© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723365

Cyber-Resilient Network Security Architecture for Al-Haramain Digital Infrastructure: Threat Detection, Zero Trust and Operational Continuity at The Grand Mosque in Makkah

ILYAS SIDDIQUI MOHAMMAD
ORCiD: 0009-0009-5904-7181

Abstract- ***The Al-Haramain operational environment increasingly depends on interconnected sensors, smart cameras, edge-computing nodes, high-capacity networks, identity services, analytics platforms and control-room applications. This dependency expands the cyber-attack surface of a safety-critical digital infrastructure in which loss of availability, manipulation of data, credential compromise or delayed detection could disrupt operational awareness. This review therefore examines cybersecurity and network resilience as its primary subject; real-time crowd monitoring is used only as the operational case through which security requirements are evaluated. An integrative review of literature published from 2020 to 2025 synthesises evidence on Internet of Things security, network intrusion detection, multi-access edge-computing security, federated anomaly detection, zero-trust architecture, microsegmentation, model integrity and cyber recovery. The study identifies principal attack paths across field devices, edge workloads, 5G and fibre transport, application services, data and machine-learning pipelines, privileged identities and third-party dependencies. It proposes a literature-derived conceptual/reference defence-in-depth architecture comprising cryptographic device identity, secure boot and attestation, default-deny segmentation, encrypted communications, continuous network and workload telemetry, adaptive threat detection, least-privilege access, protected model pipelines, immutable logging, verified backups and orchestrated recovery. Cyber resilience is assessed through mean time to detect and contain, attack-path reduction, trusted-service coverage, failover time, recovery integrity and bounded degradation. The architecture is not presented as an operationally validated cybersecurity framework; a representative cyber-range/simulation pathway is specified for future empirical testing. The review concludes that Al-Haramain digital services require an identity-centred, observable and recoverable security architecture capable of maintaining trusted minimum operations during partial compromise.***

I. INTRODUCTION

The Grand Mosque in Makkah is supported by an expanding digital ecosystem that includes connected sensing, smart cameras, edge analytics, fifth-generation communications and automated operational support. Public reports describe enhanced connectivity and artificial-intelligence capabilities for monitoring gates, movement and congestion (Saudi Press Agency 2024, 2025). These capabilities improve situational awareness, but they also create a safety-critical cybersecurity dependency. A compromised sensor, edge node, network route, service account or analytics model can affect the integrity, availability and timeliness of information presented to operators. The central research problem is therefore not how to manage crowds; it is how to secure and sustain the digital infrastructure that produces operational intelligence under attempted intrusion, component failure and deliberate manipulation.

Al-Haramain operations present a demanding network-security context because heterogeneous devices and services must exchange data across field, edge, transport, platform and control-room layers. The resulting attack surface includes device impersonation, malicious firmware, stream replay, edge-workload compromise, denial of service, lateral movement, route manipulation, credential theft, privileged misuse, software-supply-chain compromise, poisoned model updates and unauthorised data extraction. A conventional perimeter model cannot adequately protect such a distributed environment. Security must instead verify every device, workload, user and data flow; restrict communication to explicitly authorised paths; monitor

behaviour continuously; and preserve essential services when preventive controls fail.

This paper defines cyber resilience as the capacity of the Al-Haramain digital infrastructure to anticipate, withstand, detect, contain, recover from and learn from cyber disruption while maintaining a trustworthy minimum operational service. Cyber resilience is broader than intrusion prevention or classifier accuracy. It requires the architecture to limit blast radius, provide alternative communications and processing paths, expose the trust and freshness of information, restore systems from verified states and retain sufficient forensic evidence for improvement. This interpretation reflects research that treats resilience in industrial and Internet of Things environments as a combined property of protection, continuity, recovery and adaptive governance (Alrumaih et al. 2023; Li et al. 2024; Zanasi et al. 2024).

The study aims to propose a literature-based conceptual/reference cyber-resilient network security architecture for the Al-Haramain digital infrastructure, using smart sensor-based monitoring at the Grand Mosque as a bounded case application. Its objectives are to map critical digital assets and trust boundaries; identify cyber threats capable of compromising confidentiality, integrity, availability, authenticity or operational continuity; evaluate zero-trust, segmentation, detection and recovery controls suitable for distributed edge and Internet of Things environments; specify a layered defence-in-depth reference architecture; and define measurable cyber-resilience indicators together with a cyber-range/simulation validation pathway. The research questions are: RQ1, what assets, identities, data flows and trust boundaries define the Al-Haramain monitoring attack surface; RQ2, which threats create the greatest risk to network integrity and service availability; RQ3, how should preventive, detective, containment and recovery controls be integrated; and RQ4, which metrics and representative simulation tests should be used to assess whether essential digital services remain trustworthy during and after a cyber incident?

II. CYBERSECURITY CONTEXT, ATTACK SURFACE AND THREAT MODEL

Public information indicates an interconnected operational ecosystem rather than a stand-alone monitoring application. Connectivity, artificial-intelligence services, sensor readers and smart cameras form only the visible endpoints of a wider chain that may include gateways, edge clusters, fibre and wireless aggregation, identity infrastructure, analytics platforms, management interfaces, dashboards, repositories and backup systems (Saudi Press Agency 2024, 2025). Each component is a security dependency and a possible pivot point. The reference threat model consequently treats field devices, edge workloads, network fabrics, administrative identities, data stores, machine-learning artefacts and external update channels as separate trust domains. No claim is made about the confidential topology or controls currently deployed at Al-Haramain.

The most consequential risk is loss of trust in the digital service chain. An attacker need not disable the entire network to create operational harm; selective manipulation of timestamps, device identities, routes, sensor streams or model outputs may be sufficient to produce misleading information while the system appears available. Conversely, an over-aggressive containment action can preserve security policy while unnecessarily removing essential coverage. The architecture must therefore protect both information integrity and service continuity. Every operational output should carry evidence of device provenance, timestamp validity, processing state and confidence, enabling the security platform to distinguish a verified measurement, a degraded estimate and an untrusted result.

The threat model covers sensor spoofing, stream replay, malicious firmware, camera or gateway takeover, compromise of edge containers or virtual machines, exploitation of exposed services, distributed denial of service, route disruption, time-synchronisation attacks, credential theft, privileged misuse, policy tampering, software-supply-chain compromise, adversarial input, model poisoning and unauthorised disclosure. Threats are assessed according to their effect on confidentiality, integrity,

availability, authenticity, latency and recoverability. Privacy is incorporated through data minimisation, purpose limitation, edge processing, controlled retention and auditable access. The security objective is a verifiable and recoverable digital service, not maximum collection of operational data.

III. REVIEW METHODOLOGY

An integrative review was selected because the cybersecurity problem crosses Internet of Things security, network intrusion detection, multi-access edge computing, 5G security, zero trust, federated learning, software and model integrity, incident response and cyber resilience. Literature on Hajj technology and crowd analytics was included only to define the protected service, its latency requirements and the consequences of corrupted or unavailable information. The synthesis was therefore security-led: operational sources established assets and dependencies, while cybersecurity sources supplied the threat, control, containment and recovery evidence used to construct the architecture.

The review covered January 2020 to December 2025 and searched Scopus, Web of Science, IEEE Xplore, ScienceDirect, SpringerLink, ACM and Wiley. Security-focused terms included cyber resilience, network security, zero trust, microsegmentation, Internet of Things intrusion detection, network anomaly detection, multi-access edge-computing security, 5G security, device attestation, identity and access management, federated intrusion detection, model poisoning, software supply chain, incident response and disaster recovery. These were combined with Al-Haramain, Grand Mosque, smart sensor, smart camera, edge intelligence and real-time monitoring. Official Saudi sources were used only to establish publicly documented operational context, not to infer private topology, vulnerabilities, incidents or security performance.

A publication was included when it was published in English between 2020 and 2025, addressed at least one security architecture layer, and provided sufficient methodological or architectural detail for threat-control synthesis. Priority was given to work on deployable controls, distributed detection, identity-centric security, containment, recovery, adversarial robustness and continuity. Crowd-analytics studies were retained only where they established a protected asset, latency constraint, data-integrity dependency or edge-processing requirement. Purely operational crowd-management studies without a direct security implication, generic artificial-intelligence commentary, duplicates and sources lacking an identifiable contribution to security or resilience were excluded. The formal synthesis used 28 relevant sources; two 2026 public sources were retained solely as supplementary case context.

Sources were coded into seven cybersecurity themes: asset and attack-surface identification; device and workload trust; secure edge computing; segmented and redundant transport; identity and zero-trust enforcement; threat detection and collaborative analytics; and incident containment, recovery and assurance. Evidence was converted into a threat-control-resilience matrix that distinguishes prevention, detection, containment, continuity and recovery. This distinction is essential: authentication may prevent unauthorised attachment, network detection may identify malicious traffic, segmentation may restrict lateral movement, redundant processing may preserve service, and verified restoration may remove persistence. No single control or detection model provides cyber resilience on its own.

No access was made to or inferred from any private Al-Haramain topology, incident log, camera feed, security configuration, or personally identifiable dataset. As a result, the design presented here is a literature-derived conceptual/reference architecture and research synthesis, not an operationally validated cybersecurity framework or a statement regarding the precise configuration currently in use at the Grand Mosque. This point is methodologically significant because public case evidence shows that smart sensing and advanced connectivity are in place but does not reveal internal segmentation, vendor architecture, cryptographic decisions, response playbooks or measured security performance. The recommendations are therefore expressed as requirements and control patterns. Their empirical validation should be performed by authorised stakeholders in a representative cyber range or simulation that reproduces the relevant functional

layers and trust boundaries without exposing confidential production systems.

IV. LITERATURE SYNTHESIS

4.1 Protected digital assets and cyber-attack surface

The protected system is an end-to-end digital service chain, not the crowd-management process itself. Publicly described sensors, cameras and connectivity create inputs to a wider architecture of device identities, firmware, time services, gateways, edge nodes, transport networks, analytics applications, databases, operator interfaces and administrative tools (Saudi Press Agency 2024, 2025). Hajj technology studies help identify these operational dependencies (Felemban et al. 2020; Albattah et al. 2021; Abalkhail and Al Amri 2022), but the security question concerns how compromise propagates across them. A field device can be an entry point, an edge node a persistence location, a flat network a lateral-movement path, a service account a privilege-escalation route, and analytics pipeline a means of manipulating decisions. Asset inventory and trust-boundary mapping are therefore the first security controls.

The integrity of operational analytics depends on the security provenance of its inputs and processing stages. Crowd-counting research demonstrates legitimate uncertainty caused by scale, occlusion, perspective and domain variation (Fan et al. 2022; Deng et al. 2024). Cybersecurity controls must distinguish these natural errors from spoofing, replay, adversarial input and model tampering. Edge inference can reduce latency and maintain local service during backhaul disruption, but it places credentials, models and cached data in distributed nodes closer to exposed networks (Yi et al. 2024; Lin and Hu 2024). Secure boot, signed firmware, workload identity, hardware-backed attestation, encrypted local state and controlled isolation are therefore required. Model confidence without device, time, network and workload provenance is insufficient for a safety-critical digital service.

4.2 Intrusion detection and adaptive security analytics.

Whenever Internet of Things security is reviewed, it is found that the variety of devices, the limited available resources, the constantly changing nature of attacks, and the lack of high-quality datasets all make intrusion detection difficult (Ahmad and Alsmadi 2021; Khraisat and Alazab 2021; Alsoufi et al. 2021). Although machine-learning techniques are able to enhance classification beyond the use of static signatures, their practical value is determined by the need to control false positives, the problem of concept drift, computational cost, and the ability to provide explanations. Dina and Manivannan (2021) demonstrate the widespread application of machine learning in the area of network intrusion detection, while the wider body of literature cautions against assuming that accuracy achieved in laboratory settings ensures effectiveness when applied to constantly changing production networks.

At the Grand Mosque, a large number of false alarms during peak times could distract the operators and failures to spot anomalies might allow for silent degradation. Detection must therefore rely on multiple layers of telemetry and risk scoring rather than on a single model threshold. Federated learning delivers a way of sharing security knowledge without having to centralize all the raw data. Mothukuri et al. (2022) have shown how federated anomaly detection can be used in relation to Internet of Things attacks, whereas Campos et al. (2022) and Agrawal et al. (2022) have identified the benefits and challenges associated with federated intrusion detection. Since the method is suitable for use in distributed areas, the local nodes send updates to the model while keeping their own observations locally. However, federation does bring in new kinds of threats. SecFedNIDS in particular looks at the problem of poisoning in federated network intrusion detection (Zhang et al. 2022), and Sarhan et al. (2022) combine hierarchical federation with collaboration supported by blockchain. These studies lead to a careful conclusion that although collaborative learning can improve cross-zone visibility, the updates should be authenticated, kept within certain limits, based on reputation, and checked for poisoning and drift before they are used to influence safety-critical policies.

4.3 Edge, 5G, and network continuity.

While multi-access edge computing can help to decrease latency and lessen reliance on backhaul, making it apt for real-time video and sensor analytics, it also brings with it risks relating to distributed trust, virtualization, orchestration, and privacy (Ranaweera

et al. 2021). Fifth-generation networks offer programmable quality-of-service and the ability to logically segment; however, such benefits are only valuable if the management planes, identities, and the boundaries between slices or segments are secured. Since the publicly announced modernisation of the Grand Mosque includes 5G capabilities (Saudi Press Agency 2024), the architecture should treat high-capacity connectivity as a resource that contributes to resilience rather than as a single point of dependence. Priority paths should be assigned to critical sensor traffic, and operational areas should be made separate so that a volumetric attack or a compromise of a field segment does not cause the capacity needed by the whole monitoring service to be consumed. A network design that is practical will therefore integrate segmentation with redundancy. Microsegments can be used to isolate various groups such as cameras, gate sensors, edge clusters, management services, and operator workstations on the basis of function.

If a communication path is lost, redundant fibre or wireless links, different aggregation points, and local buffering can ensure that some level of service is still maintained. The key principle is that segmentation has to correspond to the operational blast radius. A security breach at one gate should not automatically grant access to model repositories, identity services, or other areas. Similarly, it should not be necessary to shut down all crowd monitoring if it is possible to isolate a smaller trust boundary. Cyber resilience is achieved as a result of having narrow trust domains and providing multiple means of sustaining essential data flows.

4.4 Zero trust, provenance, and recoverability.

Instead of relying on a continuous trust based merely on a device's network location, zero-trust research takes this approach. Li et al. (2024) apply the principles of zero trust to future industrial Internet of Things environments, Zanasi et al. (2024) suggest a flexible zero-trust architecture for industrial Internet of Things infrastructures, and Ameer et al. (2024) design a zero-trust architecture incorporating usage control for Internet of Things systems. In the case of crowd monitoring, these principles result in unique device and workload identities, clear authorization, mutual authentication, continuous posture checks, the concept of least privilege, and policy decisions that take context into account. A camera which was trusted yesterday must not continue to have unrestricted access if there has been a firmware anomaly or a certificate issue.

Cyber-resilience literature additionally emphasizes that protection must include response and restoration. Alrumaih et al. (2023) identify resilience challenges in industrial networks, while Epiphaniou et al. (2023) use digital-twin concepts to model low-resilience points in Internet of Things and cyber-physical systems. A digital twin or cyber range can be useful here as an assurance environment: operators can replay sensor breakdowns, denial-of-service conditions, credential abuse, or malicious model updates without experimenting on live pilgrimage operations. The value of the twin is not to reproduce every proprietary detail publicly, but to test whether security policies, rerouting, quarantine, backup inference, and recovery sequences preserve acceptable crowd-state visibility under realistic faults.

Table 1 Thematic synthesis of 2020-2025 literature and the case-specific security gap

Theme	Representative evidence	What the literature establishes	Gap addressed by this review
IoT and network intrusion detection	Ahmad and Alsmadi (2021); Khraisat and Alazab (2021); Alsoufi et al. (2021)	Distributed devices require behavioural detection beyond static signatures.	Detection must be tied to containment, continuity and acceptable false-positive burden.

Theme	Representative evidence	What the literature establishes	Gap addressed by this review
Zero trust and microsegmentation	Li et al. (2024); Zanasi et al. (2024); Ameer et al. (2024)	Identity-centred access and default-deny segmentation reduce implicit trust and blast radius.	Policies must support legacy devices, low latency and controlled degraded operation.
Federated security analytics	Mothukuri et al. (2022); Campos et al. (2022); Agrawal et al. (2022); Zhang et al. (2022)	Distributed learning can share threat knowledge while limiting raw-data movement.	Model-update provenance, poisoning resistance and rollback require explicit controls.
MEC and distributed network security	Ranaweera et al. (2021)	Edge networking improves responsiveness but expands distributed trust and orchestration risk.	Security architecture requires attested workloads, path diversity and zone-level isolation.
Cyber-resilience assurance	Alrumaih et al. (2023); Epiphaniou et al. (2023)	Resilience combines protection, containment, recovery, simulation and learning.	Al-Haramain validation requires attack-chain testing and verified recovery metrics.
Protected operational service	Felemban et al. (2020); Albattah et al. (2021); Saudi Press Agency (2024, 2025)	Sensors, cameras, communications and analytics create an interconnected service chain.	Assets, identities, trust boundaries and cyber dependencies require end-to-end mapping.
Analytics and edge integrity	Fan et al. (2022); Deng et al. (2024); Yi et al. (2024); Lin and Hu (2024)	Edge analytics reduces latency but distributes models, credentials and sensitive data.	Model confidence must be combined with device, time, network and workload provenance.
Cyber-resilience assurance	Alrumaih et al. (2023); Epiphaniou et al. (2023)	Resilience expands security toward containment, recovery, simulation and learning.	Grand Mosque validation needs combined cybersecurity and crowd-safety metrics.

V. CYBER-RESILIENT NETWORK SECURITY ARCHITECTURE

The synthesis supports a six-layer defence-in-depth architecture governed by zero-trust principles. The layers are trusted devices and gateways; hardened edge and multi-access edge computing; segmented and redundant 5G and fibre transport; integrated security analytics and policy enforcement; protected identity, data and model services; and incident continuity and verified recovery. Every connection is authenticated and authorised, every critical flow is observable, and no device or workload receives trust solely because of network location. Operational monitoring is the protected use case, while identity, segmentation, threat detection, containment and recoverability are the architectural core.

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723365

Fig. 1 Cyber-resilient network security architecture for Al-Haramain digital infrastructure

The second layer includes the secure edge and the multi-access edge computing component. It performs time-sensitive preprocessing, video inference, local density estimation, and data fusion right beside the area that is being monitored. The edge nodes have to function with a minimal service set, use workload identities, isolate the inference containers or virtual machines, and encrypt the data that is cached. Hardware-backed attestation is recommended for nodes which host important models. Most important of all, edge applications should have a defined degraded mode; if the upstream analytics platform becomes unreachable, the zone should still produce a local crowd indicator based on verified sensors, store in a buffer any essential events, and give a clear confidence label. In the case of a camera being found to be untrusted, the fusion logic should either reduce its weight or replace it with data from neighbouring sensors rather than just proceeding as if all the inputs were still trustworthy.

The third layer is concerned with resilient transport. Resilience can be achieved by combining fibre, using managed wireless links and taking advantage of 5G connectivity, thus gaining capacity and providing multiple routes, but the logical architecture has to be clearly split into separate sections. Each of the following — sensor access networks, edge clusters, security monitoring, administrative management, analytics services and operator interfaces — should be located in its own trust zone with communication set at default to denial. Security gateways or software-defined policies must permit only the necessary traffic flows and should impose rate limits on traffic from abnormal sources. Quality-of-service protection for important crowd telemetry must be given against less essential traffic, and alternate routes should be tested regularly. There also needs to be protection for time synchronization since replayed or altered timestamps can corrupt the calculation of movement rates even if the individual sensor values appear plausible.

The fourth layer consists of the security analytics and policy plane. A single, unified monitoring function should receive data from network detection and response, from endpoint or workload telemetry, from identity events, from certificate status, and from application logs. Although machine-learning anomaly detection can supplement the use of signatures, any automated decisions have to be able to distinguish between an observation, a recommendation for containment, and a safety-critical action. When it is highly certain that one device has been compromised, it should be automatically isolated; yet any more extensive changes that could reduce the sensing coverage must be controlled by policy rules which consider redundancy and require operator supervision. The federated learning component can share knowledge of attacks between different zones without having to transfer the raw video, but the model updates must be checked for poisoning, abnormal magnitude, and provenance before they are aggregated (Mothukuri et al. 2022; Zhang et al. 2022).

The data, model, and identity trust plane make up the fifth layer. Along with live streams, crowd-monitoring systems also make use of model weights, configuration, calibration, certificates, roles, and

© SEP 2026 | IRE Journals | Volume 10 Issue 3 | ISSN: 2456-8880
DOI: https://doi.org/10.64388/IREV10I3-1723365

historical baselines. These items must be versioned, have their integrity verified, require separate administrative responsibilities, and include recoverable backups. Model promotion has to occur via a controlled pipeline so that a candidate model is tested under representative high-density and hostile conditions before it is put into use. Where possible, privileged access should be on a just-in-time basis, be thoroughly authenticated, and be recorded. The privileges granted to service accounts should be restricted to specific functions. Privacy controls in this layer include edge filtering, masking when it is operationally suitable, minimum retention, purpose-based access, and auditable data export.

The aspect dealing with business continuity and recovery involves security orchestration linking technical events to their impact on services. For example, the response playbook in the event of a compromise at the edge should state the coverage area affected, the neighbouring sensors available, the alternative inference capacity, the network quarantine rule, and the criteria for returning the node to service. When carrying out recovery, known-good images and verified configurations should be used rather than simply reconnecting a cleaned device; parallel test environments or digital twins should be employed to run the playbooks both before the peak seasons and after major configuration changes (Epiphaniou et al. 2023). Thus, the architecture ensures that resilience becomes a normal operating condition rather than something that depends on emergency improvisation.

The cyber-resilience lifecycle is Identify, Protect, Detect, Contain, Sustain, Recover and Learn. Identify establishes assets, owners, dependencies and trust boundaries. Protect applies secure configuration, cryptographic identity, least privilege, segmentation and resilient design. Detect correlates network, endpoint, identity, certificate, application and model telemetry. Contain limits blast radius without unnecessarily removing essential coverage. Sustain maintains a defined minimum service through alternate routes, peer edge capacity and trusted degraded modes. Recover restores known-good firmware, workloads, policies and data with integrity verification. Learn updates threat models, rules, inventories and exercises from incident evidence. Availability, integrity, latency, privacy, observability and recoverability are measured throughout the lifecycle.

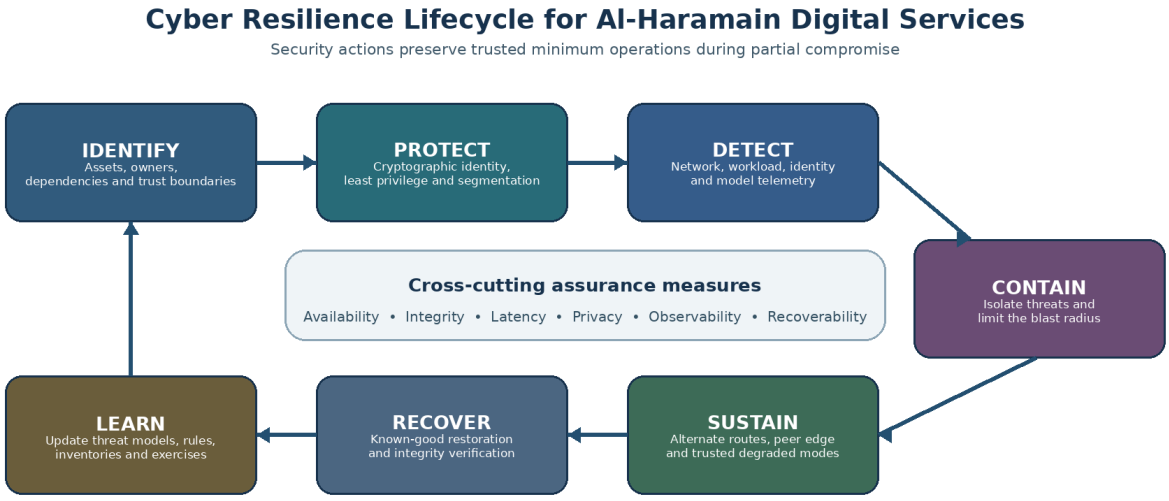

Fig. 2 Cyber-resilience lifecycle for protecting and restoring trusted Al-Haramain digital services

Table 2 maps principal cyber threats to preventive, detective, containment and recovery controls, demonstrating why a perimeter firewall alone is inadequate. Device impersonation requires cryptographic identity and attestation; denial of service requires rate control, segmentation, prioritisation and alternate paths; lateral movement requires default-deny trust zones and workload-level policy; model poisoning requires authenticated updates, robust aggregation, versioning and rollback; privileged compromise requires strong authentication, least privilege, just-in-time elevation and immutable audit; and supply-chain compromise requires signed artefacts, controlled deployment and verified restoration. Each control set defines how the service behaves after prevention fails.

Because the architecture has not been tested on the live Al-Haramain network, validation should be conducted in an authorised representative cyber range or simulation rather than inferred from literature alone. Cybersecurity and resilience measures should include asset visibility, unauthorised-path reduction, detection precision and recall, false-positive burden, mean time to detect, mean time to contain, certificate-revocation time, lateral-movement containment, failover time, recovery-point integrity and mean time to verified recovery. Service-continuity measures remain necessary but secondary: trusted coverage, information freshness, end-to-end latency and degradation during isolation. The target is bounded and observable degradation, rapid containment and restoration from known-good states rather than an unrealistic claim of zero disruption. Independent simulation should also test whether security orchestration preserves minimum trusted operations when simultaneous device, identity and network failures occur, ensuring that automated containment decisions remain reversible, auditable and proportionate to the affected service risk during peak demand. This evidence would support proportionate automation without weakening human oversight during critical incidents.

Table 2 Threat-control-resilience mapping for a smart crowd-monitoring service

Asset / function	Threat scenario	Security controls	Resilience behaviour	Validation metrics
Gate counters / sensor readers	Spoofed counts, replay, credential misuse	Unique device identity; mTLS; secure time; plausibility checks; cross-sensor fusion	Down-weight or isolate suspect source while adjacent verified sensors preserve inflow/outflow estimate	Provenance coverage; detection latency; count divergence; trusted-zone coverage
Smart cameras	Camera takeover, stream substitution, firmware compromise	Secure boot; signed firmware; certificate identity; encrypted streams; attested gateways	Quarantine affected camera and switch fusion to neighbouring cameras/counters with explicit confidence reduction	Stream integrity; isolation time; confidence loss; coverage continuity
Edge/MEC nodes	Malware, privilege escalation, model tampering	Workload identity; hardened runtime; attestation; model signatures; minimal services	Fail over to peer edge capacity or local backup model; buffer essential events	Failover time; inference freshness; resource headroom; verified restart time
5G/fibrez transport	DDoS, route disruption, unauthorized lateral movement	Micro segmentation; QoS; rate control; alternate paths; network telemetry	Preserve priority crowd flows, reroute around failed segment, isolate offending zone	Packet loss; convergence time; end-to-end latency; blast-radius reduction

Asset / function	Threat scenario	Security controls	Resilience behaviour	Validation metrics
Analytics / data fusion	Poisoned input, adversarial drift, service outage	Input trust scoring; ensemble validation; protected pipelines; rollback	Continue with verified subset and last-known safe baseline while degraded state is declared	Estimate error under isolation; model integrity; service recovery time
Identity / administration	Credential theft, excessive privilege, malicious policy change	Strong MFA; least privilege; just-in-time elevation; separation of duties; audit	Revoke session, freeze high-risk changes, restore verified policy baseline	Privilege exposure time; revocation time; unauthorized change rate
Federated security learning	Poisoned model update, malicious participant	Authenticated updates; robust aggregation; anomaly scoring; versioned rollback	Reject suspect updates and retain last verified global model	Poisoning detection rate; model-drift bound; rollback time
Operator dashboards / guidance	False visualization, data staleness, interface outage	Signed data feeds; freshness indicators; redundant displays; role-based access	Expose confidence and freshness; shift to verified alternate dashboard or manual procedure	Dashboard availability; information age; operator acknowledgement time

5.1 Proposed cyber-range/simulation validation protocol

The conceptual/reference architecture should be empirically evaluated in an authorised representative cyber range before it is treated as a validated framework. The testbed should reproduce the logical roles of field sensors and cameras, gateways, edge or multi-access edge-computing nodes, segmented 5G and fibre transport, identity services, security analytics, data and model services, operator dashboards, and backup/recovery services using synthetic or non-personal data. It does not need to duplicate confidential Al-Haramain topology; instead, it should preserve the functional dependencies, trust boundaries, redundancy paths and minimum-service assumptions required to exercise the proposed controls. Baseline runs should establish normal latency, trusted coverage, false-positive rates, failover behaviour and recovery times before adversarial scenarios are introduced, consistent with the use of digital twins for cyber-effects and resilience assurance (Epiphaniou et al. 2023).

Validation scenarios should include device impersonation and replay, compromised firmware or edge workloads, credential theft and privilege escalation, lateral movement, distributed denial of service or route disruption, time-synchronisation manipulation, poisoned model updates and supply-chain tampering. Each scenario should compare an unmitigated baseline with the proposed control stack and record mean time to detect, mean time to contain, unauthorised-path reduction, blast radius, failover time, trusted-service coverage, information freshness, false-positive burden, recovery-point integrity and mean time to verified recovery. Compound scenarios should test whether containment remains reversible and auditable while minimum trusted monitoring continues. Meeting predefined thresholds would provide empirical support for the architecture; until such testing is completed, the contribution remains a literature-based conceptual/reference architecture rather than a validated cybersecurity framework.

VI. DISCUSSION AND IMPLICATIONS

The principal finding is that cyber resilience must be engineered across the entire Al-Haramain digital infrastructure rather than added as a feature of individual monitoring devices. Research often evaluates intrusion detection, zero trust, edge computing and recovery separately, yet operational

resilience depends on their interaction. Accurate detection without segmentation permits lateral movement; segmentation without redundancy can cause self-inflicted outages; redundancy without integrity checking can replicate manipulated data; and backups without verification can restore compromised configurations. The proposed conceptual/reference architecture integrates these controls into a continuous assurance model to be tested empirically through the validation protocol described above.

Implementation should be risk-led. Critical assets and data flows must first be inventoried and assigned owners, trust zones and minimum service requirements. Security exercises should then test compound scenarios, including credential theft combined with edge-node compromise, route disruption during high demand, poisoning of a federated update, loss of time synchronisation and malicious policy change. These exercises should measure both attacker containment and continuity of trusted operations. Privacy remains a security requirement: edge inference, selective retention, encryption and controlled export reduce exposure of raw visual or device-linked data.

A phased programme should begin with asset discovery, secure configuration, cryptographic identity, privileged-access control, default-deny segmentation, central telemetry, immutable logging and verified backups. Behavioural analytics, federated detection and digital-twin exercises should follow only after this foundation is demonstrably effective. Machine learning cannot compensate for unknown assets, shared credentials, flat networks or untested recovery procedures. The framework is vendor-neutral and can guide security improvement without presuming replacement of the current infrastructure.

Governance should treat the architecture as a living cyber-assurance model. Changes to devices, network services, software, analytics models, suppliers or operator workflows require threat review, security testing, rollback planning and evidence that minimum operational services remain available. Security ownership should cover field technology, information technology, identity, applications, models and third-party dependencies. Regular control validation, incident exercises and post-event learning convert the reference architecture from a design document into an operational resilience programme.

VII. CYBERSECURITY RESEARCH AGENDA AND LIMITATIONS

The key limitation is that confidential operational details have been intentionally omitted. Even though the public sources mention the presence of smart sensing, artificial intelligence, and advanced connectivity, they do not give any information regarding actual network segmentation, the specific models of the devices, the security history, or the testing of the architecture, and as a result no operational claims can be made until such tests have been carried out by authorized stakeholders. A further limitation is the heterogeneity observed in the literature surveyed: the benchmark intrusion datasets and the typical Internet of Things environments do not correspond to the particular conditions of load, mobility, or safety that apply to the Grand Mosque.

A priority next step is to implement the authorised cyber-range/simulation protocol proposed in Section 5.1 using a representative Al-Haramain environment that models device, edge, 5G and fibre, identity, analytics and control-room layers without exposing operational systems or personal data. Experiments should measure detection and containment across realistic attack chains, including device impersonation, replay, edge compromise, denial of service, lateral movement, credential abuse, route manipulation, model poisoning and supply-chain tampering. Research is also needed on explainable multi-source threat detection, privacy-preserving telemetry, secure multi-modal fusion, adversarially robust federated learning, automated containment constrained by service coverage, and cryptographically verified recovery. The decisive research question is how much trusted service remains during compromise and how quickly verified operation can be restored.

VIII. CONCLUSION

This review reframes smart monitoring at the Grand Mosque as a cybersecurity and network-resilience problem. The protected asset is the Al-Haramain digital service chain spanning devices, edge

workloads, 5G and fibre transport, identities, analytics, data, models and operator interfaces. The proposed literature-derived conceptual/reference architecture applies cryptographic device and workload identity, secure boot and attestation, zero-trust segmentation, encrypted communications, integrated threat detection, least-privilege administration, protected software and model pipelines, immutable logging, resilient routing, trusted degraded modes and verified recovery. Crowd monitoring provides the operational case and safety consequence; it is not the paper's primary analytical focus. The architecture should not be interpreted as a validated cybersecurity framework until the representative cyber-range/simulation protocol has tested attack-path reduction, detection and containment time, blast radius, integrity of failover, service continuity and recovery from known-good states. This security-led approach provides a practical research basis for protecting the digital infrastructure that supports Al-Haramain operations without asserting access to confidential systems or current security configurations.

REFERENCES

[1] E. A. Felemban, F. U. Rehman, S. A. A. Biabani, A. Ahmad, A. Naseer, A. R. M. Abdul Majid, O. K. Hussain, A. M. Qamar, R. Falemban, and F. Zanjir, "Digital Revolution for Hajj Crowd Management: A Technology Survey," *IEEE Access*, vol. 8, pp. 208583–208609, 2020, doi: 10.1109/ACCESS.2020.3037396. Crossref

[2] W. Albattah, M. H. K. Khel, S. Habib, M. Islam, S. Khan, and K. A. Kadir, "Hajj Crowd Management Using CNN-Based Approach," *Computers, Materials & Continua*, vol. 66, no. 2, pp. 2183–2197, 2021, doi: 10.32604/cmc.2020.014227. Tech Science Press

[3] A. A. A. Abalkhail and S. M. A. Al Amri, "Saudi Arabia's Management of the Hajj Season through Artificial Intelligence and Sustainability," *Sustainability*, vol. 14, no. 21, Art. no. 14142, 2022, doi: 10.3390/su142114142. MDPI

[4] F. Gazzawe and M. Albahar, "Reducing traffic congestion in Makkah during Hajj through the use of AI technology," *Heliyon*, vol. 10, no. 1, Art. no. e23304, 2024, doi: 10.1016/j.heliyon.2023.e23304. ScienceDirect

[5] Saudi Press Agency, "Grand Mosque Gets Improved Connectivity, AI Applications through High-Tech Upgrade," Apr. 2, 2024. Saudi Press Agency

[6] Saudi Press Agency, "Two Holy Mosques General Authority Uses AI at Grand Mosque Gates," Jun. 11, 2025. Saudi Press Agency

[7] Z. Fan, H. Zhang, Z. Zhang, G. Lu, Y. Zhang, and Y. Wang, "A survey of crowd counting and density estimation based on convolutional neural network," *Neurocomputing*, vol. 472, pp. 224–251, 2022, doi: 10.1016/j.neucom.2021.02.103. ScienceDirect

[8] L. Deng, Q. Zhou, S. Wang, J. M. Górriz, and Y. Zhang, "Deep learning in crowd counting: A survey," *CAAI Transactions on Intelligence Technology*, vol. 9, no. 5, pp. 1043–1077, 2024, doi: 10.1049/cit2.12241. Wiley

[9] L. Yuan, Y. Chen, H. Wu, W. Wan, and P. Chen, "Crowd counting via Localization Guided Transformer," *Computers & Electrical Engineering*, vol. 104, Art. no. 108430, 2022, doi: 10.1016/j.compeleceng.2022.108430. ScienceDirect

[10] Z. Miao, Y. Zhang, Y. Peng, H. Peng, and B. Yin, "DTCC: Multi-level dilated convolution with transformer for weakly-supervised crowd counting," *Computational Visual Media*, vol. 9, pp. 859–873, 2023, doi: 10.1007/s41095-022-0313-5. Springer

[11] H. Gao, W. Zhao, D. Zhang, and M. Deng, "Application of improved transformer based on weakly supervised in crowd localization and crowd counting," *Scientific Reports*, vol. 13, Art. no. 1144, 2023, doi: 10.1038/s41598-022-27299-0. Nature

[12] J. Yi, F. Chen, Z. Shen, Y. Xiang, S. Xiao, and W. Zhou, "An Effective Lightweight Crowd Counting Method Based on an Encoder-Decoder Network for Internet of Video Things," *IEEE Internet of Things Journal*, vol. 11, no. 2, pp. 3082–3094, 2024, doi: 10.1109/JIOT.2023.3294727. IEEE

[13] C. Lin and X. Hu, "Efficient crowd density estimation with edge intelligence via structural

reparameterization and knowledge transfer," *Applied Soft Computing*, vol. 154, Art. no. 111366, 2024, doi: 10.1016/j.asoc.2024.111366. ScienceDirect

[14] R. Ahmad and I. Alsmadi, "Machine learning approaches to IoT security: A systematic literature review," *Internet of Things*, vol. 14, Art. no. 100365, 2021, doi: 10.1016/j.iot.2021.100365. ScienceDirect

[15] A. Khraisat and A. Alazab, "A critical review of intrusion detection systems in the internet of things: techniques, deployment strategy, validation strategy, attacks, public datasets and challenges," *Cybersecurity*, vol. 4, Art. no. 18, 2021, doi: 10.1186/s42400-021-00077-7. Springer

[16] S. Dina and D. Manivannan, "Intrusion detection based on Machine Learning techniques in computer networks," *Internet of Things*, vol. 16, Art. no. 100462, 2021, doi: 10.1016/j.iot.2021.100462. Crossref

[17] M. A. Alsoufi, S. Razak, M. M. Siraj, I. Nafea, F. A. Ghaleb, F. Saeed, and M. Nasser, "Anomaly-Based Intrusion Detection Systems in IoT Using Deep Learning: A Systematic Literature Review," *Applied Sciences*, vol. 11, no. 18, Art. no. 8383, 2021, doi: 10.3390/app11188383. MDPI

[18] P. Ranaweera, A. D. Jurcut, and M. Liyanage, "Survey on Multi-Access Edge Computing Security and Privacy," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 2, pp. 1078–1124, 2021, doi: 10.1109/COMST.2021.3062546. IEEE

[19] V. Mothukuri, P. Khare, R. M. Parizi, A. Pouriyeh, A. Dehghantanha, and G. Srivastava, "Federated-Learning-Based Anomaly Detection for IoT Security Attacks," *IEEE Internet of Things Journal*, vol. 9, no. 4, pp. 2545–2554, 2022, doi: 10.1109/JIOT.2021.3077803. IEEE

[20] E. Marmol Campos, P. Fernandez Saura, A. Gonzalez-Vidal, J. L. Hernandez-Ramos, J. Bernal Bernabe, G. Baldini, and A. Skarmeta, "Evaluating Federated Learning for intrusion detection in Internet of Things: Review and challenges," *Computer Networks*, vol. 203, Art. no. 108661, 2022, doi: 10.1016/j.comnet.2021.108661. Crossref

[21] S. Agrawal, S. Sarkar, O. Aouedi, G. Yenduri, K. Piamrat, M. Alazab, S. Bhattacharya, P. K. R. Maddikunta, and T. R. Gadekallu, "Federated Learning for intrusion detection system: Concepts, challenges and future directions," *Computer Communications*, vol. 195, pp. 346–361, 2022, doi: 10.1016/j.comcom.2022.09.012. Crossref

[22] M. Sarhan, W. W. Lo, S. Layeghy, and M. Portmann, "HBFL: A hierarchical blockchain-based federated learning framework for collaborative IoT intrusion detection," *Computers & Electrical Engineering*, vol. 103, Art. no. 108379, 2022, doi: 10.1016/j.compeleceng.2022.108379. Crossref

[23] Z. Zhang, Y. Zhang, D. Guo, L. Yao, and Z. Li, "SecFedNIDS: Robust defense for poisoning attack against federated learning-based network intrusion detection system," *Future Generation Computer Systems*, vol. 134, pp. 154–169, 2022, doi: 10.1016/j.future.2022.04.010. Crossref

[24] S. Li, M. Iqbal, and N. Saxena, "Future Industry Internet of Things with Zero-trust Security," *Information Systems Frontiers*, vol. 26, pp. 1653–1666, 2024, doi: 10.1007/s10796-021-10199-5. Springer

[25] Zanasi, S. Russo, and M. Colajanni, "Flexible zero trust architecture for the cybersecurity of industrial IoT infrastructures," *Ad Hoc Networks*, vol. 156, Art. no. 103414, 2024, doi: 10.1016/j.adhoc.2024.103414. Crossref

[26] S. Ameer, L. Praharaj, R. Sandhu, S. Bhatt, and M. Gupta, "ZTA-IoT: A Novel Architecture for Zero-Trust in IoT Systems and an Ensuing Usage Control Model," *ACM Transactions on Privacy and Security*, vol. 27, no. 3, Art. no. 22, pp. 1–36, 2024, doi: 10.1145/3671147. ACM

[27] G. Epiphaniou, M. Hammoudeh, H. Yuan, C. Maple, and U. Ani, "Digital twins in cyber effects modelling of IoT/CPS points of low resilience," *Simulation Modelling Practice and Theory*, vol. 125, Art. no. 102744, 2023, doi: 10.1016/j.simpat.2023.102744. Crossref

[28] T. N. I. Alrumaih, M. J. F. Alenazi, N. A. AlSowaygh, A. A. Humayed, and I. A. AlAblani,

"Cyber resilience in industrial networks: A state of the art, challenges, and future directions," *Journal of King Saud University - Computer and Information Sciences*, vol. 35, no. 9, Art. no. 101781, 2023, doi: 10.1016/j.jksuci.2023.101781. ScienceDirect

[29] General Authority for the Care and Management of the Grand Mosque and the Prophet's Mosque, "The General Authority for the Care and Management of the Grand Mosque and the Prophet's Mosque Adopts Smart Technologies for Visitor Counting and Enhancing Crowd Management at the Grand Mosque," Mar. 1, 2026. General Authority for the Care and Management of the Grand Mosque and the Prophet's Mosque

[30] Saudi Press Agency, "Smart Systems Streamline Crowds at Grand Mosque," *Arab News*, Jan. 14, 2026. Arab News