КІБЕРБЕЗПЕКА: освіта, наука, техніка
CYBERSECURITY: EDUCATION, SCIENCE, TECHNIQUE
№ 2 (34), 2026
ISSN 2663 – 4023
DOI 10.28925/2663-4023.2026.34.1368
UDC 004.056.5:004.738.5

Oleksandr Pliekhov
Independent Researcher, IEEE Senior Member, EPAM Systems, Inc., Newtown, PA, USA
ORCID:0009-0000-1343-4308
aleks.plekhov@gmail.com

Ganna Grynkevych
DSc, Professor, Department of Information and Cyber Security named after Professor Volodymyr Buriachok, Faculty of Information Technology and Mathematics, Borys Grinchenko Kyiv Metropolitan University, Kyiv, Ukraine
ORCID:0000-0003-1922-5165
ggrynkevych@ukr.net

Volodymyr Vasylenko
PhD, Associate Professor, Associate Professor of the Department of Computer Science, State University of Information and Communication Technologies, Kyiv, Ukraine
ORCID:0000-0001-8465-6178
oknelisavvova172@gmail.com

AN EXTENDED STRIDE THREAT MODEL FOR OFFLINE BLE MESH DISASTER COMMUNICATION APPLICATIONS

Abstract. Offline mesh messaging over Bluetooth Low Energy (BLE) is increasingly used as a communication layer of last resort when infrastructure fails. Its security is normally analysed at the BLE stack level, and rarely at the level of a deployed application in a disaster context. This paper presents a systematic threat model for that system class. We extend STRIDE with two categories the disaster context requires and standard STRIDE cannot express - Bystander harm and Physical compromise – and apply a reduced three-factor rating scheme (Damage x Reproducibility x Affected parties) with published anchors. The methodology is instantiated on a single case study – ResQMesh AI, an open-source Android BLE mesh emergency-communication platform – which fixes the scope of the mitigation mapping. The catalogue contains 29 threats and six adversary profiles, with a per-threat mitigation mapping in which 3 threats are effectively addressed, 19 partially and 7 not at all. The two coupling effects the model predicts are checked by discrete-time simulation of a 60-node flooding mesh (TTL 7, 20 runs): forged critical-priority traffic at three times per-node transmission capacity reduces the deadline delivery ratio of genuine critical messages from 0.87 to 0.59, a per-identity relay budget restores it to 0.85, and Sybil identity rotation defeats the budget entirely; driving 40% of nodes into the lowest energy mode reduces delivery from 0.88 to 0.50, with the critical relay floor recovering 3-5 percentage points. Three findings generalise beyond the case study: removing infrastructure relocates the entire trust bootstrap onto a trust-on-first-use radio exchange, making spoofing the highest-ranked class; AI-driven message prioritisation creates an amplification primitive, since forged high-priority traffic pre-empts genuine emergency traffic; and energy-aware relaying gives an adversary a low-cost lever for partitioning the mesh.

Keywords: threat modeling; STRIDE; Bluetooth Low Energy; BLE mesh; disaster communication; offline messaging; bystander privacy; denial of service; ResQMesh AI.

INTRODUCTION

Disasters disable the communication infrastructure that response depends on, under severe time pressure: across 34 earthquakes, most live extrications occur early and survival probability falls steeply with time to rescue [1]. Offline BLE mesh messaging addresses this directly, needing no additional hardware and forming a multi-hop network from the devices already present. Bridgefy was adopted during network shutdowns in several countries [2], and BitChat, released in July 2025, has seen comparable use [3].

The security record is poor, and its failures are instructive. Bridgefy, as analysed, permitted tracking, offered no authenticity and no effective confidentiality, and could be disabled network-wide by one crafted message [2]. BitChat was released with an explicit statement in its repository that it had not undergone external

security review [3], and an independent researcher reported identity-authentication weaknesses permitting impersonation of trusted contacts [4]. Neither is a failure of BLE as a radio technology; both occurred at the application and trust-establishment layer, which the BLE security literature covers least.

This is the paradox of infrastructure-absent operation. Removing servers removes the central database, the operator and the single point of mass compromise. It also removes the authority that binds identities, leaving trust-on-first-use over a link the adversary can hear and transmit on; it removes the monitoring plane, since no node can observe a flood or a Sybil population; it weakens the assumption that a user retains physical control of their device; and it implicates people who never installed the application but appear in the photographs and location reports others send.

Purpose of the article and contributions. (1) A systematic threat classification for offline BLE mesh disaster applications – 29 threats in eight categories – at the application and trust-establishment layer. (2) An extension of STRIDE with two categories required by the deployment context, plus a reduced, explicitly anchored rating scheme. (3) A mapping of the catalogue to the mitigations of a real open-source system, with per-threat effectiveness and residual gaps. (4) A simulation of the two coupling effects the model predicts – prioritisation with availability (T5.5) and energy policy with connectivity (T5.6) – including a countermeasure and its failure mode. (5) Design implications for the class.

Object, subject, novelty and practical significance. The object of the research is the security of offline BLE mesh communication applications in disaster deployment; the subject is the threat model of that class and its instantiation on a concrete system. The scientific novelty consists in (a) the extension of STRIDE with two elicitation categories that the disaster context requires, (b) the identification of three threats specific to AI-assisted mesh messaging (T2.4, T5.5, T5.6) and one specific to dual-transport architectures (T4.6), and (c) the demonstration that a defence against one of them is defeated by another (per-identity budget by Sybil identities), which couples the two highest-ranked threat classes. The practical significance is a reusable catalogue with per-threat mitigation status that a developer of this class can apply at design time, and six design implications derived from it.

ANALYSIS OF RECENT RESEARCH AND PUBLICATIONS

BLE stack and link layer. Surveys map the BLE vulnerability landscape across pairing, connection and privacy mechanisms [5, 6, 13]. Specific results include impersonation and method-confusion attacks on pairing [14], spoofing against reconnection [15], design flaws found by formal analysis [16], discovery of non-discoverable devices [17], and state-aware defence frameworks [18]. Devices remain fingerprintable from GATT profiles despite address randomisation [19, 36], standards-compliant man-in-the-middle attacks remain feasible against current consumer devices [20], and jamming is well characterised generally [21], for BLE beacons specifically [22, 23], and for constrained IoT deployments in the Ukrainian literature [35].

Comparable wireless-security lifecycles. A similar trajectory is visible in a neighbouring wireless technology: Wi-Fi authentication moved from the badly broken WEP through WPA and WPA2 to WPA3 only after each generation's weaknesses were exploited in deployment, a lifecycle documented in detail for financial-sector networks [38]. Offline BLE mesh messengers are earlier on the same curve: Bridgefy and BitChat are closer to WEP-era assumptions than to an audited, multi-generation security architecture, and this class has not yet had the hardening cycle Wi-Fi authentication received.

Threat models for BLE. Skallak and Schmidt applied STRIDE to BLE, producing a threat landscape from packet sniffing to machine-in-the-middle mapped onto the stack [24]. This is the closest prior work methodologically, and its scope is the difference: it models the protocol stack, not an application, a deployment context, or the non-users a disaster application affects.

Mesh and ad hoc networks. For multi-hop networks, the classical treatment of routing attacks applies directly to flooding meshes [25], and Douceur's result that Sybil attacks are unavoidable without a logically centralised identity authority is a structural constraint on this class rather than an implementation defect [26]. In standardised Bluetooth Mesh, the friendship mechanism permits denial of service reducing node battery life from years to days [27], and its provisioning has been subject to impersonation vulnerabilities [33]; application-layer meshes such as the one studied here do not use Bluetooth Mesh Profile, and so inherit its link-layer exposure but not its security architecture [7, 8]. Deployed messengers. Analyses of deployed messengers are ad hoc rather than systematic [2, 4]; the BitChat reports describe its July 2025 state and the upstream transport has since been revised, so we treat them as evidence about a class of failure rather than a description of the current build. Prior work by the present author on on-device prioritisation and transport selection is functional rather than adversarial, assuming benign inputs [28, 29]. Deferrable authentication, which allows a session to be established immediately and authenticated later, is a recent proposal directly relevant to the trust-on-first-use problem of this class [37].

Methodology. Methodologically, STRIDE classifies threats by walking a data-flow model [9], LINDDUN applies the same approach to privacy [10], a comparison of six methods concludes that selection should follow the analytical objective [11], and STRIDE has been adapted to IoT risk analysis in standard-compliant form [12]. We adopt STRIDE because it is data-flow driven, developer-facing and extensible by construction; its blind spots – the privacy of non-participants and the physical custody of the device – are what our two additional categories address.

Comparison with the closest prior work. Table 1 positions the present model against the four analyses it builds on.

Table 1

Comparison with prior threat analyses relevant to this system class

Criterion	Karlof & Wagner [25]	Deng et al.(LINDDUN) [10]	Albrecht et al. [2]	Skallak & Schmidt [24]	This work
Modelling level	Routing protocols of sensor networks	Software data flows (privacy)	One deployed application (Bridgefy)	BLE protocol stack	Application, trust bootstrap and deployment context
Method	Attack taxonomy	DFD-driven, LINDDUN categories	Reverse engineering, practical attacks	DFD-driven STRIDE	DFD-driven extended STRIDE (+2 categories)
Categories	Routing attacks	7 privacy types	Not categorised	6 STRIDE	8 (STRIDE + Bystander harm + Physical compromise)
Non-user parties	No	Partly (data subjects)	No	No	Yes (explicit category)
Physical custody	No	No	No	No	Yes (explicit category)
AI/prioritisation logic	No	No	No	No	Yes (T2.4, T5.5)
Energy policy as attack surface	No	No	No	No	Yes (T5.6)
Rating scheme	No	No	No	No	D x R x A with published anchors, sensitivity analysis
Mitigation mapping	Countermeasures per attack	PETs per threat type	Recommendations	Not central	Per threat, with effectiveness and residual gap
Empirical component	No	No	Yes (attacks demonstrated)	No	Simulation of two threats and one countermeasure

Gap. The literature covers the BLE stack in depth, threat modeling methodology in general, individual applications non-systematically, and the functional design of AI-assisted emergency messaging. Missing is a systematic, reusable threat model for offline BLE mesh disaster applications at the level where these systems actually fail: trust bootstrap, flooding behaviour, prioritisation logic, attachment handling, transport selection and device custody.

THEORETICAL FOUNDATIONS OF THE STUDY

Target class. A mobile application that discovers peers over BLE without infrastructure, forwards messages over multiple hops by flooding with duplicate suppression, establishes end-to-end sessions peer-to-peer, and applies on-device processing to the traffic it originates or relays.

Case study. ResQMesh AI is an open-source Android platform derived from BitChat [3, 34]; the analysis refers to the develop branch of the repository as of August 2026 (commit: [to be inserted by the author at submission]). The inherited transport provides Noise sessions in the XX handshake pattern with X25519 and ChaCha20-Poly1305 [31, 32], peer identities derived from static keys, TTL-limited relay bounded at seven hops with fragmentation and duplicate suppression, adaptive duty cycling, optional password-protected channels using Argon2id and AES-256-GCM, and an emergency wipe that clears application data. The transport is dual: alongside the BLE mesh it offers internet-backed paths (Nostr relays with ephemeral per-area keys, Tor, geohash-addressed location channels) and a Wi-Fi Aware link, with automatic transport selection. ResQMesh AI adds five security-relevant components: a two-stage message priority classifier (keyword rules, then a quantised neural model), offline speech recognition, a FEMA ICS-213 report generator, a battery-aware relay policy with four energy modes and a minimum relay probability for critical traffic, and an image-to-text scene analyser that classifies a photograph locally and sends a short description instead of the image. A radio-layer priority queue schedules critical packets ahead of routine traffic.

Deployment context. C1 – infrastructure absent: no Internet, cellular, certificate authority, time source or revocation service. C2 – ad hoc trust: peers are often strangers with no out-of-band verification channel. C3 – physical accessibility: devices are lost, seized, shared and unlocked under duress. C4 - time criticality: defences that delay critical traffic are unacceptable [1]. C5 – energy scarcity: battery exhaustion is a network-availability threat [27]. C6 – multi-hop propagation: every node relays, and a sender controls neither the path nor its visibility. C7 – non-user parties appear in attachments and reports.

Assets. A1 message confidentiality; A2 integrity; A3 sender authenticity; A4 user privacy (identity, location); A5 device availability; A6 network availability; A7 bystander privacy [30]; A8 coordination integrity. A8 is deliberately separated from A2: individually authentic messages can still compose into a false operational picture if an adversary controls their distribution across categories or geography.

Trust assumptions. T-A1: device and OS are not compromised in advance; malware and supply-chain attacks are out of scope. T-A2: the primitives above are not broken, but attacks on their composition and implementation remain in scope, since that is where deployed mesh messengers have failed [2, 4]. T-A3: the user controls their device at the moment of use, but not necessarily afterwards. T-A4: the adversary can receive and transmit on BLE channels within range of one or more nodes, and can position several radios. T-A5: nation-state-scale and quantum-capable adversaries are excluded, since including them yields no actionable design guidance here. T-A6 (transport scoping): the model covers the offline BLE path – the path that operates under C1 – plus the consequences of automatic transport selection for that path (T4.6); relay-side threats on the internet-backed paths are separate work.

Extended STRIDE Methodology and Rating Scheme. Data-flow model. Fig. 1 shows the data-flow diagram used for elicitation: four trust boundaries (sender device, BLE radio and relay nodes, receiver device, internet-backed transports), the processes that ResQMesh AI adds to the transport, the two local stores, and the three external entities – user, coordinator and bystander – the last of which is the element standard STRIDE does not model. Each threat in Table 2 is annotated at the element where it was elicited.

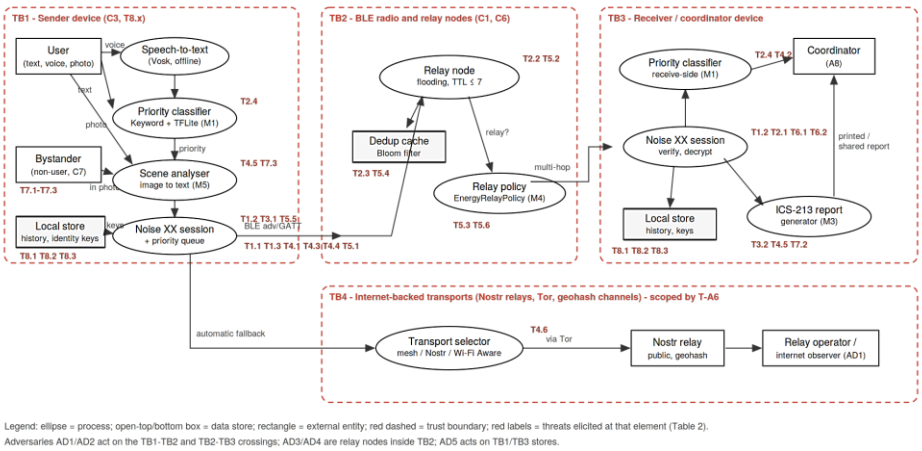

Fig. 1. Data-flow diagram of the case-study system with trust boundaries (red, dashed) and the threat identifiers elicited at each element.

Extended STRIDE. We retain the six STRIDE categories and add two. Bystander harm covers threats whose damaged party is neither a user nor a principal in the data-flow model, and therefore cannot be elicited by standard STRIDE. It is the disaster analogue of LINDDUN's disclosure and identifiability types, with the decisive difference that the affected party cannot consent, cannot be notified and cannot act, so the consent-signalling mechanisms developed for wearable cameras [30] are unavailable and on-device data minimisation is the only control. Physical compromise covers threats arising from possession of a device or coercion of its holder; STRIDE's elevation-of-privilege category assumes a logical boundary, whereas here the boundary is physical and is crossed routinely (C3). Both satisfy the standard criterion for a STRIDE category: each is elicited by a distinct question asked of each model element, and each maps to a distinct class of countermeasure.

Rating. We use a reduced DREAD variant with three factors scored 1-3 against explicit anchors, risk R = D x R x A in the range 1-27. Exploitability and Discoverability are dropped: in an open-source system discoverability is uniformly high and therefore non-discriminating, and exploitability is largely collinear with reproducibility for radio-proximate attacks. Damage: 1 = degraded convenience; 2 = loss of confidentiality, integrity or availability for some traffic; 3 = plausible contribution to physical harm, or loss of the channel itself. Reproducibility: 1 = requires rare conditions or a specific target state; 2 = repeatable with commodity hardware and moderate effort; 3 = repeatable at will with public tooling. Affected parties: 1 = a single user; 2 = a group, channel or locality; 3 = the mesh as a whole, or non-users at scale. Bands: Low 1-6, Medium 7-14, High 15-27. The scheme is ordinal, not cardinal, and expert-assigned; the anchors are published so that a second analyst can reproduce or contest each score [11], and its robustness to scoring disagreement is quantified in Section 5.4.

Adversaries. AD1 passive external observer (receives BLE traffic in range); AD2 active external attacker (crafted advertisements, connection requests, jamming; not a channel member); AD3 malicious mesh participant (joined legitimately; originates, relays, drops); AD4 compromised legitimate node (retains accumulated trust); AD5 physical attacker (has the device, or coerces its holder); AD6 advanced or nation-state adversary – out of scope (T-A5).

METHODOLOGY OF THE STUDY

The two threats that the model ranks highest among those specific to this class – priority starvation (T5.5) and induced energy scarcity (T5.6) – were simulated in a discrete-time model of a flooding mesh. Sixty nodes are placed uniformly in a unit square with communication radius 0.19 (mean degree about six; topologies that are not connected are rejected). Each node transmits at most one packet per slot from a two-level priority queue (CRITICAL before NORMAL, FIFO within level), forwards each message at most once (duplicate suppression) with TTL = 7, and applies a relay probability per energy mode. Background NORMAL traffic is generated at 0.05 messages per node per slot; four honest sources emit genuine CRITICAL messages at 0.01 per slot. The metric is the deadline delivery ratio (DDR): the mean, over genuine CRITICAL messages, of the fraction of nodes that receive the message within 120 slots of origination. Each configuration is run 20 times with fixed seeds (topology and traffic streams seeded independently); medians are reported with 10th-90th percentile bands, rounded half-up to two decimals. Because the honest-source set is sampled to exclude attacker nodes, the no-attack baselines of the two experiments differ slightly (0.87 in Table 3, 0.88 in Table 4). Parameters taken from the case-study documentation are TTL = 7, the four energy modes with NORMAL relay multipliers 1.00/0.75/0.50/0.25, and the CRITICAL relay floor of 0.20 at ULTRA_LOW_POWER; all other values are modelling assumptions. The simulator is a 150-line Python script released with the paper. The model is deliberately abstract – it has no PHY-layer collisions, no connection-interval effects and no mobility – so absolute values should not be read as predictions; the object of interest is the shape of the response and the relative effect of the countermeasure.

RESULTS OF THE STUDY

Threat Catalogue and Mitigation Mapping. Table 2 combines the catalogue with the mapping to the case-study system. Effectiveness is E (addressed under the stated assumptions), P (reduced, meaningful residual remains) or N (not addressed).

Table 2

Threat catalogue and mitigation mapping (D = damage, R = reproducibility, A = affected parties; Risk = D x R x A)

ID	Threat	Adv.	D	R	A	Risk	Existing mitigation	Eff.	Residual gap
	Spoofing								
T1.1	Sybil attack: one entity presents many identities to dominate a locality	AD2, AD3	3	3	3	27	None specific	N	No identity cost, no per-identity rate limiting
T1.2	Identity spoofing against trust-on-first-use	AD2, AD3	3	3	2	18	Noise XX mutual authentication, forward secrecy, identity binding (NoiseSessionManager)	P	First contact unverified; no documented out-of-band fingerprint workflow [37]; no independent review [4]
T1.3	Advertisement spoofing to attract connections or pollute the peer list	AD2	2	3	2	12	Session establishment required before exchange	P	Advertising layer unauthenticated by construction
	Subtotal Spoofing: 3 threats, mean risk 19.0								
	Tampering								
T2.1	Modification of message content by a relay	AD3, AD4	3	1	2	6	End-to-end AEAD (ChaCha20-Poly1305) [32]	E	Public-channel traffic not covered
T2.2	Peer/route table poisoning to attract or divert traffic	AD3, AD4	2	2	3	12	Flooding requires no trusted route state	P	Peer entries accepted without corroboration
T2.3	Deduplication-state manipulation causing re-flooding	AD3	2	2	3	12	Duplicate-suppression cache (Bloom filter), TTL ≤ 7 hops	P	Behaviour under adversarial input not characterised
T2.4	Classifier-input manipulation: crafting text to obtain a chosen category or priority	AD3	2	3	2	12	Two-stage classifier (KeywordMessageClassifier then TFLiteMessageClassif	N	Classification applied to attacker-chosen text; no per-sender limit on high-priority outcomes

ID	Threat	Adv.	D	R	A	Risk	Existing mitigation	Eff.	Residual gap
							ier via Composite MessageClassifier)		
	Subtotal Tampering: 4 threats, mean risk 10.5								
	Repudiation								
T3.1	Sender denies originating a critical message	AD3	2	3	1	6	Sender authentication within a session	P	No signed, retained record of origination
T3.2	No audit trail for post-incident reconstruction	AD3, AD4	2	3	2	12	ICS213ReportGenerator captures sender, time, category, body	P	Locally generated, not tamper-evident
	Subtotal Repudiation: 2 threats, mean risk 9.0								
	Information disclosure								
T4.1	Passive eavesdropping on broadcast and public-channel traffic	AD1	2	3	3	18	End-to-end encryption for direct messages	P	Public-channel traffic readable in range by design
T4.2	Traffic analysis despite encryption	AD1	2	2	3	12	None	N	Priority markings are themselves a side channel
T4.3	Location inference from signal strength and repeated proximity	AD1, AD2	3	2	2	12	None	N	No transmit-power management or timing randomisation
T4.4	Metadata linkability across time and place [19, 36]	AD1	2	3	3	18	Ephemeral transport identifiers	P	Application-level nickname and key fingerprint remain stable
T4.5	Personal data in attachments and exported situation reports	AD1, AD3, AD5	3	2	2	12	On-device scene analysis (VisionTFLiteClassifier, SceneToEmergencyMapper) replaces	E	Exported reports and locally retained media not covered

ID	Threat	Adv.	D	R	A	Risk	Existing mitigation	Eff.	Residual gap
							the image with a short text		
T4.6	Transport-selection disclosure: automatic routing moves traffic to public relays and geohash-addressed channels, taking it outside the local trust boundary	AD1, AD3	3	3	2	18	Ephemeral keys per geohash area; Tor for relay connectivity; mesh preferred when peers are reachable	P	Coarse location is inherent in geohash addressing; relay operators observe timing and volume; the transport actually used is not prominent to the user
	Subtotal Information disclosure: 6 threats, mean risk 15.0								
	Denial of service								
T5.1	Radio-layer jamming [21, 22]	AD2	3	2	3	18	None available at the application layer	N	Physical-layer threat; only detection and notification are feasible
T5.2	Flooding the mesh with valid-looking messages	AD2, AD3	3	3	3	27	TTL limit, deduplication, adaptive relay probability	P	No admission control, no per-sender quota
T5.3	Battery drain via sustained connection requests or induced relaying [27]	AD2, AD3	3	2	3	18	EnergyRelayPolicy (four EnergyMode tiers), adaptive duty cycling	P	Reduces self-inflicted drain; does not detect induced drain
T5.4	Saturation of the duplicate-suppression cache	AD3	2	2	3	12	Bounded cache	P	Eviction behaviour under attack unspecified
T5.5	Priority starvation: forged critical traffic pre-empts genuine critical traffic	AD3	3	3	3	27	Radio-layer priority queue (Bluetooth PacketBroadcaster, priority set by KeywordMessageClassifier)	N	Priority is sender-asserted, unauthenticated and unlimited across the mesh
T5.6	Induced energy	AD3	3	2	3	18	EnergyRel	P	Bounds damage; does

ID	Threat	Adv.	D	R	A	Risk	Existing mitigation	Eff.	Residual gap
	scarcity: driving neighbours into low-power modes to partition the mesh						ayPolicy critical relay floor ≥ 0.20 at ULTRA_LOW_POWER		not detect the inducement
	Subtotal Denial of service: 6 threats, mean risk 20.0								
	Elevation of privilege								
T6.1	Offline dictionary attack on a password-protected channel key	AD1, AD2	2	1	2	4	Argon2id key derivation, AES-256-GCM	P	Weak user-chosen passwords remain attackable; no rate limiting is possible offline
T6.2	Channel or group takeover after key compromise	AD3, AD4	2	1	2	4	Per-session keys limit blast radius	P	No revocation path without infrastructure
	Subtotal Elevation of privilege: 2 threats, mean risk 4.0								
	Bystander harm								
T7.1	Non-consenting individuals captured in attachments	any sender (no adversary required)	2	3	2	12	Image classified on device (ImageSceneAnalyzer) and discarded; only a category description is sent	E	Original media remain in the device gallery
T7.2	Disclosure of the location of vulnerable populations through aggregated reports	AD1, AD3	3	2	2	12	None	N	Aggregation in situation reports is not privacy-reviewed
T7.3	Residual identifying detail in AI-generated scene descriptions	AD1, AD3	2	2	2	8	Fixed category vocabulary limits free text	P	Fallback path emits generic labels not analysed for identifying content
	Subtotal Bystander harm: 3 threats, mean risk 10.7								
	Physical compromise								

ID	Threat	Adv.	D	R	A	Risk	Existing mitigation	Eff.	Residual gap
T8.1	Device theft or seizure during the event	AD5	3	2	2	12	OS storage encryption; emergency wipe (triple-tap)	P	Wipe requires the user to act in time; application-level at-rest encryption not documented
T8.2	Coercion of the holder to unlock or disclose keys	AD5	3	1	2	6	Emergency wipe usable as a duress action	P	No deniable state or decoy profile; refusal remains observable
T8.3	Forensic recovery after user-initiated deletion	AD5	3	2	1	6	Emergency wipe removes application data	P	Flash-level residue not characterised
	Subtotal Physical compromise: 3 threats, mean risk 8.0								

Summary. 29 threats in eight categories: three effectively addressed, nineteen partially, seven not at all. The highest-scoring threats – T1.1, T5.2 and T5.5 at 27 – are unaddressed or only partially addressed, a pattern that holds for every deployed system of this class publicly analysed to date [2, 4]. Content modification (T2.1) is the one threat the class handles well, because end-to-end authenticated encryption makes silent modification by a relay infeasible under T-A2; the residual tampering risk lies in state kept outside the authenticated envelope – peer tables, deduplication state – and in the classifier input, a target that requires no adversarial-machine-learning apparatus, only knowledge of rules that are public in an open-source system. The scores in the bystander category understate its ethical weight, which we state rather than adjust: the third factor counts scale, so a single non-consenting person scores 2, while the harm to that individual may be severe and irreversible.

Priority Starvation (T5.5). Three attacker nodes inject forged CRITICAL messages at an aggregate load L of 0-3 messages per slot, expressed relative to the per-node transmission capacity of one packet per slot (Fig. 2a). Three conditions are compared: no defence; a per-identity relay budget in which each relay forwards at most two CRITICAL messages per source identity per 50-slot window; and the same budget against an attacker that presents a fresh identity for every message (Sybil, T1.1).

Table 3

Deadline delivery ratio of genuine CRITICAL traffic under forged-priority load (median of 20 runs; 10th-90th percentile in parentheses)

Forged load L (msgs/slot)	No defence	Per-identity budget	Budget + Sybil identities
0 (baseline)	0.87 (0.71-0.97)	–	–
0.5	0.86 (0.60-0.97)	0.85 (0.68-0.96)	0.82 (0.60-0.97)
1.0	0.87 (0.69-0.95)	0.84 (0.71-0.89)	0.83 (0.63-0.92)
1.5	0.81 (0.59-0.88)	0.87 (0.67-0.94)	0.80 (0.56-0.87)
2.0	0.73 (0.42-0.86)	0.82 (0.67-0.95)	0.68 (0.42-0.86)
3.0	0.59 (0.38-0.78)	0.85 (0.67-0.94)	0.59 (0.34-0.75)

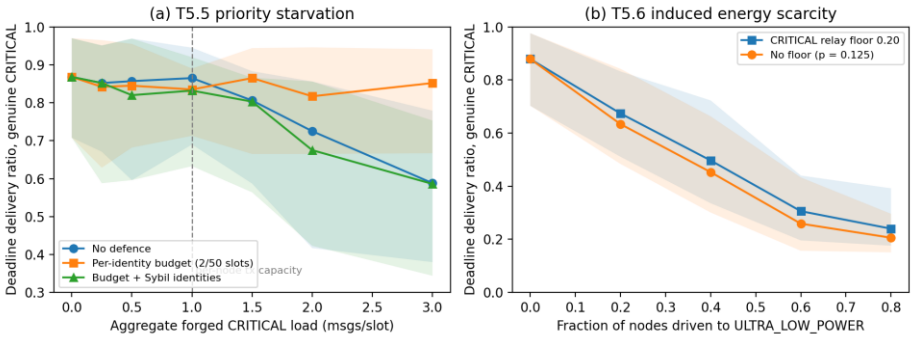

Fig. 2. (a) Deadline delivery ratio of genuine CRITICAL messages as a function of aggregate forged CRITICAL load, with and without a per-identity relay budget; (b) deadline delivery ratio as a function of the fraction of nodes driven into ULTRA_LOW_POWER, with and without the CRITICAL relay floor. Shaded bands: 10th-90th percentiles over 20 runs.

Three properties of the response are relevant to the threat model. First, the effect is a threshold, not a gradient: below per-node capacity the priority queue absorbs forged traffic without measurable cost to genuine traffic, and degradation begins only when forged CRITICAL load exceeds what a node can transmit per slot. This is consistent with the mechanism – a priority queue starves lower priorities, and forged traffic starves genuine traffic of the same priority only when the CRITICAL level itself saturates. Second, above the threshold the degradation is severe: at three times capacity, genuine CRITICAL delivery falls by a third (0.87 to 0.59), and the lower percentile falls to 0.38, meaning that in the worst topologies most of the mesh does not receive a genuine rescue request within the deadline. Third, a per-identity budget enforced at relays holds DDR at 0.82-0.87 across the entire load range, but only while the attacker keeps a stable identity: with identity rotation the budget provides no benefit at any load (0.59 at L = 3, identical to no defence). This is the coupling stated as F2 in Section 5.5: the defence against T5.5 is only as strong as the defence against T1.1.

Induced Energy Scarcity (T5.6). A fraction f of nodes is placed in ULTRA_LOW_POWER, where NORMAL relay probability is 0.125 (the 0.25 multiplier scaled by an assumed network factor of 0.5 for a mesh of this density) and CRITICAL relay probability is either the documented floor of 0.20 or, without the floor, the same 0.125 (Fig. 2b).

Table 4

Deadline delivery ratio of genuine CRITICAL traffic as a function of the fraction of nodes in ULTRA_LOW_POWER

Fraction *f* in ULTRA_LOW_POWER	With CRITICAL floor 0.20	Without floor
0.0	0.88 (0.70-0.98)	0.88 (0.70-0.98)
0.2	0.67 (0.51-0.83)	0.63 (0.48-0.84)
0.4	0.50 (0.34-0.72)	0.45 (0.30-0.66)
0.6	0.31 (0.20-0.44)	0.26 (0.16-0.43)
0.8	0.24 (0.18-0.39)	0.21 (0.15-0.30)

Delivery degrades approximately linearly up to f = 0.6 and then flattens as delivery reduces to direct neighbours: one fifth of the mesh in the lowest energy mode costs a quarter of genuine CRITICAL delivery (0.88 to 0.67), and two fifths cost 43% (0.88 to 0.50). The floor helps consistently but modestly – three to five percentage points at every level – which is the expected behaviour of a control that bounds the per-node relay probability but not the number of nodes affected. The attacker's lever is the latter: the cost of the attack is the energy needed to push nodes across a policy threshold, which is lower than the cost of draining them, and the floor does not raise it.

Sensitivity of the Rating Scheme. To quantify the effect of scoring disagreement on the catalogue's conclusions, each of the 87 factor scores in Table 2 was perturbed by ±1 independently with probability p, 2,000 times per p, and the perturbed ranking compared with the original. At p = 0.1 (one judgement in ten contested), Spearman rank correlation with the original ordering is 0.90 (10th percentile 0.83), band membership agrees for 87% of threats, and the set of High-band threats has mean Jaccard overlap 0.82 with the original. At p = 0.2 the corresponding values are 0.81, 76% and 0.69; at p = 0.33 they are 0.71, 65% and 0.57. The ranking is therefore stable under realistic disagreement, but individual band assignments are not: a single ±1 on one factor moves 26 of 29 threats across a band boundary, because the multiplicative scale has only nine distinct values. The paper's

conclusions are accordingly drawn from the ordering and from the High set – which at p = 0.1 retains the three top-ranked threats in 70% of perturbations – and from the simulated behaviour of T5.5 and T5.6, not from the exact band of any single threat.

Findings and Design Implications.

F1. Removing infrastructure relocates the trust problem rather than removing it. The design eliminates the server-side surface entirely and concentrates the remainder at one point: the trust-on-first-use exchange. Every publicly analysed system of this class has failed there [2, 4]. The security of an offline mesh messenger is therefore determined almost entirely by its identity model – the one part that cannot be borrowed from standard practice, because standard practice assumes an authority.

F2. Prioritisation and availability are coupled. Priority marking exists to prevent a rescue request from queuing behind routine traffic, and the case-study system implements it at both the classification and the radio-scheduling layer. But a priority attribute asserted by the sender and honoured by every relay is an amplification primitive: an adversary who marks traffic critical obtains preferential relaying, scheduling and coordinator attention at the direct expense of genuine critical traffic (T2.4, T5.5). The simulation makes the shape of the effect concrete: no cost below per-node capacity, then a loss of a third of genuine critical delivery at three times capacity (Section 5.2). Any mechanism that allocates a scarce shared resource on an unauthenticated, sender-asserted attribute converts that attribute into an attack surface. The functional literature does not address this [28, 29].

F3. Energy policy and connectivity are coupled. Reducing relay probability as battery declines is correct engineering against self-inflicted drain. Against an adversary it is a lever with mechanical advantage (T5.6): the attacker need not exhaust a device, only push it across a policy threshold, after which it withdraws from relaying voluntarily. A relay-probability floor for critical traffic bounds the worst case but does not remove the incentive: in simulation the floor recovers three to five percentage points while the attack costs about 20 points per fifth of the mesh affected (Section 5.3).

F4. The two added categories are the least defended. They contain no top-scoring threats but the highest proportion of unmitigated ones – the predictable consequence of their absence from standard elicitation: what is not asked about is not designed for. The case-study system's image-to-text pipeline is an instructive exception: introduced for bandwidth reasons, it is also a strong bystander-privacy control, because data minimisation adopted for efficiency is the correct privacy architecture.

Design implications. (1) Make the identity model the first design decision: provide an out-of-band verification path and show verification status persistently in the interface. (2) Never honour a sender-asserted priority without a cost – rate-limit critical traffic per identity, enforced at relays rather than only at the originator. In simulation this holds genuine critical delivery at 0.82-0.87 under any forged load, at a cost to genuine traffic of at most 0.03 at low load, and converts T5.5 from a free attack into one requiring Sybil identities, which defeats it (Section 5.2). The budget therefore has value only together with an identity cost. (3) Design bystander privacy in from the first release, extending on-device minimisation to retained copies and exported reports. (4) Treat physical compromise as the normal case: application-level encryption at rest, identity rotation, bounded retention and a rapid wipe path are baseline requirements here. (5) Make the transport in use visible – in a dual-transport system the user's expectation of locality is a security property that automatic selection silently violates (T4.6). (6) Instrument for local attack detection: a node cannot see the mesh, but an anomalous local rate of critical-priority messages, connection requests or new identities is observable and actionable. Automated on-device anomaly detection of this kind is already demonstrated for infrastructure networks, where AI agents flag incident patterns and misconfigurations without a standing analyst [39]; adapting that approach to the connectivity and power budget of a mesh node, detection that runs locally with nothing to escalate to, is a direct extension of this implication.

CONCLUSIONS AND PROSPECTS FOR FURTHER RESEARCH

Limitations and Future Work. The catalogue is analytical and the scores are expert judgements against published anchors; Section 5.4 quantifies, but does not remove, the resulting sensitivity, and a second independent rater would strengthen the result. The simulation is abstract: it omits PHY-layer collisions, connection scheduling, mobility and real energy accounting, and it validates the two coupling effects on a model of the mechanism rather than on the implementation. The catalogue rests on one case study, so components absent from it are unrepresented. Advanced and quantum-capable adversaries (T-A5) and OS or supply-chain compromise (T-A1) are excluded by assumption; the latter is worth naming, since sideloading during a shutdown is a realistic vector and reproducible builds are the corresponding control. The mapping reflects the system as documented, not as audited.

Three directions follow. T5.5 and T5.6 should be reproduced on a physical testbed of five to ten devices, where connection-interval and collision effects are expected to lower the starvation threshold below the idealised one-packet-per-slot capacity. A per-identity priority budget with relay-side enforcement should be implemented and its latency and energy cost measured on devices, together with an identity-cost mechanism, since Section 5.2 shows the budget alone is insufficient. The internet-backed transports excluded by T-A6 deserve a companion model, since a dual-transport deployment will use them whenever connectivity returns. Encoding the budget and its enforcement as versioned, testable policy, the approach already demonstrated for automated network-security policy in infrastructure settings [40], would keep the relay logic auditable as it moves from simulation to an implemented defence.

Offline BLE mesh messaging is being adopted for situations in which its failure is least recoverable, and its security analyses have consistently arrived after deployment rather than before it. This paper contributes a systematic threat model for the class: eight categories, twenty-nine threats, six adversary profiles, and a mitigation mapping showing nineteen threats partially addressed, seven unaddressed and three effectively addressed. Simulation confirms the two coupling effects the model predicts: forged critical-priority traffic above per-node capacity removes up to a third of genuine critical delivery, a per-identity budget restores it but is defeated by Sybil identities, and withdrawing two fifths of the mesh into the lowest energy mode removes 43% of delivery while the documented relay floor recovers only three to five points. The two categories added to STRIDE exist because the disaster context falsifies two assumptions ordinary threat modeling makes silently – that the parties affected by a system are its users, and that a device stays in its owner's hands. The two coupling effects are the most transferable result, because both arise from features introduced to improve outcomes and both are invisible to a purely functional evaluation. Further work will reproduce them on physical devices, implement a per-identity budget coupled to an identity cost, and extend the model to the internet-backed transports.

REFERENCES

1. Macintyre, A. G., Barbera, J. A., & Smith, E. R. (2006). Surviving collapsed structure entrapment after earthquakes: A "time-to-rescue" analysis. Prehospital and Disaster Medicine, 21(1), 4–17. https://doi.org/10.1017/S1049023X00003253
2. Albrecht, M. R., Blasco, J., Jensen, R. B., & Mareková, L. (2021). Mesh messaging in large-scale protests: Breaking Bridgefy. In K. G. Paterson (Ed.), Topics in Cryptology – CT-RSA 2021 (Lecture Notes in Computer Science, Vol. 12704, pp. 375–398). Springer. https://doi.org/10.1007/978-3-030-75539-3_16
3. Permissionless Technologies. (2026). bitchat for Android [Software repository, README]. GitHub. https://github.com/permissionlesstech/bitchat-android
4. Radocea, A. (2025, July 7). Identity is a Bitchat challenge (MITM flaw). Supernetworks. https://www.supernetworks.org/pages/blog/agentic-insecurity-vibes-on-bitchat
5. Barua, A., Al Alamin, M. A., Hossain, M. S., & Hossain, E. (2022). Security and privacy threats for Bluetooth Low Energy in IoT and wearable devices: A comprehensive survey. IEEE Open Journal of the Communications Society, 3, 251–281. https://doi.org/10.1109/OJCOMS.2022.3149732
6. Koulouras, G., Katsoulis, S., & Zantalis, F. (2025). Evolution of Bluetooth technology: BLE in the IoT ecosystem. Sensors, 25(4), 996. https://doi.org/10.3390/s25040996
7. Lacava, A., Zottola, V., Bonaldo, A., Cuomo, F., & Basagni, S. (2022). Securing Bluetooth Low Energy networking: An overview of security procedures and threats. Computer Networks, 211, 108953. https://doi.org/10.1016/j.comnet.2022.108953
8. Wang, L., Li, J., & Li, M. (2024). A rapid flooding approach based on adaptive delay and low-power sleep for BLE mesh networks. IEEE Access, 12, 65323–65332. https://doi.org/10.1109/ACCESS.2024.3398348
9. Shostack, A. (2014). Threat modeling: Designing for security. Wiley.
10. Deng, M., Wuyts, K., Scandariato, R., Preneel, B., & Joosen, W. (2011). A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering, 16(1), 3–32. https://doi.org/10.1007/s00766-010-0115-7
11. Naik, N., Jenkins, P., Grace, P., Naik, D., Prajapat, S., & Song, J. (2024). A comparative analysis of threat modelling methods: STRIDE, DREAD, VAST, PASTA, OCTAVE, and LINDDUN. In N. Naik, P. Jenkins, S. Prajapat, & P. Grace (Eds.), Contributions presented at the International Conference on Computing, Communication, Cybersecurity and AI (C3AI 2024) (Lecture Notes in Networks and Systems, Vol. 884, pp. 271–280). Springer. https://doi.org/10.1007/978-3-031-74443-3_16

12. Danielis, P., Beckmann, M., & Skodzik, J. (2020). An ISO-compliant test procedure for technical risk analyses of IoT systems based on STRIDE. In 2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC) (pp. 499–504). IEEE. https://doi.org/10.1109/COMPSAC48688.2020.0-203
13. Cäsar, M., Pawelke, T., Steffan, J., & Terhorst, G. (2022). A survey on Bluetooth Low Energy security and privacy. Computer Networks, 205, 108712. https://doi.org/10.1016/j.comnet.2021.108712
14. von Tschirschnitz, M., Peuckert, L., Franzen, F., & Grossklags, J. (2021). Method confusion attack on Bluetooth pairing. In 2021 IEEE Symposium on Security and Privacy (SP) (pp. 1332–1347). IEEE. https://doi.org/10.1109/SP40001.2021.00013
15. Wu, J., Nan, Y., Kumar, V., Tian, D. J., Bianchi, A., Payer, M., & Xu, D. (2020). BLESA: Spoofing attacks against reconnections in Bluetooth Low Energy. In 14th USENIX Workshop on Offensive Technologies (WOOT '20). USENIX Association. https://www.usenix.org/conference/woot20/presentation/wu
16. Wu, J., Wu, R., Xu, D., Tian, D. J., & Bianchi, A. (2022). Formal model-driven discovery of Bluetooth protocol design vulnerabilities. In 2022 IEEE Symposium on Security and Privacy (SP) (pp. 2285–2303). IEEE. https://doi.org/10.1109/SP46214.2022.9833777
17. Tucker, T., Searle, H., Butler, K., & Traynor, P. (2023). Blue's Clues: Practical discovery of non-discoverable Bluetooth devices. In 2023 IEEE Symposium on Security and Privacy (SP) (pp. 3098–3112). IEEE. https://doi.org/10.1109/SP46215.2023.10179358
18. Che, X., He, Y., Feng, X., Sun, K., Xu, K., & Li, Q. (2024). BlueSWAT: A lightweight state-aware security framework for Bluetooth Low Energy. In Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24) (pp. 2087–2101). ACM. https://doi.org/10.1145/3658644.3670397
19. Celosia, G., & Cunche, M. (2019). Fingerprinting Bluetooth-Low-Energy devices based on the Generic Attribute Profile. In Proceedings of the 2nd International ACM Workshop on Security and Privacy for the Internet-of-Things (pp. 24–31). ACM. https://doi.org/10.1145/3338507.3358617
20. Greß, H., Krüger, B., & Tischhauser, E. (2025). The newer, the more secure? Standards-compliant Bluetooth Low Energy man-in-the-middle attacks on fitness trackers. Sensors, 25(6), 1815. https://doi.org/10.3390/s25061815
21. Pirayesh, H., & Zeng, H. (2022). Jamming attacks and anti-jamming strategies in wireless networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 24(2), 767–809. https://doi.org/10.1109/COMST.2022.3159185
22. Sokolov, V., Skladannyi, P., & Astapenya, V. (2023). Bluetooth low-energy beacon resistance to jamming attack. In IEEE 13th International Conference on Electronics and Information Technologies (ELIT) (pp. 270–274). IEEE. https://doi.org/10.1109/ELIT61488.2023.10310815
23. Sokolov, V., Skladannyi, P., & Korshun, N. (2023). ZigBee network resistance to jamming attacks. In IEEE 6th International Conference on Information and Telecommunication Technologies and Radio Electronics (UkrMiCo) (pp. 161–165). IEEE. https://doi.org/10.1109/UkrMiCo61577.2023.10380360
24. Skallak, C., & Schmidt, S. (2024). Indescribably blue: Bluetooth Low Energy threat landscape. In Proceedings of the 9th International Conference on Internet of Things, Big Data and Security (IoTBDS) (pp. 339–350). SciTePress. https://doi.org/10.5220/0012737300003705
25. Karlof, C., & Wagner, D. (2003). Secure routing in wireless sensor networks: Attacks and countermeasures. Ad Hoc Networks, 1(2–3), 293–315. https://doi.org/10.1016/S1570-8705(03)00008-8
26. Douceur, J. R. (2002). The Sybil attack. In P. Druschel, F. Kaashoek, & A. Rowstron (Eds.), Peer-to-Peer Systems (Lecture Notes in Computer Science, Vol. 2429, pp. 251–260). Springer. https://doi.org/10.1007/3-540-45748-8_24
27. Álvarez, F., Almon, L., Hahn, A.-S., & Hollick, M. (2019). Toxic friends in your network: Breaking the Bluetooth mesh friendship concept. In Proceedings of the 5th ACM Workshop on Security Standardisation Research (SSR '19) (pp. 1–12). ACM. https://doi.org/10.1145/3338500.3360334
28. Pliekhov, O., & Babii, K. (2025). Wi-Fi Direct in Android: Creating seamless device-to-device communication. Sustainable Engineering and Innovation, 7(2), 477–492. https://doi.org/10.37868/sei.v7i2.id539
29. Pliekhov, O. (2026). AI-driven message prioritization for offline BLE mesh networks in disaster response scenarios [Preprint]. SSRN. https://doi.org/10.2139/ssrn.6428398
30. Wang, X., Peng, K., Yi, X., & Li, H. (2026). Mind the gap: Mapping wearer-bystander privacy tensions and context-adaptive pathways for camera glasses. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26). ACM. https://doi.org/10.1145/3772318.3791848

31. Langley, A., Hamburg, M., & Turner, S. (2016). Elliptic curves for security (RFC 7748). IETF. https://doi.org/10.17487/RFC7748
32. Nir, Y., & Langley, A. (2018). ChaCha20 and Poly1305 for IETF protocols (RFC 8439). IETF. https://doi.org/10.17487/RFC8439
33. Bluetooth SIG. (2021). Bluetooth SIG statement regarding the "Impersonation Attack in Bluetooth Mesh Provisioning" vulnerability (CVE-2020-26560). https://www.bluetooth.com/learn-about-bluetooth/key-attributes/bluetooth-security/reporting-security/
34. Pliekhov, O. (2026). ResQMesh AI Platform [Software repository]. GitHub. https://github.com/AleksPlekhov/ai-mesh-emergency-communication-platform
35. Oliinyk, Y. S., Platonenko, A. V., Cherevyk, V. M., Vorokhob, M. V., & Shevchuk, Y. (2025). Methods of information security in IoT technologies. Cybersecurity: Education, Science, Technique, 3(27), 100–108. https://doi.org/10.28925/2663-4023.2025.27.705
36. Locatelli, P., Perri, M., Jimenez Gutierrez, D., Lacava, A., & Cuomo, F. (2023). Device discovery and tracing in the Bluetooth Low Energy domain. Computer Communications, 202, 42–56. https://doi.org/10.1016/j.comcom.2023.02.008
37. Fischlin, M., & Sanina, O. (2024). Fake it till you make it: Enhancing security of Bluetooth Secure Connections via deferrable authentication. In Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24) (pp. 4762–4776). ACM. https://doi.org/10.1145/3658644.3670360
38. Grynkevych, G., Vasylenko, V., & Rudin, D. (2025). Data protection in Wi-Fi networks of financial institutions: WEP to WPA3 evolution and economic impact. International Scientific Journal "Internauka", (8). https://doi.org/10.25313/2520-2057-2025-8-12018
39. Grynkevych, G., Vasylenko, V., & Rudin, D. (2025). AI agents for automated network incident detection and misconfiguration identification. International Scientific Journal "Internauka", (12). https://doi.org/10.25313/2520-2057-2025-12-12019
40. Grynkevych, G., Vasylenko, V., & Rudin, D. (2026). Adaptive network security automation: DevSecOps, ML detection, and policy-as-code. International Scientific Journal "Internauka", (1). https://doi.org/10.25313/2520-2057-2026-1-12020

КІБЕРБЕЗПЕКА: освіта, наука, техніка
CYBERSECURITY: EDUCATION, SCIENCE, TECHNIQUE
№ 2 (34), 2026
ISSN 2663 – 4023

Плєхов Олександр
дослідник, старший член IEEE (IEEE Senior Member)
EPAM Systems, Inc., Ньютаун, Пенсільванія, США
ORCID:0009-0000-1343-4308
aleks.plekhov@gmail.com

Гринкевич Ганна Олександрівна
доктор технічних наук, доцент, професор кафедри інформаційної та кібернетичної безпеки імені професора Володимира Бурячка
Київський столичний університет імені Бориса Грінченка, Київ, Україна
ORCID:0000-0003-1922-5165
ggrynkevych@ukr.net

Василенко Володимир
кандидат технічних наук, доцент, доцент кафедри комп'ютерних наук
Державний університет інформаційно-комунікаційних технологій, Київ, Україна
ORCID:0000-0001-8465-6178
oknelisavvova172@gmail.com

РОЗШИРЕНА МОДЕЛЬ ЗАГРОЗ STRIDE ДЛЯ ОФЛАЙН-ЗАСТОСУНКІВ АВАРІЙНОГО ЗВ'ЯЗКУ НА ОСНОВІ BLE MESH-МЕРЕЖ

Анотація. Офлайн-обмін повідомленнями через Bluetooth Low Energy (BLE) дедалі частіше застосовується як комунікаційний рівень останньої надії, коли інфраструктура виходить з ладу. Безпеку таких систем зазвичай аналізують на рівні стеку BLE і майже ніколи – на рівні розгорнутого застосунку в умовах катастрофи. У статті представлено систематичну модель загроз для цього класу систем. Методологію STRIDE розширено двома категоріями, яких вимагає контекст катастрофи і які стандартний STRIDE не виражає, – шкода стороннім особам та фізична компрометація; для впорядкування загроз застосовано спрощену трифакторну схему оцінювання (Damage x Reproducibility x Affected parties) з опублікованими анкерами. Методологію інстанційовано на одному кейсі - ResQMesh AI, відкритій Android-платформі аварійного зв’язку на BLE mesh, – що визначає межі узагальнення зіставлення з механізмами захисту. Каталог містить 29 загроз і шість профілів порушника; за результатами зіставлення з механізмами захисту 3 загрози закриті ефективно, 19 – частково, 7 – не закриті. Два передбачені моделлю ефекти зв’язності перевірено дискретною симуляцією mesh-мережі з 60 вузлів (TTL 7, 20 прогонів): підроблений трафік критичного пріоритету з інтенсивністю, утричі більшою за пропускну здатність вузла, знижує частку своєчасної доставки справжніх критичних повідомлень з 0,87 до 0,59; бюджет ретрансляції на ідентичність відновлює її до 0,85, а ротація Sybil-ідентичностей повністю нівелює бюджет; переведення 40% вузлів у найнижчий енергорежим знижує доставку з 0,88 до 0,50, а нижня межа ретрансляції критичного трафіку повертає 3-5 в.п. Три висновки мають загальне значення: відмова від інфраструктури переносить усе початкове встановлення довіри на обмін trust-on-first-use через радіоканал, що робить підміну особи найвищим за рангом класом загроз; пріоритезація повідомлень на основі ШІ утворює примітив підсилення, оскільки підроблений високопріоритетний трафік витісняє справжній аварійний; енергоощадне ретранслювання дає порушникові дешевий важіль для сегментації мережі.

Ключові слова: моделювання загроз; STRIDE; Bluetooth Low Energy; BLE mesh; аварійний зв’язок; офлайн-обмін повідомленнями; приватність сторонніх осіб; відмова в обслуговуванні; ResQMesh AI.

СПИСОК ВИКОРИСТАНИХ ДЖЕРЕЛ

1. Macintyre, A. G., Barbera, J. A., & Smith, E. R. (2006). Surviving collapsed structure entrapment after earthquakes: A “time-to-rescue” analysis. Prehospital and Disaster Medicine, 21(1), 4–17. https://doi.org/10.1017/S1049023X00003253
2. Albrecht, M. R., Blasco, J., Jensen, R. B., & Mareková, L. (2021). Mesh messaging in large-scale protests: Breaking Bridgefy. In K. G. Paterson (Ed.), Topics in Cryptology – CT-RSA 2021 (Lecture Notes in Computer Science, Vol. 12704, pp. 375–398). Springer. https://doi.org/10.1007/978-3-030-75539-3_16

3. Permissionless Technologies. (2026). bitchat for Android [Software repository, README]. GitHub. https://github.com/permissionlesstech/bitchat-android
4. Radocea, A. (2025, July 7). Identity is a Bitchat challenge (MITM flaw). Supernetworks. https://www.supernetworks.org/pages/blog/agentic-insecurity-vibes-on-bitchat
5. Barua, A., Al Alamin, M. A., Hossain, M. S., & Hossain, E. (2022). Security and privacy threats for Bluetooth Low Energy in IoT and wearable devices: A comprehensive survey. IEEE Open Journal of the Communications Society, 3, 251–281. https://doi.org/10.1109/OJCOMS.2022.3149732
6. Koulouras, G., Katsoulis, S., & Zantalis, F. (2025). Evolution of Bluetooth technology: BLE in the IoT ecosystem. Sensors, 25(4), 996. https://doi.org/10.3390/s25040996
7. Lacava, A., Zottola, V., Bonaldo, A., Cuomo, F., & Basagni, S. (2022). Securing Bluetooth Low Energy networking: An overview of security procedures and threats. Computer Networks, 211, 108953. https://doi.org/10.1016/j.comnet.2022.108953
8. Wang, L., Li, J., & Li, M. (2024). A rapid flooding approach based on adaptive delay and low-power sleep for BLE mesh networks. IEEE Access, 12, 65323–65332. https://doi.org/10.1109/ACCESS.2024.3398348
9. Shostack, A. (2014). Threat modeling: Designing for security. Wiley.
10. Deng, M., Wuyts, K., Scandariato, R., Preneel, B., & Joosen, W. (2011). A privacy threat analysis framework: Supporting the elicitation and fulfillment of privacy requirements. Requirements Engineering, 16(1), 3–32. https://doi.org/10.1007/s00766-010-0115-7
11. Naik, N., Jenkins, P., Grace, P., Naik, D., Prajapat, S., & Song, J. (2024). A comparative analysis of threat modelling methods: STRIDE, DREAD, VAST, PASTA, OCTAVE, and LINDDUN. In N. Naik, P. Jenkins, S. Prajapat, & P. Grace (Eds.), Contributions presented at the International Conference on Computing, Communication, Cybersecurity and AI (C3AI 2024) (Lecture Notes in Networks and Systems, Vol. 884, pp. 271–280). Springer. https://doi.org/10.1007/978-3-031-74443-3_16
12. Danielis, P., Beckmann, M., & Skodzik, J. (2020). An ISO-compliant test procedure for technical risk analyses of IoT systems based on STRIDE. In 2020 IEEE 44th Annual Computers, Software, and Applications Conference (COMPSAC) (pp. 499–504). IEEE. https://doi.org/10.1109/COMPSAC48688.2020.0-203
13. Cäsar, M., Pawelke, T., Steffan, J., & Terhorst, G. (2022). A survey on Bluetooth Low Energy security and privacy. Computer Networks, 205, 108712. https://doi.org/10.1016/j.comnet.2021.108712
14. von Tschirschnitz, M., Peuckert, L., Franzen, F., & Grossklags, J. (2021). Method confusion attack on Bluetooth pairing. In 2021 IEEE Symposium on Security and Privacy (SP) (pp. 1332–1347). IEEE. https://doi.org/10.1109/SP40001.2021.00013
15. Wu, J., Nan, Y., Kumar, V., Tian, D. J., Bianchi, A., Payer, M., & Xu, D. (2020). BLESA: Spoofing attacks against reconnections in Bluetooth Low Energy. In 14th USENIX Workshop on Offensive Technologies (WOOT '20). USENIX Association. https://www.usenix.org/conference/woot20/presentation/wu
16. Wu, J., Wu, R., Xu, D., Tian, D. J., & Bianchi, A. (2022). Formal model-driven discovery of Bluetooth protocol design vulnerabilities. In 2022 IEEE Symposium on Security and Privacy (SP) (pp. 2285–2303). IEEE. https://doi.org/10.1109/SP46214.2022.9833777
17. Tucker, T., Searle, H., Butler, K., & Traynor, P. (2023). Blue's Clues: Practical discovery of non-discoverable Bluetooth devices. In 2023 IEEE Symposium on Security and Privacy (SP) (pp. 3098–3112). IEEE. https://doi.org/10.1109/SP46215.2023.10179358
18. Che, X., He, Y., Feng, X., Sun, K., Xu, K., & Li, Q. (2024). BlueSWAT: A lightweight state-aware security framework for Bluetooth Low Energy. In Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24) (pp. 2087–2101). ACM. https://doi.org/10.1145/3658644.3670397
19. Celosia, G., & Cunche, M. (2019). Fingerprinting Bluetooth-Low-Energy devices based on the Generic Attribute Profile. In Proceedings of the 2nd International ACM Workshop on Security and Privacy for the Internet-of-Things (pp. 24–31). ACM. https://doi.org/10.1145/3338507.3358617
20. Greß, H., Krüger, B., & Tischhauser, E. (2025). The newer, the more secure? Standards-compliant Bluetooth Low Energy man-in-the-middle attacks on fitness trackers. Sensors, 25(6), 1815. https://doi.org/10.3390/s25061815
21. Pirayesh, H., & Zeng, H. (2022). Jamming attacks and anti-jamming strategies in wireless networks: A comprehensive survey. IEEE Communications Surveys & Tutorials, 24(2), 767–809. https://doi.org/10.1109/COMST.2022.3159185
22. Sokolov, V., Skladannyi, P., & Astapenya, V. (2023). Bluetooth low-energy beacon resistance to jamming attack. In IEEE 13th International Conference on Electronics and Information Technologies (ELIT) (pp. 270–274). IEEE. https://doi.org/10.1109/ELIT61488.2023.10310815

23. Sokolov, V., Skladannyi, P., & Korshun, N. (2023). ZigBee network resistance to jamming attacks. In IEEE 6th International Conference on Information and Telecommunication Technologies and Radio Electronics (UkrMiCo) (pp. 161–165). IEEE. https://doi.org/10.1109/UkrMiCo61577.2023.10380360
24. Skallak, C., & Schmidt, S. (2024). Indescribably blue: Bluetooth Low Energy threat landscape. In Proceedings of the 9th International Conference on Internet of Things, Big Data and Security (IoTBDS) (pp. 339–350). SciTePress. https://doi.org/10.5220/0012737300003705
25. Karlof, C., & Wagner, D. (2003). Secure routing in wireless sensor networks: Attacks and countermeasures. Ad Hoc Networks, 1(2–3), 293–315. https://doi.org/10.1016/S1570-8705(03)00008-8
26. Douceur, J. R. (2002). The Sybil attack. In P. Druschel, F. Kaashoek, & A. Rowstron (Eds.), Peer-to-Peer Systems (Lecture Notes in Computer Science, Vol. 2429, pp. 251–260). Springer. https://doi.org/10.1007/3-540-45748-8_24
27. Álvarez, F., Almon, L., Hahn, A.-S., & Hollick, M. (2019). Toxic friends in your network: Breaking the Bluetooth mesh friendship concept. In Proceedings of the 5th ACM Workshop on Security Standardisation Research (SSR '19) (pp. 1–12). ACM. https://doi.org/10.1145/3338500.3360334
28. Pliekhov, O., & Babii, K. (2025). Wi-Fi Direct in Android: Creating seamless device-to-device communication. Sustainable Engineering and Innovation, 7(2), 477–492. https://doi.org/10.37868/sei.v7i2.id539
29. Pliekhov, O. (2026). AI-driven message prioritization for offline BLE mesh networks in disaster response scenarios [Preprint]. SSRN. https://doi.org/10.2139/ssrn.6428398
30. Wang, X., Peng, K., Yi, X., & Li, H. (2026). Mind the gap: Mapping wearer-bystander privacy tensions and context-adaptive pathways for camera glasses. In Proceedings of the 2026 CHI Conference on Human Factors in Computing Systems (CHI '26). ACM. https://doi.org/10.1145/3772318.3791848
31. Langley, A., Hamburg, M., & Turner, S. (2016). Elliptic curves for security (RFC 7748). IETF. https://doi.org/10.17487/RFC7748
32. Nir, Y., & Langley, A. (2018). ChaCha20 and Poly1305 for IETF protocols (RFC 8439). IETF. https://doi.org/10.17487/RFC8439
33. Bluetooth SIG. (2021). Bluetooth SIG statement regarding the "Impersonation Attack in Bluetooth Mesh Provisioning" vulnerability (CVE-2020-26560). https://www.bluetooth.com/learn-about-bluetooth/key-attributes/bluetooth-security/reporting-security/
34. Pliekhov, O. (2026). ResQMesh AI Platform [Software repository]. GitHub. https://github.com/AleksPlekhov/ai-mesh-emergency-communication-platform
35. Oliinyk, Y. S., Platonenko, A. V., Cherevyk, V. M., Vorokhob, M. V., & Shevchuk, Y. (2025). Methods of information security in IoT technologies. Cybersecurity: Education, Science, Technique, 3(27), 100–108. https://doi.org/10.28925/2663-4023.2025.27.705
36. Locatelli, P., Perri, M., Jimenez Gutierrez, D., Lacava, A., & Cuomo, F. (2023). Device discovery and tracing in the Bluetooth Low Energy domain. Computer Communications, 202, 42–56. https://doi.org/10.1016/j.comcom.2023.02.008
37. Fischlin, M., & Sanina, O. (2024). Fake it till you make it: Enhancing security of Bluetooth Secure Connections via deferrable authentication. In Proceedings of the 2024 ACM SIGSAC Conference on Computer and Communications Security (CCS '24) (pp. 4762–4776). ACM. https://doi.org/10.1145/3658644.3670360
38. Grynkevych, G., Vasylenko, V., & Rudin, D. (2025). Data protection in Wi-Fi networks of financial institutions: WEP to WPA3 evolution and economic impact. International Scientific Journal "Internauka", (8). https://doi.org/10.25313/2520-2057-2025-8-12018
39. Grynkevych, G., Vasylenko, V., & Rudin, D. (2025). AI agents for automated network incident detection and misconfiguration identification. International Scientific Journal "Internauka", (12). https://doi.org/10.25313/2520-2057-2025-12-12019
40. Grynkevych, G., Vasylenko, V., & Rudin, D. (2026). Adaptive network security automation: DevSecOps, ML detection, and policy-as-code. International Scientific Journal "Internauka", (1). https://doi.org/10.25313/2520-2057-2026-1-12020

Отримано редакцією журналу / Received: 22.02.26
Прорецензовано / Revised: 03.03.26
Схвалено до друку / Accepted: 24.09.26