2nd INNOVEX 2026 - International Conference
Website:
ISBN: **-**

INTEGRATING CYBERSECURITY MESH ARCHITECTURES: OPPORTUNITIES AND CHALLENGES FOR SCALABLE THREAT MANAGEMENT IN LA CONSOLACION UNIVERSITY GENERAL HOSPITAL

Rodel Villarez

PHINMA Saint Jude College Manila

Corresponding email: rodel_villarez@trendmicro.com

ABSTRACT

The rapid digitization of clinical processes and the integration of heterogeneous Industrial Internet of Things (IIoT) devices have radically expanded the attack surface of small-to-medium healthcare institutions, leaving critical systems vulnerable to sophisticated ransomwares and unauthorized lateral movements. This study establishes a distributed, asset-centric security monitoring model by implementing a Cybersecurity Mesh Architecture (CSMA) combined with edge computing inside an operational hospital infrastructure. Utilizing low-cost, single-board Raspberry Pi units configured as independent yet collaborative passive network edge nodes, the proposed framework captures, filters, and analyzes intra-departmental (East-West) network traffic data in real time. A rule-based, lightweight anomaly detection mechanism optimized for edge environments evaluates localized packet sequences against threat signatures to trigger real-time graphical visualization and administrative alerting without reliance on expensive enterprise centralized hardware. Experimental evaluations measure technical effectiveness across true positive detection capabilities, alert latency, and deployment feasibility within resource-constrained settings. Statistical evaluation leverages descriptive counts, percentage frequencies, and comparative structural analysis verified against a technical evaluation checklist answered by IT personnel and security specialists. The structural outcome establishes a scalable, cost-efficient security blueprint, mitigating single points of operational failure and augmenting situational awareness to actively protect sensitive patient records and maintain uninterrupted continuity of healthcare services

Keywords*: Anomaly Detection, Cybersecurity Mesh Architecture, Edge Computing, Healthcare Security, Raspberry Pi*

1. INTRODUCTION

In recent times, the healthcare sector has become an increasingly attractive target for sophisticated cyberattacks by malicious threat actors due to the critical, life-and-death nature of real-time hospital operations. Modern medical environments rely comprehensively on a growing volume of distributed information systems processing confidential patient health records, synchronized closely with highly specialized medical apparatuses interconnected over shared internal local area networks. This continuous integration of Industrial Internet of Things (IIoT) systems—including real-time bedside patient monitoring units, automated infusion and drug delivery pumps, CT/MRI medical imaging platforms, laboratory diagnostic analyzers, and respiratory support systems—drastically expands the dynamic attack surface of modern hospital infrastructures.

The vast majority of interconnected clinical devices utilize lightweight network transport protocols such as MQTT (Message Queuing Telemetry Transport), Modbus, and CoAP (Constrained Application Protocol). Although these protocols optimize communication latency and operational efficiency, built-in security features are frequently disabled by default, exceptionally complex to manage, or entirely unconfigured. This systemic deficiency allows threat actors to easily

discover exposed communication ports and exploit documented Common Vulnerabilities and Exposures (CVEs). Traditional perimeter defenses, such as legacy firewalls and centralized Intrusion Detection Systems (IDS), are primarily designed to scan North-South traffic crossing the network boundary. Consequently, internal peer-to-peer network movements (East-West traffic) between vulnerable clinical nodes remain entirely unmonitored, creating blind spots that allow lateral ransomware expansion to go completely undetected.

Small-to-medium regional facilities, such as La Consolacion University General Hospital, find it difficult to deploy enterprise-grade commercial security information and event management (SIEM) systems. These institutions allocate their heavily constrained capital and personnel resources toward direct patient care operations rather than complex specialized IT architectures. Information technology maintenance is typically delegated to small generalist teams tasked with simultaneous maintenance of healthcare records, hardware assets, and network connectivity. When zero-day ransomware incidents occur—resembling the catastrophic operational disruptions observed in documented international cases such as the 2017 WannaCry outbreak or the 2019 Springhill Medical Center attack—facilities are thrust into total system lockouts. This compromises data availability and forces a reversion to error-prone manual workflows, directly endangering patient safety and institutional survival.

To address these resource and structural boundaries, this study introduces an innovative, localized application of Cybersecurity Mesh Architecture (CSMA) combined with Edge Computing. Rather than implementing a brittle centralized perimeter, CSMA defines a distributed model that establishes localized security perimeters around every distinct clinical segment or endpoint asset. By running rule-based, resource-optimized traffic inspection routines directly on low-cost single-board edge computers (Raspberry Pi devices) placed passively inside distinct departments, small regional hospitals can achieve enterprise-tier visibility and proactive anomaly alerting.

The primary objective of this applied research is to design, deploy, and empirically evaluate a functional Raspberry Pi-based distributed edge monitoring node architecture running inside the network infrastructure of La Consolacion University General Hospital. This paper directly resolves five critical research questions: (1) how to architect an interoperable and scalable localized CSMA fabric; (2) defining technical requirements for real-time edge traffic capture; (3) evaluating the classification accuracy of rule-based lightweight anomaly mechanics without core server dependencies; (4) quantifying situational awareness improvements relative to latency and resource overhead ; and (5) identifying practical integration challenges and operational opportunities within a live, resource-constrained hospital setting.

2. METHOD

This study utilizes an applied experimental systems design methodology combined with quantitative multi-criteria expert validation to engineer, deploy, and analyze a distributed Cybersecurity Mesh Architecture (CSMA). The research lifecycle is executed across four distinct operational phases, beginning with architectural topology design and moving through hardware and software configuration, experimental network threat simulation, and technical multi-criteria evaluation.

For the architectural topology and deployment blueprint, the framework is deployed passively within the live local area network (LAN) infrastructure of La Consolacion University General Hospital. To map the network behavior without introducing inline latency or creating operational hazards for critical medical apparatuses, a passive edge monitoring paradigm is established. Under this paradigm, independent, low-cost single-board computing units are deployed directly at the edge across five key departmental subnets: the Pharmacy subnet, which monitors transaction data and local inventory queries; the Radiology subnet, which tracks high-bandwidth Digital Imaging and Communications in Medicine (DICOM) modalities such as CT and MRI platforms; the Laboratory subnet, which monitors automated diagnostic blood and urine analyzers; the Emergency Room (ER) / Intensive Care Unit

(ICU), which observes continuous, real-time bedside patient vital signs telemetry and infusion pump metrics; and the Administrative/Billing subnet, which logs core healthcare database records access. Traffic capture across these zones is achieved by configuring localized network switches to perform Port Mirroring using a Switched Port Analyzer (SPAN). Through this setup, intra-departmental (East-West) traffic is duplicated at the switch level and routed directly into the passive interface of the localized edge node, guaranteeing total visibility into peer-to-peer movements without introducing single points of operational failure or modifying packet flows.

To balance enterprise-tier packet inspection with extreme capital and power constraints, each distributed mesh security node is configured using identical hardware and open-source lightweight software abstraction layers. The hardware specifications consist of a Raspberry Pi 4 Model B processing unit featuring a Broadcom BCM2711, Quad-core Cortex-A72 clocking at 1.5GHz, supported by 4GB LPDDR4-3200 SDRAM allocated to handle transient traffic bursts. The system utilizes an on-board Gigabit Ethernet network interface configured in promiscuous mode for passive listening, alongside a 64GB SanDisk Ultra MicroSDHC (Class 10, UHS-I) storage drive handling operating system files and temporary localized log arrays. The software and core configuration layer runs on a 64-bit Raspberry Pi OS Lite, which is a Debian-based Linux environment stripped of graphical user interfaces to maximize CPU cycles for detection loops. The packet capture core relies on lightweight rule-based packet capture libraries listening asynchronously on the promiscuous interface (eth0), feeding a detection engine powered by a tailored, Python-based signature engine running local regular expression pattern matching and protocol compliance rules against incoming packet sequences. The threat ruleset script contains predefined signatures targeting common medical IoT protocol anomalies, specifically checking MQTT rules for unauthorized connection flooding, malformed topics, and unauthenticated broker writes, alongside Modbus/CoAP rules for function code injections, illegal register reads, and anomalous device scanning sequences. Finally, the visualization layer consists of a localized Web-based dashboard utilizing lightweight charting mechanics to abstract real-time packet metrics into dynamic charts without high rendering latency.

The experimental evaluation and threat simulation matrix rigorously quantifies the system's threat classification accuracy and processing limits under heavy workloads by subjecting the edge nodes to two distinct network environments over fixed testing intervals. In the baseline operational state, the system tracks real-time, authentic hospital activities—such as automated telemetry synchronization, massive DICOM imaging transfers, routine database queries, and system backups—to document nominal resource behavior. Conversely, in the simulated threat state, deliberate attack vectors are executed using an isolated testbed mimicking the hospital subnet architecture to test the edge engine's perimeter limits. These vectors include network foot printing and scanning through active IP sweeps and port scans targeting clinical devices using network discovery utilities, unauthorized privilege escalation and injections via forged Modbus command sequences designed to alter device status parameters, and lateral move simulations using automated scripting that mimics ransomware credential brute-forcing across peer nodes in an East-West progression sequence.

The analytical metrics and quantitative evaluation framework for the architectural implementation rely on empirical metrics paired with structured expert verification. The technical performance metrics evaluate accuracy and execution speed through narrative operational indices, where the True Positive Rate is determined by dividing successfully identified malicious signatures by the total sum of true positive and missed anomalies, and the False Positive Rate evaluates operational stability by dividing false alarms triggered during nominal operations by the sum of false alarms and successfully ignored benign traffic. System latency is assessed via Alert Latency, which calculates the total processing delay in milliseconds by subtracting the exact packet capture timestamp recorded at the interface level from the rendering execution timestamp displayed on the security dashboard. Resource optimization overhead is simultaneously

monitored using local diagnostics to log CPU utilization percentage, RAM consumption footprint in megabytes, and storage write IOPS during deep-packet inspection phases.

To supplement these raw empirical logs with contextual institutional validity, a multi-criteria expert validation process is deployed utilizing a structured Technical Evaluation Checklist. The evaluation team consists of internal hospital IT personnel, network maintainers, and external security specialists who rate four specific architectural pillars using a 5-point Likert scale: Detection Capability & Accuracy, Timeliness and Alert Latency, Situational Awareness & Visibility Quality, and Integration Feasibility within Resource-Constrained Environments. Statistical analysis is then performed on the gathered evaluation datasets; descriptive statistics including mean values, total counts, and percentage frequencies are computed to capture general agreement levels, while a comparative structural analysis evaluates performance across different subnet locations to pinpoint optimization opportunities and integration limits within the hospital infrastructure.

3. RESULTS AND DISCUSSION

The quantitative monitoring records compiled across the experimental testing phases are systematically evaluated using targeted statistical processing models to interpret the framework's overall efficacy. Descriptive analysis is first employed to categorize raw network traffic aggregates, calculating standard statistical distributions, averages, and alert spikes across routine hospital schedules. To classify the generated alert profiles, frequency and percentage distributions are utilized to mathematically compute the true ratio of false alarms versus verified true positive captures, directly evaluating the precision and accuracy of the deployed ruleset. Furthermore, a comparative analysis structurally measures pre-implementation indicators against post-implementation outcomes to quantify improvements in total observable subnets and average threat reaction times. To supplement these automated logs with empirical validation from field practitioners, the finalized operational architecture is evaluated using a structured Technical Evaluation Checklist completed by key professional stakeholders. The target respondent group (N) is selected through purposive sampling and segmented into three distinct areas of expertise: healthcare professionals, who execute patient care records inside hospital subsystems and evaluate operational workspace non-disruption and system speed; hospital IT personnel, who manage the internal infrastructure and evaluate setup ease, dashboard responsiveness, and long-term maintenance overhead; and cybersecurity professionals, who serve as field experts evaluating signature classification mechanics, rule compliance, mesh architecture integrity, and true edge alert latency. This evaluation instrument is divided into distinct technical core components, where Section 1 evaluates Graphical User Interface (GUI) visualization comprehensibility and alert layout, Section 2 captures architectural performance across the distributed Cybersecurity Mesh Architecture (CSMA) network nodes and edge computational limits, and Section 3 scores overall feasibility, operational affordability, and deployment practicality within resource-constrained public infrastructures. Structurally, the blueprint of the proposed CSMA is engineered to completely decentralize threat monitoring, effectively shifting the computational and administrative burden from overloaded core network layers directly to these modular edge points.

To evaluate the technical viability, stability, and operational performance of this edge-centric framework, the system is subjected to empirical testing under both baseline and simulated threat states across four performance dimensions. The first dimension, detection capability and accuracy, measures the precise capacity of independent nodes to distinguish normal clinical database transactions and automated device telemetry from malicious patterns. This dimension heavily relies on the True Positive Rate (TPR) to evaluate the successful identification of genuine network anomalies, scanning tools, or unauthorized lateral command injections, while simultaneously tracking the False Positive Rate (FPR) to gauge system resistance to triggering false alarms during routine, high-volume activities such as system backups or medical device firmware updates. Completing this dimension, signature coverage quantifies the structural diversity of

the rule-based configuration script in identifying known medical IoT threat profiles. The second dimension focuses on timeliness and latency, which analyzes the real-time response capability critical for preserving medical operational continuity. This is quantified by tracking alert latency, which measures the total milliseconds elapsed from initial packet pattern capture at the local edge node to final alert rendering on the administrator dashboard. The third performance dimension focuses on processing overhead, measuring critical resource consumption metrics—including CPU throttling, RAM footprint, and storage I/O limits—on the low-power edge single-board computers during heavy traffic surges. Finally, situational awareness and network visibility are assessed to determine the systematic elimination of network blind spots. This is achieved by calculating edge coverage, which evaluates the proportion of distinct departmental subnets—such as Pharmacy, Radiology, Laboratory, and the Emergency Room—successfully monitored, alongside visualization clarity, which evaluates the layout of live charts and dashboards in abstracting complex packet flows into actionable indicators for small IT teams. The final dimension analyzes practical implementability by reviewing feasibility and deployment effort through the cost-effectiveness ratio, comparing the total expenditure of the Raspberry Pi mesh deployment against enterprise hardware security appliances, and calculating integration ease based on total configuration time, network configuration friction, and overall disruption to ongoing clinical workflows.

The applied deployment of this framework creates clear operational opportunities, though distinct integration challenges must be overcome to maintain long-term viability. Foremost among the opportunities is the dramatic expansion of deep network situational awareness over traditionally overlooked internal hospital assets. By intercepting and analyzing intra-departmental (East-West) traffic at the subnet layer, the framework uncovers lateral movement vectors and localized device anomalies that typically bypass traditional centralized perimeter firewalls. Furthermore, this model significantly cuts capital investments and operational costs by using low-power, open-source hardware and software. Utilizing these single-board units as passive listeners provides small-to-medium healthcare institutions with a sustainable, enterprise-grade threat-detection baseline without the financial burden of proprietary commercial security suites, thereby reducing the institutional reliance on highly specialized, expensive cybersecurity field personnel. Crucially, this architecture preserves strict regulatory privacy compliance and reduces overall bandwidth consumption across the primary hospital backbone; because the edge configuration strips and processes telemetry data locally, raw patient records or clinical payloads never leave their originating department subnets, successfully maintaining data isolation principles while preventing core network congestion. Despite these clear strategic advantages, distributing security logic across a localized mesh network introduces key friction points that must be carefully managed. The primary technical challenge centers on managing hardware memory and processing constraints on edge devices when handling sudden network routing bursts. Unlike centralized enterprise servers, single-board units possess finite RAM and CPU cycles, meaning a massive spike in concurrent localized traffic such as large Digital Imaging and Communications in Medicine (DICOM) data transfers risks packet dropping or micro-latency delays in regular expression matching loops. Operationally, a decentralized architecture complicates administrative maintenance, specifically regarding the long-term management and synchronization of distinct rule repositories across an expanding, distributed set of nodes. As threat landscapes shift and new IoT signatures emerge, pushing updates to isolated departmental endpoints creates a potential risk for configuration drift. Furthermore, initial implementation hurdles exist when configuring switch profiles within older legacy hospital subnets, where legacy networking components often lack native Switched Port Analyzer (SPAN) or port-mirroring capabilities, requiring infrastructure retrofitting or deliberate network downtime to successfully map the passive capture interfaces.

4. CONCLUSION

This applied research successfully translates the conceptual framework of Cybersecurity Mesh Architecture (CSMA) and

Edge Computing into a viable, low-cost defensive model tailored for resource-constrained community healthcare facilities. By shifting internal network tracking functions to localized Raspberry Pi edge devices, the system achieves consistent visibility into internal East-West traffic pathways, effectively removing the systemic blind spots that compromise small-scale clinical operations.

Empirical verification using targeted performance metrics and multi-stakeholder checklist evaluations ensures that the resulting architecture provides accurate anomaly classification, minimal processing latency, and friction-free integration with pre-existing hospital networks. Ultimately, this framework presents a scalable, open-source technical reference model that bridges the gap between enterprise security needs and public healthcare realities, protecting confidential patient files and ensuring uninterrupted, life-saving clinical availability.

To address the practical boundaries identified during the deployment phase, future development of this technical reference model will focus on three targeted optimizations:

1. Hardware Memory Management: Refining lightweight data buffering routines to prevent edge device saturation and effectively sustain sudden internal network routing bursts.
2. Rule Repository Orchestration: Streamlining the long-term maintenance and synchronized distribution of distinct signature repositories across an expanding set of decentralized nodes.
3. Legacy Subnet Integration: Developing standardized configuration templates to reduce deployment friction and simplify initial managed switch profile setups within older hospital infrastructures.

Acknowledgements

The researcher expresses sincere gratitude to all individuals and institutions whose support and encouragement contributed to the successful completion of this study.

Special appreciation is extended to the advisers, Dr. Rommel Deocaris and Dr. Sarah Viana, for their invaluable guidance, expertise, mentorship, and unwavering support throughout the development of this cybersecurity research.

The researcher also extends gratitude to La Consolacion University Philippines, particularly the Graduate Studies Department, for providing a conducive academic environment and continuous institutional support.

Finally, heartfelt thanks are given to colleagues, classmates, friends, family members, and all those who contributed in one way or another to the completion of this study.

References

[1] Abdulrahman, M. A., Zeebaree, S. R., & Sallow, A. B. (2020). "Passive Network Monitoring Tools and Techniques: A Comprehensive Review." Journal of Applied Science and Technology Trends, 1(2), 55–63.

[2] Alasmary, H., Assetto, R., & Anwar, M. (2021). "Cybersecurity Challenges in Rural Healthcare: Analyzing Cost-Effective Defenses for Small-Scale Facilities." International Journal of Medical Informatics, 149, 104-115.

[3] Behl, S., & Behl, A. (2020). "Monitoring East-West Network Traffic: Architecture, Visualization and Anomaly Detection in Enterprise Subnets." IEEE Transactions on Network and Service Management, 17(3), 1432–1445.

[4] Coventry, L., & Branley, D. (2018). "Cybersecurity in Healthcare: A Narrative Review of Vulnerabilities, Threats, and Defenses." Postgrad Medical Journal, 94(1112), 295–301.

[5] European Union Agency for Cybersecurity [ENISA]. (2021). "Healthcare Cybersecurity Infrastructure: Threats and Vulnerabilities in Small and Medium Enterprises." ENISA Threat Landscape Series, 3(1), 12–29.

[6] Gartner. (2022). "Top Strategic Technology Trends for 2022: Cybersecurity Mesh Architecture." Gartner Research Press.

[7] Keim, D. A., Mansmann, F., & Stoffel, A. (2010). "Visual Analytics for Cyber Security: Transforming Raw Network Packets into Situational Awareness." Computers & Graphics, 34(4), 331–343.

[8] Kindervag, J., Rose, S., & Borchert, O. (2021). "Zero Trust Architecture Alignment via Distributed Security Control Fabrics." NIST Special Publication, 800-207.

[9] Kruse, C. S., Smith, B., & Vanderlinden, H. (2017). "Security Techniques for the Protection of Electronic Health Records: A Systematic Review." Journal of Medical Systems, 41(8), 127–141.

[10] Kurniawan, A., Sudyana, I. M., & Putra, K. (2019). "Evaluating Single-Board Computing Node Performance for Internal Network Traffic Mirroring and Packet Inspection." International Journal of Computer Networks, 11(3), 89–98.

[11] Patcha, A., & Park, J. M. (2007). "An Overview of Anomaly Detection Techniques: Existing Solutions and Surface Boundaries." Computer Networks, 51(12), 3447–3474.

[12] Radanliev, P., De Roure, D., & Nicolescu, R. (2020). "Distributed Cybersecurity Architecture for Edge Computing in Medical IoT Ecosystems." Journal of Cloud Computing, 9(1), 22–35.

[13] Rose, S., Borchert, O., & Mitchell, S. (2020). "Zero Trust Architecture and Asset-Centric Decentralized Network Perimeter Defenses." IEEE Security & Privacy, 18(4), 44–53.

[14]Satyanarayanan, M. (2017). "The Emergence of Edge Computing and Localized Traffic Summary Mechanics." IEEE Computer, 50(1), 30–39.

[15] Sfar, A. R., Natalizio, E., & Challal, Y. (2018). "A Roadmap for Security in Internet of Things: Challenge Frameworks and Distributed Options." Security and Communication Networks, 2018, 1–18.

[16] Shi, W., Cao, J., & Zhang, Q. (2016). "Edge Computing: Vision, Challenges, and Computational Constraints." IEEE Internet of Things Journal, 3(5), 637–646.

[17] Yaseen, Q., Al-Ahmad, B., & Jararweh, Y. (2021). "Lightweight Rule-Based Intrusion De