; retry through an alternative route if required.
Preserve a bounded local event record and return to low-power operation.
```

Algorithm 3: Tier-2 camera confirmation and evidence packaging

```
Receive a valid event trigger containing sector, bearing and urgency.
Wake the camera subsystem and verify storage, temperature and motor state.
Point the camera to the target sector; acquire a short burst or clip.
Run object detection/tracking and calculate class-specific confidence.
Correlate the visual result with Tier-1 evidence and recent events.
Build an evidence package containing metadata, representative image and model version.
Transmit metadata/thumbnail first; send full evidence when bandwidth permits.
Return the camera to its home position, update health state and enter standby.
```

Performance Evaluation Framework

Performance evaluation should proceed in three stages: controlled laboratory testing, representative pilot deployment and longitudinal operational assessment. Laboratory tests establish component behaviour and security functions. Pilot deployment exposes the system to terrain, weather, animal movement and communication variability. Longitudinal assessment determines whether

batteries, enclosures, calibration, software support and operator workflows remain sustainable.

Figure 5 and Table 6 present a multidimensional framework. Detection quality is necessary but insufficient. A system that detects accurately but delivers alerts too late, consumes batteries rapidly, cannot recover after a gateway failure or overwhelms operators is not operationally effective. Economic and maintenance measures are equally important because border infrastructure must be sustained beyond the demonstration period.

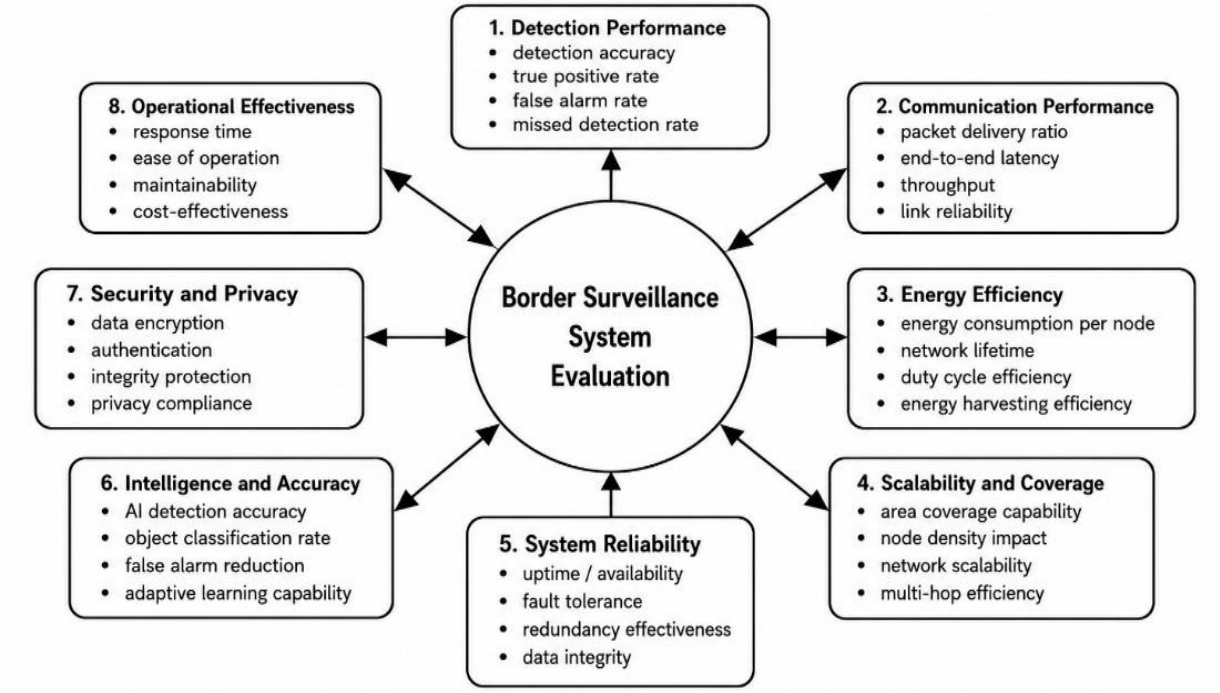

Figure 5: Multidimensional evaluation framework

Table 6: Recommended Performance Measures For Prototype and Field Evaluation

Category	Measures	Purpose
Sensing	True positive, false positive, false negative, precision, recall and class confusion	Establishes whether relevant events are detected and correctly characterised
Localisation	Position error, sector accuracy and direction-of-travel error	Determines whether response assets can be directed effectively
Latency	Sense-to-trigger, trigger-to-image, image-to-alert and alert-to-acknowledgement	Separates delays introduced by each tier
Network	Packet-delivery ratio, message age, route changes, retransmissions and outage duration	Measures communication completeness and timeliness
Energy	Average current, energy per event, duty cycle, first-node death and coverage degradation	Tests endurance and maintenance burden
Edge AI	Inference time, per-class precision/recall, model size, thermal behaviour and power	Tests whether visual confirmation is feasible on the selected platform
Availability	Service uptime, degraded-mode continuity and recovery time	Assesses resilience to node, gateway and backhaul failure
Security	Authentication failures, replay rejection, update success, revocation time and incident detection	Measures whether controls work under attack and maintenance
Human factors	Alert volume, acknowledgement time, operator error and trust	Evaluates usability and alarm fatigue
Economics	Capital cost, lifecycle cost, maintenance visits and cost per monitored kilometre	Supports procurement and scaling decisions
Governance	Access-log review, retention compliance, incident closure and accountability	Tests responsible and lawful operation

Scenario-Based Design Validation

Scenario analysis was used to test the internal coherence of the architecture. The scenarios do not constitute measured results; they identify expected state transitions and potential weaknesses that must be verified experimentally.

Table 7: Scenario-Based Validation of Expected System Behaviour

Scenario	Evidence progression	Expected system response	Residual concern
Animal crosses a remote sector	PIR and acoustic event; camera identifies likely animal	Event is downgraded or logged as benign; no critical dispatch	Low-light misclassification could still create a false alert; retain confidence and evidence.
Person moves through restricted corridor	Multiple Tier-1 nodes correlate; camera detects person and direction	Critical event with location, evidence and repeated acknowledgement	Occlusion or deliberate camouflage may reduce visual confidence; multisensor evidence remains important.
Vehicle enters an unofficial track	Vibration/acoustic/magnetic pattern followed by visual vehicle detection	High-priority alert and predicted route shown on dashboard	Fast movement increases latency sensitivity and requires sector-to-sector event linking.
Tier-2 camera fails after Tier-1 trigger	WSN event arrives but camera health check fails	Alert is escalated with reduced confidence and maintenance flag	Operators must understand that absent imagery is a system limitation, not proof of no intrusion.
Backhaul is unavailable	Gateway receives events but cannot reach central services	Local broker/dashboard and queue continue; alternative backhaul attempted	Storage capacity, clock integrity and later reconciliation must be tested.
Node transmits implausible repeated alerts	Integrity may be valid but behaviour is anomalous	Node is rate-limited, isolated or challenged; neighbouring evidence is requested	A compromised valid credential requires behavioural detection, not encryption alone.
Captured node replays an old critical alarm	Sequence/time checks detect stale event	Replay rejected, security event logged and identity investigated	Clock drift and restart behaviour must not accidentally accept stale messages.

Deployment Framework for African Border and Forest Environments

Deployment should begin with the operational problem rather than the purchase of technology. Authorities should identify priority corridors, incident types, response resources, legal constraints and acceptable alert delays. A physical survey must document terrain, vegetation, seasonal water, settlements, livestock routes, solar exposure, radio propagation, road access and locations suitable for protected gateways. Baseline observation should estimate current patrol effort, incident reporting delay, false-alarm sources and communication availability.

The first deployment phase should cover a small but difficult sector rather than only an easily accessible demonstration site. Nodes should be installed at representative elevations and vegetation densities, and the pilot should include night operation, rain or dust where seasonally relevant, backhaul interruption and deliberate device failure. Acceptance criteria must be agreed before the pilot. A dashboard that displays data is not sufficient evidence of operational success.

Rollout should proceed by sector. Each sector requires an asset register, radio plan, cryptographic enrolment record, maintenance schedule, spare inventory and named technical owner. Local technicians need practical training in battery replacement, sensor calibration, enclosure repair, camera alignment, credential revocation, log interpretation and safe recovery. Vendor support contracts should specify software-update periods, vulnerability disclosure, data-export rights and interface documentation.

Power design should be conservative. Tier-1 nodes may operate on primary batteries or hybrid battery-solar arrangements, while Tier-2 camera nodes require larger solar arrays, charge controllers and protected storage. Lightning, surge, heat and theft must be considered. Where a gateway depends on cellular service, coverage should be measured at the intended antenna height across different times and seasons rather than inferred from a consumer coverage map.

Border communities and legitimate users must be considered. Sensors should not be deployed without understanding customary paths, farms, grazing routes and lawful cross-border interaction. Repeated alerts caused by normal community life will undermine the system and may produce harmful responses. Engagement, clearly defined restricted zones and human review are therefore technical risk controls as well as governance requirements.

Lifecycle Cost and Maintainability

The purchase price of a sensor node represents only a fraction of lifecycle cost. Deployment requires surveying, poles or burial, enclosures, civil work, gateways, antennas, solar equipment, secure provisioning, servers, connectivity, operator training, spare parts, periodic visits and disposal. Camera cleaning, vegetation clearance, battery degradation and software updates can dominate recurring expenditure.

Cost should be reported per monitored kilometre, per effective detection sector and per year of service. A low-cost node that fails frequently may be more expensive than a robust node with higher acquisition cost. Likewise, a sophisticated model that requires specialist retraining may be unsustainable where local expertise is unavailable. Maintainability should be a weighted procurement criterion alongside detection performance.

Modularity reduces lock-in and simplifies replacement. Standard identifiers, documented message schemas, exportable logs and replaceable sensing modules allow the institution to evolve the system without discarding the entire architecture. The gateway should not depend on a proprietary cloud interface for basic local operation.

The principal benefit of the architecture is selective escalation. Tier-1 sensors provide broad, energy-efficient awareness; Tier-2 edge nodes improve evidential quality; and Tier 3 coordinate's human response. This reduces the need to stream every observation continuously and may lower false-alarm burden. The separation of responsibilities also permits incremental deployment: an organisation can begin with sensing and operator dashboards, then introduce visual confirmation or advanced analytics after maintenance capacity has been established.

The architecture remains subject to significant limitations. Determined intruders may avoid known sensors, mask acoustic signatures, exploit terrain, damage equipment or jam communication. Cameras may fail under dust, darkness, rain, foliage or deliberate concealment. Machine-learning models can misclassify unfamiliar objects and may encode bias from unrepresentative training data. WSN nodes can suffer clock drift, battery failure and route instability. No design can guarantee complete border coverage or eliminate the need for patrols and intelligence from communities.

The present study is conceptual and does not report field measurements. Its contribution is the integration of requirements and the specification of what should be tested. Future work should implement a prototype using a current secure low-power node alongside the legacy MICAz reference, evaluate different routing policies in NS-3, OMNeT++ or a physical testbed, and measure latency, energy, coverage and false-alarm performance under representative terrain.

Research is also required on adaptive sensor fusion, adversarially robust edge vision, privacy-preserving analytics, secure localisation, jamming detection, low-cost thermal imaging and autonomous health diagnosis. Comparative studies should report complete lifecycle cost and maintenance effort, not only laboratory accuracy. Evidence from African borders, forest reserves and rural security corridors is particularly limited and should be developed through ethically governed field partnerships.

A further research direction is the integration of mobile nodes. UAVs can be dispatched to validate a high-confidence event or restore temporary communication, but their endurance, weather sensitivity, regulation and detectability must be considered. Mobile gateways and patrol devices can also collect buffered events from disconnected clusters. Such extensions should preserve the principle that automation supports, rather than bypasses, accountable human decision making.

CONCLUSION

This paper has developed a WSN-based border-security architecture organised around a three-tiered hierarchical surveillance strategy. The proposed design recognises that expansive and sparsely monitored borders cannot be secured effectively through a single sensor, camera or communication technology. Persistent awareness requires the coordinated performance of low-power sensing, event-triggered edge intelligence, secure communications, geospatial platform services, human decision making and maintainable field infrastructure. The first tier distributes multimodal MICAz-class sensor nodes across remote sectors to detect and correlate motion, acoustic, vibration, seismic and related phenomena. The second tier uses Raspberry Pi-based motorised cameras to acquire visual evidence and execute lightweight object classification only when an event warrants the additional energy and bandwidth. The third tier integrates gateways, message brokers, dashboards, audit records and authorised response. This progressive improvement of evidence is intended to conserve power, reduce nuisance alarms and provide operators with a clearer basis for action. The study further demonstrates that cybersecurity and resilience must be designed across the entire lifecycle. Unique device identity, authenticated and encrypted communication, anti-replay controls, signed updates, gateway hardening, segmentation, revocation,

local buffering and auditable operator action are not optional enhancements. They are prerequisites for trusting a system whose failure or manipulation could have serious security consequences. No empirical performance has been fabricated. Instead, the paper provides algorithms, scenario analysis and a multidimensional evaluation set covering detection, localisation, latency, packet delivery, energy, edge inference, availability, security, human factors, economics and governance. These measures establish a defensible agenda for simulation, prototyping and field testing. For Nigeria and comparable Sub-Saharan African environments, the most viable route is phased and context-sensitive deployment. Priority sectors should be surveyed, difficult conditions included in pilots, essential local operation preserved during outages, and maintenance responsibilities assigned before scale-up. A modest architecture that local personnel can secure and sustain will provide greater public value than an elaborate demonstration that deteriorates after installation. The proposed hierarchy therefore offers a practical research and engineering foundation for extending conventional surveillance while retaining human judgement, legal accountability and operational realism.

REFERENCES

Africa Center for Strategic Studies. (2025). Region in focus: The Sahel. Africa Center for Strategic Studies.

Akyildiz, I. F., Su, W., Sankarasubramaniam, Y., & Cayirci, E. (2002). Wireless sensor networks: A survey. Computer Networks, 38(4), 393–422. https://doi.org/10.1016/S1389-1286(01)00302-4

Akkaya, K., & Younis, M. (2005). A survey on routing protocols for wireless sensor networks. Ad Hoc Networks, 3(3), 325–349. https://doi.org/10.1016/j.adhoc.2003.09.010

Arfaoui, I., Boudriga, N., Trimeche, K., & Abdallah, W. (2017). WSN-based border surveillance systems using estimated known crossing paths. In Proceedings of the 15th International Conference on Advances in Mobile Computing & Multimedia (pp. 182–190). ACM. https://doi.org/10.1145/3151848.3151869

Bhadwal, N., Madaan, V., Agrawal, P., Shukla, A., & Kakran, A. (2019). Smart border surveillance system using wireless sensor network and computer vision. In 2019 International Conference on Automation, Computational and Technology Management (ICACTM). IEEE.

Crossbow Technology. (2005). MICAz wireless measurement system: MPR2400 user manual and datasheet.

El Khediri, S., Khan, R. U., Nasri, N., & Kachouri, A. (2021). Energy efficient adaptive clustering hierarchy approach for wireless sensor networks. International Journal of Electronics, 108(1), 67–86. https://doi.org/10.1080/00207217.2020.1793415

European Telecommunications Standards Institute. (2024). Cyber security for consumer Internet of Things: Baseline requirements (ETSI EN 303 645 V3.1.3).

Farghly, M., Elgammal, A., & Elkhatib, M. (2021). Intelligent border using wireless sensor network. Military Technical College, Cairo, Egypt.

Hammoudeh, M., Al-Fayez, F., Lloyd, H., Newman, R., Adebisi, B., & Bounceur, A. (2017). A wireless sensor network border monitoring system: Deployment issues and routing protocols. IEEE Sensors Journal, 17(8), 2572–2582. https://doi.org/10.1109/JSEN.2017.2672501

He, T., Stankovic, J. A., Lu, C., & Abdelzaher, T. (2003). SPEED: A stateless protocol for real-time communication in sensor networks. In Proceedings of the 23rd International Conference on Distributed Computing Systems (pp. 46–55). IEEE. https://doi.org/10.1109/ICDCS.2003.1203451

Heinzelman, W. R., Chandrakasan, A., & Balakrishnan, H. (2000). Energy-efficient communication protocol for wireless microsensor networks. In Proceedings of the 33rd Hawaii International Conference on System Sciences. IEEE. https://doi.org/10.1109/HICSS.2000.926982

Hwang, B. R., & Shin, A. S. (2015). Design of border surveillance and control system based on wireless sensor network. KIPS Transactions on Computer and Communication Systems, 4(1), 11–14.

IEEE. (2020). IEEE standard for low-rate wireless networks (IEEE Std 802.15.4-2020). Institute of Electrical and Electronics Engineers.

Internet Engineering Task Force. (2018). The Transport Layer Security (TLS) protocol version 1.3 (RFC 8446). https://doi.org/10.17487/RFC8446

Jayachandran, J., & Vimaladevi, K. (2025). A dual-layer secure and energy-efficient model for border surveillance using sea-lion-inspired strategy in wireless sensor networks. Scientific Reports.

Kanoujiya, A., Jha, R., Rayal, S., & Khan, A. (2020). Border security system using Arduino, ultrasonic sensors and IoT. International Research Journal of Engineering and Technology, 7(5), 3293–3299.

Laouira, M. L., Abdelli, A., Ben Othman, J., & Kim, H. (2021). An efficient WSN based solution for border surveillance. IEEE Transactions on Sustainable Computing, 6(1), 54–65. https://doi.org/10.1109/TSUSC.2019.2894511

LoRa Alliance. (2020). LoRaWAN L2 1.0.4 specification. LoRa Alliance Technical Committee.

Marina, M. K., & Das, S. R. (2001). On-demand multipath distance vector routing in ad hoc networks. In Proceedings of the 9th IEEE International Conference on Network Protocols (pp. 14–23). IEEE. https://doi.org/10.1109/ICNP.2001.992756

National Institute of Standards and Technology. (2020). IoT device cybersecurity capability core baseline (NISTIR 8259A). https://doi.org/10.6028/NIST.IR.8259A

OASIS. (2019). MQTT version 5.0. OASIS Standard.

Okumu, W., & Ikelegbe, A. (Eds.). (2010). Militias, rebels and Islamist militants: Human insecurity and state crises in Africa. Institute for Security Studies.

Perkins, C. E., & Royer, E. M. (1999). Ad-hoc on-demand distance vector routing. In Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (pp. 90–100). IEEE. https://doi.org/10.1109/MCSA.1999.749281

Raspberry Pi Ltd. (2026). Raspberry Pi 4 Model B product brief and specifications.

Rehan, W., Fischer, S., & Rehan, M. (2017). A critical review of surveys emphasizing on routing in wireless sensor networks: Concepts, operational challenges and prospects for future applications. Sensors, 17(8), 1713. https://doi.org/10.3390/s17081713

Rescorla, E. (2018). The Transport Layer Security (TLS) protocol version 1.3 (RFC 8446). RFC Editor. https://doi.org/10.17487/RFC8446

Singh, R., & Singh, S. (2022). Smart border surveillance system using wireless sensor networks. International Journal of System Assurance Engineering and Management, 13(2), 880–894. https://doi.org/10.1007/s13198-021-01291-9

Texas Instruments. (2022). CC2420: 2.4 GHz IEEE 802.15.4/ZigBee-ready RF transceiver datasheet.

United Nations Institute for Disarmament Research. (2024). Banditry violence in Nigeria's North West: Insights from affected communities. UNIDIR.

United Nations Office on Drugs and Crime. (2024). Impact of transnational organized crime on stability and development in the Sahel: Transnational organized crime threat assessment. UNODC.

Ultralytics. (2023). YOLOv8 documentation and model overview.

Yaseen, M. (2024). What is YOLOv8: An in-depth exploration of the internal features of the next-generation object detector. arXiv. https://doi.org/10.48550/arXiv.2408.15857

Yu, Y., Govindan, R., & Estrin, D. (2001). Geographical and energy aware routing: A recursive data dissemination protocol for wireless sensor networks. UCLA Computer Science Department Technical Report UCLA/CSD-TR-01-0023.

Zigbee Alliance. (2017). Zigbee specification. Connectivity Standards Alliance.