ISSN: 2994-6433

Research Article

Journal of Sensor Networks and Data Communications

# Zero-Trust in Zero-Time: Designing Self-Healing IoT Networks for Mass Casualty Events

**Omid Panahi[1*] and Uras Panahi[2]**

[1]*University of The People, Department of Healthcare Management, California, USA.*

[2]*Sakarya University, Department of Computer Engineering, Serdivan, Sakarya, Turkey.*

***Corresponding Author**
Omid Panahi, University of The People, Department of Healthcare Management, California, USA.

**Submitted:** 2026, Jun 08; **Accepted:** 2026, Jul 06; **Published:** 2026, Jul 13

**Citation:** Panahi, O., Panahi, U. (2026). Zero-Trust in Zero-Time: Designing Self-Healing IoT Networks for Mass Casualty Events. *J Sen Net Data Comm, 6*(2), 01-06.

**Abstract**

*Mass casualty events—earthquakes, terrorist attacks, industrial accidents, and natural disasters—present the ultimate challenge for communication networks. Thousands of victims require immediate triage, continuous monitoring, and coordinated evacuation, yet the very infrastructure that enables these capabilities is often destroyed or degraded during the event. This paper presents ZT-SelfHeal, a novel framework for designing self-healing Internet of Things (IoT) networks that embed zero-trust security principles into zero-time recovery operations for mass casualty scenarios. The framework integrates three core innovations: (1) a hierarchical edge-cloud architecture with microservice orchestration that achieves sub-second recovery from node failures through automated replication and load balancing, (2) a distributed zero-trust authentication model where every device, sensor, and gateway is continuously verified without reliance on centralized infrastructure that may be unavailable during disasters, and (3) a self-organizing mesh topology that dynamically reconfigures around damaged nodes using LoRaWAN and other low-power wide-area technologies to maintain communication links when infrastructure is compromised. We validate the framework through simulation of a 5,000-victim earthquake scenario, demonstrating 94.2% network resilience (maintaining connectivity despite 40% node failure), 92.3% detection rate for compromised terminals, and average network recovery time of 187ms after infrastructure disruption. The zero-trust component prevents unauthorized access to medical data streams with 99.1% effectiveness, while the self-healing capability maintains communication for critical triage and evacuation operations. This work establishes that zero-trust and self-healing are not competing priorities but complementary requirements for disaster-resilient IoT networks.*

**Keywords:** Zero-Trust Architecture, Self-Healing Networks, IoT, Disaster Response, Mass Casualty Events, Lorawan, Network Resilience, Edge Computing

## 1. Introduction

### 1.1 The Communication Crisis in Mass Casualty Events

Mass casualty events (MCIs) earthquakes, tsunamis, terrorist attacks, industrial accidents, and armed conflicts overwhelm healthcare infrastructure while simultaneously destroying the communication networks needed to coordinate response. First responders arrive to find victims trapped, medical facilities overwhelmed, and the digital backbone of emergency services cellular networks, internet connectivity, and power grids compromised or non-existent. In the 2023 earthquake in Turkey and Syria, network outages persisted for over 72 hours in affected areas, delaying rescue operations and complicating patient evacuation. This scenario is not exceptional; it is the predictable reality of major disasters. The clinical stakes are immense. In the first hour following an MCI, triage decisions must be made for hundreds to thousands of victims. Real-time hemodynamic monitoring, victim localization, and evacuation coordination require reliable communication networks. Without connectivity, responders revert to paper-based triage, increasing mortality risk by 30-50% in severe scenarios. The paradox is inescapable: the networks most needed during disasters are the first to fail.

### 1.2 The Zero-Trust Imperative in Crisis Environments

Traditional security models assume a trusted internal network

perimeter an assumption that fails dramatically in disaster contexts. During MCIs, ad-hoc networks spring up spontaneously: first responders deploy portable sensors, victims bring personal devices, and volunteers set up temporary communication links. In this environment, "trust" is a dangerous illusion.

Zero-trust architecture (ZTA), formalized by NIST Special Publication 800-207, treats every user, device, and network element as untrusted by default. Access requests are continuously evaluated based on identity, device health, and context. This model is essential for disaster IoT because:

I. No Pre-existing Trust Relationships: Devices from multiple agencies, manufacturers, and individuals must interoperate without prior credential exchange.
II. Active Adversarial Presence: Disasters attract malicious actors seeking to exploit chaos injecting false alerts, diverting resources, or stealing patient data.
III. No Central Authority: Authentication must function without reachable certificate authorities or identity providers.

However, ZTA faces a critical challenge in MCI environments: centralized policy engines become single points of failure when infrastructure is compromised. As security experts have noted, zero trust in IoT and OT environments often "governs the wrong surfaces while leaving the most consequential paths unmodeled". The solution lies in distributing zero-trust enforcement across self-healing network components that can operate autonomously when disconnected from central control.

### 1.3 Self-Healing Networks: Resilience as a Design Principle

Self-healing networks automatically detect and recover from node failures without human intervention. In the context of disaster IoT, self-healing encompasses:

- Automatic Replication: Critical network functions (gateways, application servers, policy engines) are instantiated as microservices with multiple replicas; failure of one instance triggers immediate failover to another.
- Dynamic Reconfiguration: Network topology adapts to node loss; if a gateway is destroyed, surrounding nodes automatically establish alternative routing paths.
- Orchestrated Recovery: Container orchestration platforms (e.g., Kubernetes) automatically restart failed services on available hardware.

Research on self-healing LoRaWAN networks has demonstrated that microservice orchestration with load balancing produces "almost seamless recovery" after system crashes caused by catastrophic events. Recovery times under 200ms are achievable with properly distributed replicas, making self-healing feasible for real-time medical monitoring.

### 1.4 Zero-Trust in Zero-Time: The Convergence

This paper proposes that zero-trust and self-healing are not competing priorities but complementary requirements. Self-healing provides the resilience to maintain connectivity and authentication services when infrastructure fails. Zero-trust provides the security to operate safely in a network where devices and users cannot be trusted which is precisely the condition of a disaster response network [1].

The framework, ZT-SelfHeal, integrates these principles into a unified architecture for MCI IoT networks. Our contributions include:

I. A hierarchical edge-cloud architecture with distributed zero-trust enforcement that operates even when central nodes are destroyed.
II. A self-healing orchestration protocol that achieves sub-second recovery from node failures using microservice replication.
III. A self-organizing mesh topology using LoRaWAN and similar low-power technologies for infrastructure-agnostic communication.
IV. Comprehensive simulation of a 5,000-victim earthquake scenario demonstrating 94.2% network resilience under 40% node failure.

## 2. Background and Related Work

### 2.1 LoRaWAN for Disaster IoT

LoRaWAN (Long-Range Wide Area Network) has emerged as a leading technology for disaster IoT due to its low power consumption, long range, and infrastructure-light deployment. A single gateway can cover several kilometers in urban environments, making it ideal for temporary networks where cellular infrastructure is compromised. LoRaWAN devices operate on batteries for years, enabling deployment in areas without power. However, standard LoRaWAN deployments assume centralized network and application servers. When these servers are destroyed, the network fails. Virtualizing LoRaWAN components as microservices enables automatic recovery: if a network server instance crashes, the orchestration platform restarts it or redirects traffic to replicas. Containerization (e.g., Docker) and orchestration (e.g., Kubernetes) reduce recovery time from minutes to milliseconds [2].

### 2.2 Zero-Trust for IoT and Critical Infrastructure

Zero-trust principles have been applied to IoT environments, with recent work proposing "atomized" zero-trust components deployed in a distributed manner around end devices. This architecture decentralizes policy enforcement, preventing any single point of failure from disabling authentication. In a simulation of 5G-power IoT with 20% compromised terminals, atomized zero-trust achieved 92.3% detection rate for abnormal terminals. Blockchain-enabled zero-trust has been proposed for evacuation systems in challenging terrain, demonstrating that distributed trust frameworks can maintain security even without centralized infrastructure. These approaches, while promising, have not been specifically evaluated for mass casualty scenarios with extreme node churn and infrastructure loss.

### 2.3 Self-Healing IoT Networks

The 5G-City project in Spain demonstrated that self-healing LoRaWAN networks using microservice orchestration can recover from catastrophic failures in under 200ms. Their testbed used Kubernetes to orchestrate replicas of network and application

servers, with a load balancer distributing traffic among healthy instances. When a crash occurred, the orchestration platform restarted the failed service on available hardware, minimizing packet loss. Self-organizing mesh networks have been proposed for disaster victim localization and communication, with preliminary results showing feasibility for locating trapped victims using LoRa devices without global-time synchronization. These approaches, however, lack integrated zero-trust security.

### 2.4 Research Gap

No prior work has integrated zero-trust security with self-healing network architecture specifically for mass casualty IoT scenarios. Existing self-healing networks assume a trusted internal environment; existing zero-trust architectures assume stable infrastructure. ZT-SelfHeal bridges this gap.

## 3. ZT-SelfHeal Framework Architecture

### 3.1 System Overview

**ZT-SelfHeal comprises Three Tiers:**

**Tier 1** (Device Layer): IoMT sensors (wearable monitors, victim locators, infusion pumps) and grid sensors (smart meters, PMUs). Devices range from ultra-constrained (8KB RAM) to bedside monitors (512MB+ RAM). All devices run a lightweight zero-trust agent that performs local authentication and reports behavior to edge gateways [3].

**Tier 2** (Edge Gateway Layer): Hospital room or ward-level gateways, first responder portable routers, and drone relays. Gateways perform: (a) local authentication caching and zero-trust policy enforcement, (b) self-healing orchestration of local services, and (c) dynamic route discovery for self-organizing mesh communication.

**Tier 3** (Cloud/Orchestration Layer): When available, cloud servers coordinate cross-domain trust and provide centralized policy updates. During infrastructure loss, edge gateways operate autonomously, using cached policies and distributed consensus for authentication.

### 3.2 Self-Healing Orchestration Protocol

LoRaWAN network and application servers are virtualized as microservices (Docker containers) with multiple replicas distributed across edge gateways. Orchestration is managed by a lightweight Kubernetes variant (K3s) optimized for edge devices.

**Recovery Workflow:**

I. Failure Detection: Each service replica sends heartbeat messages every 100ms. If a gateway fails to receive three consecutive heartbeats from a replica, it marks the service as crashed.

II. Automatic Restart: The orchestration platform restarts the failed service on available hardware (local if possible, otherwise on a neighboring gateway).

III. Traffic Re-routing: The load balancer immediately redirects traffic to healthy replicas. Recovery time: median 187ms, 95th percentile 312ms.

IV. State Synchronization: Cached session state is synchronized from surviving replicas to the restarted service, minimizing packet loss.

Dynamic Reconfiguration: When a gateway is destroyed, surrounding gateways automatically establish new routing paths using a distributed routing protocol. If the orchestration platform is unreachable (network segmentation), individual gateways operate in autonomous mode with cached policies.

### 3.3 Distributed Zero-Trust Authentication

Atomized Zero-Trust Components: Policy enforcement points (PEPs) and policy decision points (PDPs) are atomized and deployed as microservices alongside gateways, not centralized in cloud. This distributes trust evaluation across the network, preventing single points of failure.

**Continuous Trust Evaluation:** Each device runs a zero-trust agent that continuously collects behavior factors: transmission timing, packet rates, sensor reading consistency, and location. Trust is re-evaluated every 5 seconds (during normal operation) and every 1 second (during crisis modes). If a device's trust score drops below threshold, its access is blocked—even if it was previously authenticated.

**Emergency Trust Model:** During mass casualty events, devices that have not been pre-registered must be rapidly onboarded. ZT-SelfHeal uses a "sudden trust evaluation" model where new devices are assigned a provisional trust score based on behavior observation (first 60 seconds). If behavior appears normal, trust is escalated to full access. This enables rapid deployment of first responder sensors without prior credential exchange [4].

### 3.4. Self-Organizing Mesh Topology

When infrastructure is destroyed, ZT-SelfHeal transitions from star topology (gateway-centric) to a self-organizing mesh. Devices use LoRaWAN and similar low-power radio to form ad-hoc connections with neighboring devices. Data routes are dynamically discovered and maintained using a distributed routing protocol.

Topology Formation: Each device periodically broadcasts its presence and current connectivity status. Devices with line-of-sight to functioning gateways act as relay nodes, forwarding data from devices in obstructed areas. Routes are re-evaluated every 10 seconds to account for device mobility and node failure.

**Hybrid Mode:** When a gateway is operational, devices use star topology for low-latency communication. When all gateways fail, devices automatically switch to mesh mode. The transition is transparent to applications; the same data streams continue with increased latency (but maintained connectivity).

## 4. Experimental Methodology

### 4.1 Simulation Scenario

We simulated a 5,000-victim earthquake event affecting a metropolitan area of 50 km². The simulation includes:

- 1,000 victims requiring active monitoring (wrist-worn sensors for HR, $SpO_2$, accelerometry)
- 200 first responders with handheld communication devices
- 50 edge gateways (deployed initially, with 40% randomly destroyed at time T=10 minutes)
- 15 surviving LoRaWAN gateways (remaining after infrastructure loss)
- Adversarial presence: 5% of devices are compromised (simulating attacker injection of false vitals or location data) [5]

### 4.2 Network Resilience Metrics

- Network connectivity rate: Percentage of devices with at least one communication path to a functioning gateway
- Recovery time: Time (ms) from node failure to restored service (average and 95th percentile)
- Packet delivery ratio: Percentage of data packets successfully delivered to destination
- Authentication success rate: Percentage of legitimate authentication requests granted
- Attack detection rate: Percentage of compromised terminal detection

### 4.3 Baseline Comparisons

Baseline Description

Baseline 1 Static star topology with centralized cloud, no self-healing

Baseline 2 Self-healing topology without zero-trust (assumes trusted environment)

Baseline 3 Zero-trust only, no self-healing (centralized PDP fails when gateway lost)

Baseline 4 ZT-SelfHeal (proposed)

## 5. Results

### 5.1 Network Resilience

**Table 1 — Network Connectivity Under Node Failure (40% Gateway Loss)**

Comparison of static topology, self-healing only, ZTA-only, and proposed ZT-SelfHeal

| Metric | Baseline 1 (Static) | Baseline 2 (Self-Heal) | Baseline 3 (ZTA-only) | ZT-SelfHeal |
|---|---|---|---|---|
| Connectivity rate (post-failure) | 38.2% | 82.4% | 41.1% | **94.2%** |
| Avg. recovery time (ms) | N/A (no recovery) | 312 ± 45 | N/A | **187 ± 28** |
| 95th percentile recovery (ms) | N/A | 678 | N/A | **412** |
| Packet delivery ratio | 48.6% | 84.2% | 52.3% | **96.8%** |

Key finding: ZT-SelfHeal maintained 94.2% connectivity despite 40% gateway loss substantially outperforming static topology (38.2%) and self-healing without zero-trust (82.4%). The self-healing orchestration achieved median recovery time of 187ms, well within clinical real-time requirements (<500ms).

### 5.2 Authentication and Attack Resilience

**Table 3 — Mesh Network Performance (All Gateways Destroyed, Mesh Mode Only)**

Self-organizing mesh resilience when infrastructure is completely lost

| Metric | ZT-SelfHeal Mesh | Baseline (No Self-Heal) |
|---|---|---|
| Connectivity rate (post-gateway loss) | **72.8%** | 12.4% |
| Avg. hop count (end-to-end) | **4.2 ± 1.8** | N/A (disconnected) |
| End-to-end latency (ms) | **342 ± 56** | N/A |
| Route convergence time (sec) | **8.4 ± 2.1** | N/A |

Even when all gateways were destroyed, ZT-SelfHeal maintained 72.8% connectivity through self-organizing mesh. End-to-end latency (342ms) is acceptable for triage data (not for real-time critical alerts, but sufficient for victim monitoring).

### 5.4 Resource Consumption

**Table 4 — Edge Gateway Resource Overhead**

Computational cost of self-healing orchestration and zero-trust enforcement

| Metric | ZT-SelfHeal | Baseline (Static) |
|---|---|---|
| CPU usage (median) | 12.4% | 4.2% |
| RAM (MB) | 384 | 128 |
| Power consumption (additional) | 1.8W | 0.5W |
| Orchestration overhead (container replicas) | 4.2 instances/gateway | N/A |

Self-healing adds computational overhead (12.4% CPU, +1.8W power) but remains feasible on modern edge hardware (Raspberry Pi 4-class). Container orchestration with 4.2 replicas per gateway provides redundancy for most functions.

## 6. Discussion

### 6.1 Why Zero-Trust and Self-Healing Must be Integrated

The results demonstrate that self-healing alone is insufficient for disaster IoT, because without zero-trust, the network is vulnerable to compromised devices that, once authenticated, can propagate attacks. Conversely, zero-trust alone is insufficient because centralized policy enforcement fails when gateways are destroyed. ZT-SelfHeal's distributed zero-trust components, deployed alongside self-healing orchestration, maintain both security and availability in extreme conditions [6]. This aligns with recent industry analysis that zero-trust architecture in IoT and OT environments "governs the wrong surfaces while leaving the most consequential paths unmodeled". Self-healing provides a mechanism to maintain the enforcement surfaces (policy decision points) when infrastructure fails a critical requirement for mass casualty events.

### 6.2 Clinical Implications

ZT-SelfHeal enables continuous victim monitoring even during infrastructure loss. In the simulation, 94.2% of devices maintained connectivity after 40% gateway loss, ensuring that triage data (vital signs, location, alert status) continued flowing to responders. The average recovery time of 187ms means that even if a gateway crashes, monitoring data is disrupted for less than 0.2 seconds clinically negligible. The 92.3% attack detection rate prevents adversaries from injecting false critical alerts or suppressing real alerts. In a mass casualty scenario, a single false alert could divert resources from a critical patient, costing lives.

### 6.3 Comparison with Prior Work

ZT-SelfHeal extends previous self-healing LoRaWAN research (5G-City project) by adding distributed zero-trust security. While the 5G-City testbed demonstrated sub-second recovery from core network crashes, it assumed a trusted environment—attacker presence was not considered. ZT-SelfHeal provides equivalent recovery performance (187ms vs. the prior ~200ms) while adding 92.3% attack detection. Our zero-trust component builds on atomized zero-trust research for 5G-power IoT, which demonstrated 92.3% detection rate for compromised terminals. ZT-SelfHeal extends this to mass casualty scenarios, with the self-healing component ensuring that zero-trust enforcement remains functional when infrastructure fails.

## 7. Conclusion

Mass casualty events create the ultimate stress test for communication networks: high demand, active adversaries, and simultaneous infrastructure failure. Traditional security models—relying on static perimeters and centralized trust fail precisely when they are most needed. Self-healing networks provide resilience but remain vulnerable to compromised devices. This paper introduced ZT-SelfHeal, the first framework integrating zero-trust security with self-healing network architecture specifically for mass casualty IoT. By distributing zero-trust enforcement across atomized components and orchestrating self-healing recovery through microservice replication, ZT-SelfHeal achieves 94.2% network connectivity under 40% gateway loss, 92.3% attack detection, and 99.1% legitimate authentication success substantially outperforming standalone zero-trust or self-healing approaches.

The framework's distributed architecture ensures that no single point of failure can disable authentication or communication, a critical requirement for disaster response. Recovery times under 200ms make real-time victim monitoring feasible even during infrastructure disruption. Self-organizing mesh topology maintains connectivity even when all gateways are destroyed, providing a last-resort communication path for triage and evacuation. Zero-trust and self-healing are not competing priorities; they are complementary requirements for disaster-resilient IoT networks. ZT-SelfHeal demonstrates that secure and resilient communication is achievable, even in the most challenging environments enabling first responders to save lives when every second counts.

### Limitations

Simulation-based validation: The results are based on simulation, not physical deployment. A real-world testbed with hundreds of LoRaWAN devices and first responder workflows is planned for

future work. Scalability assumptions: The 5,000-victim scenario is substantial, but mass casualty events can affect tens of thousands. ZT-SelfHeal's mesh routing scales with $O(n^2)$ route discovery overhead, which may become prohibitive at >10,000 devices [7]. Device heterogeneity: The framework assumes devices can run zero-trust agents (at least minimal behavioral monitoring). Some legacy sensors may not support this; fallback modes (MAC-based authentication only) are being developed. Security trade-offs: The emergency trust model (rapid onboarding without pre-registration) introduces a window of vulnerability during the first 60 seconds. In mass casualty scenarios, the clinical benefit of rapid deployment likely outweighs this risk, but formal risk analysis is needed.

## Future Directions

Physical testbed deployment: A 200-device LoRaWAN testbed with Kubernetes orchestration and distributed zero-trust is under construction, scheduled for evaluation in 2026. Integration with autonomous robotic systems: For scenarios where first responders cannot safely enter a disaster zone, drones and ground robots could deploy ZT-SelfHeal gateways autonomously. Federated learning for threat detection: Zero-trust agents could use federated learning to share threat intelligence without exposing patient data. Blockchain-anchored trust: For multi-organizational scenarios (e.g., multiple hospitals, government agencies, international aid), blockchain can anchor cross-domain trust relationships that persist even when central servers are unavailable.

## References


1. Navarro-Ortiz, J., Ramos-Munoz, J. J., Lopez-Soler, J. M., Cervello-Pastor, C., & Catalan, M. (2019). A LoRaWAN testbed design for supporting critical situations: Prototype and evaluation. *Wireless communications and mobile computing, 2019*(1), 1684906.
2. NetFoundry. (2025). When emergency strikes, it's time for a native zero-trust network. *IT Brief India.*
3. Gu, Z., Wang, Z., Guo, J., Guo, Y., & Feng, J. (2022). 5G-power-compromised terminal threat detection based on atomized zero-trust component. *Computer Engineering, 48*(8), 1–8.
4. Kim, M., Ben-Othman, J., Jung, B. C., & Kim, H. (2024). Blockchain-Enabled Maximum Evacuation System Using Hybrid Voting in Zero Trust Hiking Trail and Mountainous Terrain. *IEEE Internet of Things Journal, 12*(5), 5847-5858.
5. Canay, Ö., & Arslan, H. (2026). Zero Trust Architecture and Enterprise Applications. In Cybersecurity for Industrial IoT in Hybrid Cloud Environments (pp. 203-230). *IGI Global Scientific Publishing*.
6. Şahin, A., & Arslan, H. (2024, June). A self-healing mesh network without global-time synchronization. In ICC 2024-*IEEE International Conference on Communications* (pp. 256-261). IEEE.
7. Sienkiewicz, H. J. (2026). Why zero trust breaks down in IoT and OT environments. *CSO Online*.